

Intel[®] Itanium[®] 2: The Key to Choice in IT Infrastructure

An Insight 64 White Paper Sponsored by the Itanium[®] Solutions Alliance

Executive Summary

Over the past decade buyers of low end and mid range servers voted with their pocketbooks for servers based on x86 processors; these systems now account for over 90 percent of all servers priced below \$50,000. This shift from proprietary to industry-standard architectures greatly simplified the lives of end users, system suppliers and independent software vendors. End users could acquire software that ran without change on systems from dozens of suppliers. System OEMs no longer had to develop their own processors and support proprietary versions of UNIX that ran on their unique CPU architectures. ISVs could focus their development efforts on just one version of their software, confident in the knowledge that version would allow them to address 90 percent of the available market for their products.

While low end and mid-range markets moved to industry standards, proprietary RISC-based architectures continued to dominate sales of servers priced upwards of \$50,000. X86-based systems lacked the scalability and robustness demanded by users of high end systems deployed in mission-critical applications. Intel[®] designed its Itanium[®] 2 processors to support a variety of industry-standard software environments while meeting high end system requirements for scalability and reliability. Just as Intel[®] Xeon[®] processors gave low-end and mid range system buyers the ability to mix-and-match hardware and software from a variety of suppliers, Intel Itanium 2 processors provide new alternatives for buyers of high end servers and free them from the constraints of proprietary RISC architectures.

The recent introduction of the Dual-Core Intel Itanium 2 9000 series processors, along with the maturation of the Intel Itanium 2 processor ecosystem, signal that systems based on Intel Itanium 2 processors are ready for prime time. This document explores the state of the Intel Itanium 2 ecosystem and examines how these processors have changed the way system suppliers design and build high-end servers in the twenty-first century.

Prolog

Sixty years ago a group of engineers at the University of Pennsylvania fired up the ENIAC,¹ regarded as the first electronic computer system. The event marked the birth of the information technology industry that today impacts every aspect of modern life in mature economies around the world, and generates sales of over one trillion dollars on a worldwide basis. Early system suppliers like Univac and IBM had to design and manufacture all the major elements – processors, memory, peripherals, programming tools, operating systems, and applications – themselves. The end-users who purchased those systems were faced with all-or-nothing choices, and had to live with their selected supplier's hardware, peripherals, software, and support.

The IT ecosystem that today provides the hardware and software stacks used in contemporary systems evolved over many decades. A variety of *de facto* and *de jure* industry standards – Unix, Linux, Windows, NFS, SQL, x86, PCI, PCI-Express, Java, and TCP/IP, to name a few – emerged along the way. These standards allowed developers to focus on selected elements within the stack, leveraging R&D investments and enabling dramatic economies of scale. Industry standards also gave end-users the ability to mix and match hardware and software components from multiple suppliers in order to create IT environments that better matched their needs. IT users could move workloads from one supplier's systems to another's, with little if any need for software revisions or retraining.

Intel's x86 architecture played a pivotal role in facilitating the industry's shift from proprietary to open environments. The rapid adoption of MS-DOS and Windows-based systems in the 1980's and 1990's led customers to expect that their preferred PC software applications would run on systems with Intel inside, regardless of the brand on the system itself. The increasing performance of Intel's chips allowed customers to move ever more demanding applications onto industry-standard systems, and they rapidly displaced more expensive workstations and servers based on proprietary RISC processors. This process began with workstations and workgroup servers for file and printer sharing, progressed into the datacenter where at first they handled basic web services and finally tackled application and database tasks in mission-critical environments. As with desktops, the ability to mix and match operating software, middleware and applications from thousands of software vendors, and run that software on systems, large and small, from thousands of hardware suppliers, gave IT users a degree of flexibility and a range of choices unknown in the era of proprietary servers. x86-based servers, some with 64-bit extended architectures, have become so well accepted that they now account for over 90 percent of the servers sold on a worldwide basis.

