WHITE PAPER

PRIMEQUEST Servers: Mission-Critical Servers for Linux and Windows

March 2006

IDEAS Custom Consulting Services

PREPARED FOR

Fujitsu

TABLE OF CONTENTS

Executive Summary

A year ago, Fujitsu expanded its portfolio of robust enterprise servers with the introduction of the PRIMEQUEST family. Fujitsu had observed that a growing number of customers were attracted to Linux and Windows operating environments and to servers employing Intel processors. The Itanium 2-powered PRIMEQUEST line addresses that customer desire with high performance servers designed for mission-critical Linux and Windows environments.

With the addition of an eight-processor model 420, Fujitsu's PRIMEQUEST family now offers Itanium-based servers that span the range from midrange to high end. The PRIMEQUEST models complement Fujitsu's SPARC/Solaris PRIMEPOWER systems in mission-critical environments. Currently available with Madison 9M chips, the midrange PRIMEQUEST 420 joins the PRIMEQUEST 440 with 16 processors and the PRIMEQUEST 480, which scales up to 32 processors.

Employing standard Intel processors certainly should not imply that PRIMEQUEST is merely another Itanium-based platform. Quite to the contrary, Fujitsu has leveraged the expertise of the engineers who designed its mainframe and RISC platforms to create servers that are unique in their assurance of mission-critical operation. This paper highlights the innovative design that sets PRIMEQUEST apart from its competitors.

Introduction

With a long history of developing mainframe and enterprise-class UNIX servers, Fujitsu is respected as a leader in supplying customers with robust systems to satisfy their business-critical computing needs.

Fujitsu continues to support its loyal customers with enhanced PRIMEPOWER systems and remains committed to SPARC64. In fact, Sun Microsystems has turned to Fujitsu to develop a jointly marketed future family of SPARC-based Solaris servers.

In addition, Fujitsu recognizes the strong market demand for high-volume industry-standard solutions, especially since the introduction of 64-bit addressing extensions. PRIMERGY systems run Windows and Linux on servers powered by Intel Xeon and AMD Opteron processors. The PRIMERGY line spans a range from innovative blade solutions to scalable midrange systems.

INTEL REMAINS COMMITTED TO ITANIUM

Some vendors, who have not embraced Itanium 2, attempt to paint Itanium as a niche product. They assert that since Itanium shipments will probably not match the high volumes of Xeon, that Intel has in lost interest in enhancing its high-end processor.

Clearly, that is not true. Intel has emphatically reaffirmed that it will continue to enhance both its high-volume and enterprise chip families. Intel recognizes that different chip designs are needed to address the different requirements of enterprise servers and commodity-like industry-standard servers. Admittedly, smaller industry-standard servers will ship in much higher volumes, but the enterprise servers carry higher prices and have the potential to be a very lucrative market segment for Intel.

Intel is an enthusiastic supporter of the industry consortium "Itanium Solutions Alliance" and has committed billions of dollars of continued funding for Itanium 2 and its follow-ons. Fujitsu is also a founding member of the Itanium Solutions Alliance and is committed to demonstrating that Itanium is ideal for business-critical computing-intensive environments.

In particular, Fujitsu had noticed that many customers were seeking to deploy enterprise applications on Linux. The fact that open source Linux code is thoroughly examined gives those customers the comfort that Linux can provide a secure environment. For example, governmental organizations have embraced the open, collaborative environment represented by Linux as an alternative to potential vendor lock-in. Furthermore, because Linux and Windows run on multiple platforms, there is a large developer community for those open and broadly deployed operating environments, thus assuring an extensive application portfolio.

Fujitsu's PRIMEQUEST extends the company's industry-standard solutions to mission-critical enterprise environments. This Itanium 2-based family leverages Fujitsu's expertise in designing highly available, powerful servers. PRIMEQUEST gives customers the opportunity to run enterprise applications on Linux or Windows to complement their use of industry-standard entry and midrange models.

In April 2005, Fujitsu announced its PRIMEQUEST family, starting with the high-end PRIMEQUEST 440 and 480 models. Now, the new PRIMEQUEST 420 extends Fujitsu's mission-critical Itanium servers into the midrange. Able to configure up to eight processor chips, the initial PRIMEQUEST 420 models support from one to eight Itanium Madison processors, providing customers with a highly available midrange server. The PRIMEQUEST product line, consisting of 420, 440, and 480 models, offers scalability up to 32 Itanium processor chips. Fujitsu has indicated that the PRIMEQUEST family will incorporate dual-core Montecito chips when they become available, thus extending the processing capacity up to 64 processor cores.

PRIMEQUEST models 440 and 480 have been shipping since the summer of 2005. Fujitsu customers who have already deployed PRIMEQUEST report they are pleased with its capabilities. However, there remain many potential customers who are not familiar with PRIMEQUEST's capabilities.

Certainly, Fujitsu computing systems have a strong reputation worldwide. Fujitsu has supplied computers in Japan for over 50 years, and remains the leading provider of computer systems in that country. The Fujitsu-Siemens partnership has a strong presence in much of Europe. In North America, Fujitsu once played a significant role as a mainframe supplier. However, as mainframes evolved into specialized roles, customers turned their focus to suppliers of high-volume platforms. Now, customers once again seek highly reliable systems. PRIMEQUEST addresses that need for business-critical computing. This paper endeavors to highlight the advantages of Fujitsu's family of Itanium servers to those customers who may not be familiar with PRIMEQUEST.

Industry Standard, but with a Unique Approach

From the customer perspective, choosing industry-standard platforms provides the freedom to consider multiple vendors' offerings. At the same time, vendors face the challenge of differentiating their systems while adhering to industry standards.

For Fujitsu, PRIMEQUEST's differentiation focuses on providing a robust server suitable for mission-critical environments. Certainly, high-end servers must be able to deliver the performance promised by high-end processors. Leveraging the

engineers who created its mainframes, supercomputers, and 128-way SPARC SMPs, Fujitsu incorporated low latency, high-bandwidth system buses into PRIMEQUEST. Beyond designing for performance, the engineers designed for high availability through redundancy. Fujitsu-designed chipsets enable address buses, data buses, and even memory to be configured as mirrored.

The PRIMEQUEST Mirror Mode architecture offers two options – Standard Mirror Mode and Extended Mirror Mode – each of which configure redundant hardware to operate in lock step. As a result, failure of any of these components does not cause a PRIMEQUEST partition to fail. Standard Mirror Mode provides redundant system buses. Extended Mirror Mode even duplicates memory to protect against multiple memory-bit errors or failure of an entire memory module.

Hot-plug capabilities also contribute to a PRIMEQUEST server's availability. Many PRIMEQUEST components are, by design, hot pluggable. These components include most I/O devices; fans and fan trays; power supplies; the System Management Board; the operator panel; and the keyboard/video/mouse (KVM) unit, among others.

And, of course, for further high availability, PRIMEQUEST servers may be clustered. In addition to the traditional clustering of two or more physically separate servers, PRIMEQUEST models 440 and 480 support clustered partitions within the same server.

The high availability features of PRIMEQUEST set it apart from other servers that compete in the enterprise market. This paper highlights some of the advantages that PRIMEQUEST holds over these competing servers.[1]

PRIMEQUEST Family Overview

Fujitsu's PRIMEQUEST family consists of three models of scalable, shared-memory, symmetric multiprocessors (SMPs). The systems are comprised of modular building blocks: system boards (SBs), I/O units (IOU), crossbar interconnect fabrics, and System Management Modules. Each system board, contains four processor chips and 32 memory DIMM slots. The system boards connect into the crossbar switches to allow all memory and I/O to be accessed by any processor. Figure 1 (next page) depicts a block diagram of the PRIMEQUEST family. The largest member of the family, PRIMEQUEST 480, can contain all eight system boards (as illustrated in the figure), scaling to 32 Itanium 2 processor chips. PRIMEQUEST 440 comprises four systems boards and four IOUs, creating up to a 16-processor system. The recently announced model 420 consists of up to two system boards and two IOUs, allowing midrange systems to be configured with one to eight Itanium 2 processor chips.