Although x86-based servers revolutionized the world of high volume servers, the machines that comprise the remaining ten percent of server unit shipments (but account for 46 percent of server system sales) have remained largely the proprietary fiefdoms of a handful of established system suppliers. Customers deploying these larger, more complex systems in their IT infrastructures (a group that includes virtually all large enterprises) find the costs and risks associated with moving applications to incompatible systems usually outweigh the costs of remaining with the original supplier. Until just a few years ago, there was no way to avoid this dilemma; even the most powerful x86 processors on the market lacked the performance, scalability and flexibility to power the large systems that handle the most demanding applications. Multi-core design techniques have boosted x86 performance and 64-bit extensions have improved their ability to handle larger memories, but today's x86 processors still lack the reliability and availability features system designers need to build the mainframe systems that populate the datacenters of large enterprises.

¹ ENIAC was an acronym for Electronic Numerical Integrator and Computer

Fortunately, five years ago Intel delivered the first members of the Intel® Itanium® processor family, a chip architecture designed specifically for the requirements of mainframe system suppliers. Few companies anywhere have the financial resources, organizational skills and industry clout to construct an ecosystem around an entirely new architecture. Intel had the resources and persevered as it constructed a vibrant infrastructure. At the July 2006 launch of the Dual-Core Intel Itanium 2 processor 9000 series, eight leading suppliers of high end servers² displayed a few of the systems they designed around the latest Intel Itanium 2 processor family, and dozens of ISVs demonstrated a sampling of the more than 8000 operating system, middleware and applications that now run on Itanium® 2-based platforms.

Although the bits and bytes that form the basis of Intel Itanium 2 architecture obviously differ in many regards from those that underlie x86 processors, the business model for Intel's Itanium 2 processor builds on the lessons learned during the x86 architecture's rise to ubiquity. Specifically:

1. Intel develops and markets a family of general purpose processors.
2. OEMs integrate those processors into their systems, thus facilitating software compatibility across hardware system vendor lines.
3. ISVs see a far larger market for their software (the collective shipments of all OEMs marketing compatible products) compared with the more balkanized markets inherent in the world of RISC-based architectures.
4. End-users can separate their hardware and software acquisition strategies, since most OEM systems can support standard versions of most ISV software products. Users can combine a heterogeneous hardware vendor environment with software homogeneity, or they can stick with a single hardware vendor while retaining the ability to mix and match software operating environments.
5. The collective growth of the installed bases of all the OEMs adopting the architecture makes that architecture increasingly attractive to ISVs and enhances its value to end-users. (The flip side of this phenomenon is that the stagnant or shrinking bases of systems based on RISC-based architectures become increasingly unattractive to ISVs.)
6. Go to 1. (Repeat)

It has taken far longer than Intel or its partners ever imagined to attain critical mass with regard to the Intel Itanium 2 ecosystem. The launch of the Dual-Core Intel Itanium 2 processor 9000 series (formerly codenamed Montecito) marks the end of the beginning of this process. Intel's Itanium 2 has garnered the strong support of eight out of the top ten server suppliers, along with hundreds of software providers dealing with all elements of the software stack. In the remainder of this document we explore the state of Intel's Itanium 2 processor ecosystem and examine how these processors have changed the way system suppliers design and build high-end servers in the twenty-first century.

² Bull, Fujitsu, Fujitsu Siemens Computers, Hitachi, HP, NEC, SGI and Unisys

A Wide Array of Robust, Industry-Standard System Designs

Itanium 2-based systems come in shapes and sizes to fit any serious computing requirement. Looking for small? How about a 1U dual-processor rack-mount system? The low end configurations based on the latest Dual-Core Itanium 2 Series 9000 processors incorporate four processor cores, each with two hardware threads. Need something a little more powerful? Few systems (based on any processor architecture) can match the multi-rack 60 terabyte, 256 socket Itanium 2-based systems one Itanium-based system OEM developed for the most demanding applications.

Unlike the x86 server market, where system suppliers often differentiate their systems via peripheral content or add-on RAS features like hot-swappable fans and power supplies, manufacturers of Itanium-based systems apply their deep understanding of system architecture to the system design task. Many of the engineers who design these Itanium-based systems formerly designed proprietary platforms built around proprietary processors, proprietary chipsets and proprietary I/O controllers. Now that Intel provides the CPUs, the OEM's engineers can focus on system design issues that impact the cost, capability and competitiveness of their company's products. Can they build a faster crossbar switch that improves memory latency or I/O performance? How should they connect processor modules in large systems? What's the best way to connect blades so they can adapt to varying workloads?