[1] This document was prepared during March 2006 and covers PRIMEQUEST and competitor models that were available at that time. Over time, as new processor chips, including Montecito, become available, those upgraded chips will be incorporated into new models. Although the details of the upgraded models may differ from those discussed in this document, each vendor's overall product family will likely remain consistent with what is discussed in this document.

FIGURE 1

Block Diagram of PRIMEQUEST Architecture

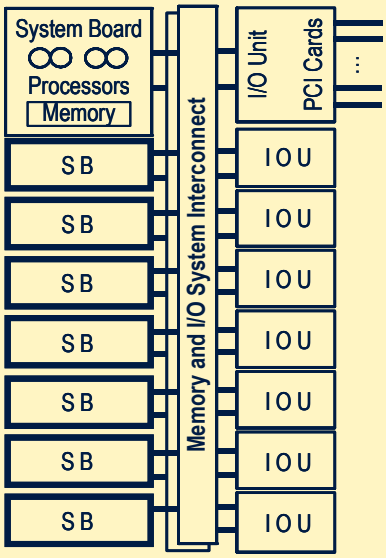

At first glance, Figure 1 seems similar to high-level block diagrams of other modularly scalable SMPs. What is not apparent from the block diagram is that Fujitsu has leveraged its supercomputing and mainframe expertise to create a crossbar switch that delivers higher bandwidth at lower latencies than competitive systems. For example, a 32-processor PRIMEQUEST 480 has 75 GB/s peak I/O bandwidth compared to the 48 GB/s peak I/O bandwidth of a 32-processor-core IBM p5-590. The interconnect chipset of an eight-processor PRIMEQUEST 420 delivers 25.6 GB/s crossbar bandwidth while HP's sx1000 chipset only provides 8 GB/s bandwidth for an eight-processor HP rx7620. Of course, customers do not really need to understand all the chipset specifications. What they do appreciate is that Fujitsu's experienced engineering team has designed PRIMEQUEST to deliver higher system performance than other modular implementations that may use similar processor chips.

Note from the figure that the system boards and I/O units are physically separate from each other. As will be covered later, when splitting the system into partitions, system boards and I/O units can be independently assigned to different partitions, a flexibility not offered by some competitors.

Careful observation of Figure 1 implies that the PRIMEQUEST system interconnect is actually implemented as multiple crossbar fabrics. Indeed, the multiple interconnect fabrics is an important reason why customers can depend on PRIMEQUEST systems for highly available, mission-critical computing.

The interconnect block of Figure 1 actually consists of multiple address crossbar fabrics and data crossbar fabrics. In Fujitsu's Standard Mirror Mode, the global address crossbars are run in a lock-step duplicated manner.

To put this concept into perspective, the global address crossbar broadcasts the address of requested data among all memory, I/O, and system boards. An undetected error could result in data being written in the wrong location, which could later develop into a serious system failure or corrupted data. PRIMEQUEST actually implements two complete address buses. Each system board has separate chips to interface to each of the address buses. Standard Mirror Mode operation treats these two address buses as duplicate copies of the same address bus transactions. Intermittent errors, or complete failure, of one bus (or the associated chipset) will not impact system operation since the duplicate, mirrored bus has an independent copy of the correct address information.[2]

Other vendors may implement various error detecting mechanisms on their address interconnect. However, Fujitsu has taken the lead as the first to offer full redundancy of the mirrored address bus interconnect as the standard operating mode.

Note that Fujitsu's hardware design also duplicates data buses in addition to address buses. Extended Mirror Mode redundantly configures that data

[2] For environments with exceptional memory and I/O traffic, non-mirrored operation can be selected to permit both address buses to be used to maximize throughput. But for most workloads, the total anticipated load on the two buses is so low, that only using one, with the other as a mirrored copy, degrades overall performance by less than 1%, according to Fujitsu projections. Thus, Standard Mirror Mode is the default mirroring option on PRIMEQUEST servers.

interconnect in a manner similar to address bus mirroring. However, before exploring other aspects of Extended Mirror Mode, it is useful to look at the memory capacity of PRIMEQUEST servers.