The use of a common processor architecture allows designers to innovate in many different dimensions, with the confidence that whatever scheme they adopt, novel or traditional, there will be a vast array of software applications ready to run on the day they ship their first systems. Table 1 provides a snapshot of the major system offerings from key Itanium 2-based system OEMs. Some sell their systems by the rack, others by the rack-mounted chassis. A few offer bladed configurations. Many have developed Itanium-based systems with an eye toward interoperability with or migration from earlier proprietary systems. The breadth of these products illustrates better than any marketing brochure ever could the intellectual and financial resources manufacturers have invested in Itanium-based products. They're not just casually interested in these products; they're committed to them, and in some cases have bet the future of their companies on them.

We could write a book about the range of designs OEMs have pursued, but in the interests of brevity, we'll only note a few of the distinctive features these suppliers have built into their Itanium 2-based systems.

Bull's NovaScale line includes configurations as small as 2U/4 CPU systems that target dense high performance computing (HPC) applications and as large as a 40U/32 CPU arrangements aimed at database and data mining tasks. Bull developed its own "scalability switch" to complement Intel's 8870 chipset. This combination enables modular expandability from 2 to 32 CPU systems. Bull ported its proprietary GCOS 7 OS to Intel Itanium 2 architecture, and created a software emulator that allows customers to run binary software from its earlier proprietary GCOS line without recompilation. The resulting NovaScale 7000 series can mix partitions running GCOS, Linux, and Windows Server 2003 in a single rack, thus simplifying the server consolidation problem for Bull's customers.

Fujitsu's PRIMEQUEST line comes in 8, 16 and 32-processor versions; physical memory scales to 2,048 GB (2 terabytes). Fujitsu's proprietary chipset lets it link up to eight 4P system boards to a non-blocking crossbar switch with an aggregate throughput of 136.4 GB/second. Fujitsu's "System Mirror" assures application availability by preventing failures in the crossbar switch and memory from impacting system operation. Its "Flexible I/O" architecture lets it dynamically reallocate I/O resources to specific hardware partitions, based on system requirements. Fujitsu sells its PRIMEQUEST systems in Asia and the US, but its partner **Fujitsu Siemens Computers** markets these systems in Europe. As the European high-end systems market consolidated over the past

decade, Siemens acquired the architectures and customers of RCA, Nixdorf and Wang. Now Fujitsu Siemens Computers plans to migrate BS2000 (the surviving Siemens OS) to Intel's Itanium 2 architecture, giving its customers a straightforward upgrade path into the open systems world.

Hitachi used its technological expertise to craft **BladeSymphony**, a bladed implementation of a contemporary mainframe. Hitachi starts with a basic atomic unit of a dual processor Itanium 2-based blade that contains up to 32 GB of memory; its blade chassis accommodates up to eight blades. Most blade systems simply rely on some sort of cluster arrangement (i.e. gigabit Ethernet or InfiniBand) to connect blades, but Hitachi goes further. Its proprietary Node Controller (the chipset on each blade that ties the blade to the chassis) can link the blades in each chassis into SMP configurations with 2, 4, 6 or 8 CPUs per SMP. A single chassis can operate as eight two-way blades or two eight-way blades. BladeSymphony lets end-users deploy a common hardware infrastructure that can be tailored for scale-up applications like database management or application processing, or scale-out tasks like web services. BladeSymphony also includes two-way and four-way Xeon blades, but only the Intel Itanium 2 processor blades can pull off the cool SMP reconfiguration trick.

HP offers the broadest line of Itanium 2-based systems, a not-surprising fact given HP's position as the co-developer of the Intel Itanium architecture and designer of some early Intel Itanium processors. HP realized earlier than most that the future belonged to systems based on industry standards, and joined with Intel to insure that powerful and scalable industry standard processors would be available when it designed its next generation of high-end systems. A lot of water has passed under HP's bridge since the 1994 announcement of its plan to partner with Intel on a line of 64-bit processors. The industry landscape has changed in many ways, but HP has remained true to its original plan, and now finds itself in the final stages of the migration of its HP 9000, **AlphaServer** and Non-Stop Kernel migrations to Intel's Itanium 2 architecture. HP could serve as the poster child for the shift to industry-standard based servers. It ported the HP-UX operating system to Intel Itanium 2 architecture, and provided software emulators that allowed users to run unmodified PA-RISC binaries on Itanium-based systems.³ It ported the OpenVMS OS it acquired via its Compaq acquisition to run on Intel Itanium 2 architecture as well. HP designed its last PA-RISC CPUs to be compatible with Intel's Itanium 2 processors from electrical, mechanical and thermal perspectives, so it could leverage its high-end sx1000/sx2000 chipset designs across both architectures. HP's latest Integrity SuperDome configurations can simultaneously run HP-UX, Linux, Windows and OpenVMS in different system partitions. As a result, HP is well positioned to offer its customers some of the most flexible systems in the industry when it comes to server consolidation.