PRIMEQUEST Configurations

Before customers can select which platform best satisfies their requirements, they need to narrow down their choices to a "short list" of potential solutions that offer adequate capabilities to meet their computing needs. For enterprise computing, that means finding servers that can configure enough processors, memory, and I/O to handle the application workload. As shown in Table 1, PRIMEQUEST models 420, 440, and 480 offer configurations that span the range from midrange to high end. All offer powerful processors, large amounts of fast DDR2 memory, and abundant PCI slots and internal disk bays to attach necessary I/O.

Table 1. PRIMEQUEST Capacities

PRIMEQUEST Model	Processor GHz	Processors	Max. Memory (DDR2)	Max. PCI Slots (with Expansion)	Internal Disk
PRIMEQUEST 420	Itanium 2 1.5, 1.6 GHz	1-8	256 GB	4 PCI-X 2 PCI-Express (18 PCI-X/E total)	8 drives 584 GB
PRIMEQUEST 440	Itanium 2 1.5, 1.6 GHz	1-16	512 GB	16 PCI-X (64 PCI-X)	16 drives 2.35 TB
PRIMEQUEST 480	Itanium 2 1.5, 1.6 GHz	1-32	1 TB	32 PCI-X (128 PCI-X)	32 drives 4.7 TB

As shown in Table 2, IBM System p5 servers do not match the amount of memory available in similar PRIMEQUEST configurations. For example, IBM's p5-550 and p5-550Q only support 64 GB of DDR2 memory, while PRIMEQUEST offers four times as much. And, the p5-590 only supports its maximum memory with the slower DDR1 modules; DDR2 capacity for the p5-590 is only one eighth as much.

Table 2. IBM System p5 Supported Memory Capacity

IBM System p5 Model	Processor, GHz	Max. No. Chips/Cores	Maximum Memory (DDR2 except where noted)
p5-550	POWER5+: 1.65, 1.9 GHz	2 / 4	64 GB
p5-550Q	POWER5+: 1.5 GHz	4 / 8	64 GB
p5-560Q	POWER5+: 1.5 GHz	8 / 16	128 GB
p5-570	POWER5+: 1.9, 2.2 GHz	8 / 16	512 GB (for 2.2 GHz) 256 GB (for 1.9 GHz)
p5-590	POWER5: 1.65 GHz	16 / 32	1 TB (DDR1) 128 GB (DDR2)

Table 3 indicates that HP Integrity servers also do not match the amount of memory offered in PRIMEQUEST configurations. Note that the Integrity platforms that HP has been shipping (with the sx1000 chipset) only supported slower PC133 SDRAM. HP has recently announced new models (with the sx2000 chipset) that support DDR2, but even those new models cannot match PRIMEQUEST memory capacity. Similarly, PRIMEQUEST configures more memory per processor core compared to

WHAT ABOUT MONTECITO?

The "Madison 9M" processor chips currently used in PRIMEQUEST remain quite competitive with all RISC alternatives. Soon, a dramatically new Itanium implementation will become available – now codenamed "Montecito." Each Montecito chip will contain two processor cores and each core will execute two instruction stream "threads."

PRIMEQUEST models are designed for incremental upgrade and will support systems containing both Madison and Montecito system boards.

When Montecito becomes available the performance of the PRIMEQUEST will be dramatically boosted. Customers installing PRIMEQUEST now can look forward to substantial performance gains while still enjoying the mission-critical high-availability attributes that set PRIMEQUEST apart from other enterprise servers.

the Unisys ES7000 models of Table 4. Note also that the Unisys Itanium platform, ES7000 model 400, only scales to a maximum of 16 Itanium chips in each partition, half the maximum single-system image offered on the PRIMEQUEST 480.