NEC's Itanium 2-based system line reflects the decades of high-end system experience the company has amassed, along with almost a decade of experience with Itanium-based systems. The company's 16-way AzusA platform, the largest to be built around the first generation Intel Itanium processor (Merced), served as a development platform for many development projects at major ISVs in the early days of Intel's Itanium program. NEC's latest **Express 5800/1000** series incorporates the kinds of features only a company with years of mainframe expertise would appreciate. Fully configured systems include eight cell boards (each with four dual-core Intel Itanium 2 processors and 128 GB memory) and eight I/O Router modules (each with 8 PCI-X slots for AsAmA2). All these modules connect to a pair of redundant crossbar switches with an aggregate throughput of over 280 GB/second. An autonomic service processor monitors system operation and isolates failing components before they can impact other parts of the system. NEC's

³ Of course, the efficiency of software emulation varies greatly by application. The technique has proven quite useful for many system management tasks, but is less helpful in computationally intensive areas.

systems run Linux, HP-UX, Windows and ACOS, the OS that runs on NEC's line of mainframes based on its proprietary architecture.

SGI recently recast the **Altix** series and moved to a blade orientation at the same time it incorporated the latest Dual-Core Intel Itanium 2 processor 9000 series. The company's new Altix 4700 system accommodates up to 512 10U blades with anywhere from 512 to 1024 Intel Itanium 2 processors. A system this large fills at least 16 racks, even more if you want to load the system with the maximum 60 TB of main memory it's designed to handle. Unlike other Intel Itanium 2 blade implementations, SGI separates functions across blades. Processor blades contain *only* processors and the NUMALink controllers that tie them to the NUMAflex fabric. Memory blades contain *only* memory -- up to 48 GB per blade. I/O blades support a variety of PCI-X, PCI-Express and graphics options. SGI argues that its approach does not force customers who need vast amounts of physical memory to buy extra "cells" that contain processors they might not need, and vice versa. Customers able to get by with *only* 38 Itanium 2 processors and 456 GB of RAM can acquire SGI's technology via its cost-optimized Altix 450 system. Same blades, smaller racks.

Unisys, a company with roots dating back to the ENIAC discussed earlier in this document, was an early adopter of Intel Itanium 2 technology. Its **ES7000/one** server allows users to combine up to eight system cells, each with four Intel Itanium 2 or Xeon processors and 64 GB of main memory for a total of 32 processors and half a terabyte of memory. It is the only system on the market today capable of running Intel Xeon and Itanium 2 processors in the same server. This flexibility allows customers to integrate their enterprises through standardization and maximize the return on investment by investing in what they need when they need it. Each cell contains the Unisys Cellular Multi Processor (CMP) crossbar, which connects memory, processors and I/O within the cell and facilitates communications between the cells via the Flexbar. Given the company's heritage, the Unisys customer base includes users who depend on modern versions of the OS2200 environment Univac introduced in 1964, and the MCP (Master Control Program) Burroughs introduced in 1962. The company recently announced its intent to migrate its OS2200 and MCP software to Intel technology, in order to give its customers a more flexible environment. After Unisys completes this development activity, its customers will have the ability to run Linux, Windows, OS2200 and MCP applications on a single system.

Table 1: Snapshot of OEM Itanium® 2-based System Offerings as of September 2006

Supplier	Model	Package	Form Factor	Maximum Modules	Sockets/ Module	Maximum Sockets	Maximum Memory	Memory Arch	Chipset	Interconnect Technology
Bull	NovaScale 5245	40U Chassis	Board	6	4	24	384 GB	NUMA	8870	Bull
Bull	NovaScale 5165	19U Chassis	Board	4	4	16	256 GB	NUMA	8870	Scalability
Bull	NovaScale 5085	8U Chassis	Board	2	4	8	128 GB	NUMA	8870	Switch
Bull	NovaScale 9320/7320	Rack	Board	8	4	32	512 GB	NUMA	NUMA	"
Bull	NovaScale 9160/7160	Rack	Board	4	4	16	256 GB	NUMA	NUMA	"
Bull	NovaScale 3045C	2U Chassis	Board	1	4	4	128 GB	Flat	Proprietary	--
Bull	NovaScale 3045	4U Chassis	Board	1	4	4	128 GB	Flat	Proprietary	--
Bull	NovaScale 3025	4U Chassis	Board	1	2	2	64 GB	Flat	Proprietary	--
Fujitsu / Fuji. Siemens	PRIMEQUEST 580	Rack	Board	8	4	32	2048 GB	Flat	Fujitsu-developed	X bar Switch
Fujitsu / Fuji. Siemens	PRIMEQUEST 540	Rack	Board	4	4	16	1024 GB	Flat	Fujitsu-developed	X bar Switch
Fujitsu / Fuji. Siemens	PRIMEQUEST 520	12U Chassis	Board	2	4	8	256 GB	Flat	Fujitsu-developed	X bar Switch
Hitachi	BladeSymphony	10U Blade	Blade	8	2	16	128 GB	Flat	Node Controller	8-way SMP
HP	Superdome	Rack	Cell	16	4	64	2048 GB	NUMA	sx2000	X bar Switch
HP	Integrity rx8640	17U Chassis	Cell	4	8	16	256GB	NUMA	sx2000	X bar Switch
HP	Integrity rx7640	10U Chassis	Cell	2	4	8	128GB	NUMA	sx2000	X bar Switch
HP	Integrity rx6600	7U Chassis	Cell	1	4	4	192 GB	Flat	sx1000	--
HP	Integrity rx3600	4U Chassis	Board	1	2	2	96 GB	Flat	zx2	--
HP	Integrity rx2620	2U Chassis	Board	1	2	2	24 GB	Flat	ZX1	--
HP	Integrity rx1620	1U Chassis	Board	1	2	2	16 GB	Flat	ZX1	--
HP	Integrity BL60p	6U Blade	Blade	1	2	2	8 GB	Flat	ZX1	Backplane
NEC	Express 5800/1320Xf	Rack	Cell	8	4	32	1024 GB	NUMA	A^3(Acube)	X Bar Switch
NEC	Express 5800/1080Rf	8U	Cell	2	4	8	128 GB	NUMA	A^3(Acube)	Backplane
SGI	Altix 4700	10U Blade	Blade	1	1	512	60 TB	NUMA	NUMAlink	NUMAflex
SGI	Altix 450	10U Blade	Blade	1	2	38	456 GB	NUMA	NUMAlink	NUMAflex
Unisys	ES7000/one	3U Chassis	Cell	8	4	32	512 GB	Directory	CMP	CMP FleXbar

Source: Vendor Websites

Industry-Standard Hardware Means ISVs Never Have to Say they're Sorry

In days of yore, suppliers who based their systems on proprietary processors needed to deliver and support proprietary operating system(s) that supported these processors. Some environments – like IBM's I5OS or Z/OS – remain truly unique and incompatible with any other operating system. Others, like AIX and Solaris, were variations on a common UNIX theme, but still required human intervention whenever an organization needed to move from one vendor environment to another. Although most contemporary applications are written in high level languages like C, C++ or Java, a migration project involving a small number of applications can involve hundreds of software modules and require man-years of effort to convert and test.

Independent software developers (ISVs) expend the lion's share of their R&D budgets supporting the migration of their software products from one environment to another. They flock to the high volume environments where the largest number of potential customers reside – Windows or Linux on x86 systems, or Solaris on SPARC or x86 – but the obvious choices drop off rapidly from there. Successful ISVs shun RISC-based environments (hardware or software) with small market shares because the opportunities look so more promising with architectures that have already acquired a major share of their segments. This fact places system suppliers with small market shares at a decided disadvantage from those that appear more attractive to ISVs.

X86-based systems have long benefited from this phenomenon. Given the vast number of systems compatible with the x86 standard, it's an obvious target for ISVs. This never-ending supply of new software increases the value of the x86 platform, and makes other platforms less useful to end-users. Those who purchase x86 systems can obtain ready-to-run software already ported and tested on the platform. X86 system suppliers may choose to bundle copies of Windows or Linux with the systems they market, primarily to simplify their customer's acquisition efforts, but software availability rarely becomes an issue in the customer's buying decision.

The annual shipment volume of systems based on the Intel Itanium 2 architecture has grown rapidly, and will soon surpass that of any of the traditional RISC architectures. This result stems not from the success of any single Itanium-based system supplier, but rather from the collective shipments of all these suppliers. Although we would not expect Itanium-based system shipments will ever surpass the aggregated annual volume of x86-based system shipments, it may very well come in second in this horse race, thus making it an attractive target for ISVs on an ongoing basis. Most software vendors already treat the Itanium-based versions of their packages the same way they treat those adapted for x86-based systems. Their software packages are available off-the-shelf, ready to run on Itanium 2-based hardware.

Most ISV organizations are far too savvy to suggest publicly that they prefer one hardware architecture over another, but many note privately that they often achieve their best raw performance, or their best price/performance on Itanium 2-based platforms. They also appreciate the economic benefits they obtain when they can amortize the expense of their software migration efforts over sales to customers of many different vendors, instead of needing to expend effort on each individual vendor's environment.

A Win-Win-Win Situation

We have discussed how the move to Intel Itanium 2 processors allows hardware system suppliers to innovate in areas of more direct concern to their current and prospective customers, and how it also allows ISVs to leverage their R&D efforts, so that they can spend more time adding features their customers ask for, instead of merely porting the existing package to yet another platform. The advantages that accrue to the Intel Itanium 2 ecosystem partners would be worthless if the end-users who purchase these systems, and depend on them for their most critical computing tasks, did not benefit

directly as well. Fortunately, end-users win the most from the shift to systems based on industry-standard Intel Itanium 2 processors.

First and foremost, Intel Itanium 2 processors bring to the high end of the server market the same practical advantages that x86 architecture brought to low end and mid range servers. Customers can select hardware vendors without any concern about the software available for those vendors' systems. They can even test novel hardware arrangements (based on industry-standard CPU architectures) with their existing software applications, and thus determine quickly whether the new arrangement improves their overall capabilities. Alternately, they can evaluate most new software packages on hardware systems already in place within their organizations.

As customers build out their high-end architectures with Itanium-based systems, they can be confident that systems and software they acquire in years to come will interoperate with the systems they obtain today, in much the same way as today's x86-based systems have incorporated backward compatibility with earlier systems, independent of the supplier of those systems. That same binary compatibility across system vendor lines will facilitate workload consolidation and scheduling, since tasks can be moved from one system to another, based solely on system loading and configuration, without regard to the specific system label on the boxes in the datacenter.

Last but not least, customers using high end systems based on industry-standard architectures gain a new and useful tool to optimize their IT hardware purchasing decisions. If their incumbent system supplier fails to meet their needs, they can move to a different supplier with minimal disruption to existing software environments. Such a shift in purchasing arrangements can be disruptive to both the buyer and seller, but it becomes especially difficult, expensive and risky when the shift requires a transition in software environments as well as commercial terms and conditions.

Let's recap what we have discussed in this paper:

- Until recently, buyers of high end servers had no choice but to use systems based on RISC architectures – existing industry-standard processors just were not up to the task.
- The advent of Intel Itanium 2 processors and the maturation of the Intel Itanium 2 processor ecosystem gives buyers of high end servers a new choice, analogous to the choice x86-based processors gave buyers of low end and mid range servers ten years ago.
- Itanium 2-based systems from many OEMs, large and small, give end users the ability to make hardware decisions independently of software, and vice versa.
- End-user environments based on common processor architectures available in systems from many independent suppliers allow a wider selection of software, and facilitate easier workload consolidation and/or scheduling.

The history of the computer industry clearly demonstrates that proprietary approaches invariably yield to those based on widely accepted *de facto* and/or *de jure* industry standards. No one company can innovate as successfully and consistently as a collection of companies working toward a common goal. Industry standards allow software developers to focus their energies on a small number of platforms, leveraging their efforts and increasing their financial returns. Regardless of any short term intermediate swings in performance, price/performance or performance/watt metrics, the end point is clear. Those who recognize and act on this trend it will ultimately prevail over those that deny its inevitability.