Table 3. HP Integrity Supported Memory Capacity

HP Integrity Model	Processor, GHz	Max. No. Processors	Max. Memory
rx7640 (sx2000)	Itanium 2: 1.6 GHz	8	128 GB (DDR2)
rx8640 (sx2000)	Itanium 2: 1.6 GHz	16	256 GB (DDR2)
Superdome (sx2000)	Itanium 2: 1.6 GHz	64	1 TB (DDR2)
rx7620-16 (sx1000)	Itanium 2: 1.5, 1.6 GHz Itanium 2: 1.1 GHz (mx2)	8 16 (mx2)	128 GB (SDRAM)
rx8620-32 (sx1000)	Itanium 2: 1.5, 1.6 GHz Itanium 2: 1.1 GHz (mx2)	16 32 (mx2)	256 GB (SDRAM)
Superdome (sx1000)	Itanium 2: 1.6 GHz Itanium 2: 1.1 GHz (mx2)	64 128 (mx2)	1 TB (SDRAM)

Table 4. Unisys ES7000 Supported Memory Capacity

Unisys Model	Processor, GHz	Max. No. Chips/Cores	Maximum Memory
ES7000 model 400	Itanium 2: 1.3, 1.5, 1.6 GHz Xeon: 2.8, 3.0 GHz	32 chips total for all partitions Max. 16 chips per partition	256 GB per partition
ES7000 model 600	Xeon EM64T: 2.8, 3.3 GHz	32 / 64	512 GB

So, why did Fujitsu design PRIMEQUEST models to support more memory than its competitors? The additional memory can be used either to satisfy memory-hungry applications or to provide memory mirroring for mission-critical environments.

Many applications have a voracious appetite for memory. Some, such as Business Intelligence applications, search vast amounts of data to discover trends and data relationships. Other applications service large numbers of concurrent users, each of whom has his/her own data needs. Adding memory is often the easiest way to improve performance. Customers should examine their applications to understand how they would benefit from the larger memory capacity of PRIMEQUEST. In addition, they should benchmark PRIMEQUEST against other enterprise servers to demonstrate the benefit brought by DDR2 and Fujitsu's own high-performance interface chips. For workloads that need large amounts of memory per processor core, the PRIMEQUEST models offer configurations unmatched by competition.

Other customer environments may not be constrained by memory capacity but require high availability attributes beyond the typical single-error-correct double-error-detect (SECDED) Error Correcting Code (ECC) or survival of single memory-chip failure (often described as chip kill tolerance). For these mission-critical enterprise environments, Fujitsu's Extended Mirror Mode offers memory and data bus redundancy.

Data Mirroring

As common in the industry, Fujitsu has incorporated ECC and chip kill to help minimize the chance that a memory module failure could crash the system. Nonetheless, there are mission-critical customer environments that demand even higher levels of availability. In addition to the standard address bus mirroring, Fujitsu offers optional mirroring of the data crossbar and even the memory itself.

When configured as mirrored, resources are paired and perform the identical function in parallel. Should one in the pair fail, the other resource will still provide a valid result, allowing the system to continue without interruption. As shown in Table 5, Standard Mirror Mode starts by mirroring the global address crossbars that serve as the address interconnect between all system boards and I/O units. Extended Mirror Mode arranges the four global data crossbars as two duplicate halves and configures memory so that each memory location is duplicated in two separate memory modules.

Table 5. Mirror Mode Capabilities Offered on PRIMEQUEST Servers

	Standard Mirror Mode	Extended Mirror Mode
Address Bus	Mirrored	Mirrored
Data Bus	Independent	Mirrored
System Board Bandwidth	12.8 GB/sec.	6.4 GB/sec.
IOU Bandwidth	6.4 GB/sec.	3.2 GB/sec.
Memory	Independent	Mirrored
Memory Capacity	100%	50%
Performance	>99%	>95%

PRIMEQUEST's data crossbar consists of four separate data paths. Each is protected by single-bit ECC, and has dedicated interface chips on the system and I/O boards. For customers who want an extra degree of protection, Extended Mirror Mode configures the four data buses as two mirrored pairs. When the data buses are mirrored, the system board interface circuitry is also run in a mirrored configuration, thus protecting the system from failures in the interface circuitry as well as on the buses themselves.

More importantly, Extended Mirror Mode duplicates data stored in main memory. As previously noted, memory ECC not only corrects single bit errors but also reconstructs all four bits from a defective memory chip. However, if multiple errors occur, the contents of that memory location would not be recoverable. Memory mirroring duplicates data in two separate sets of memory chips. Multiple bit failures that might corrupt one memory location are highly unlikely to occur simultaneously at the mirrored location. And, mirroring the interface circuitry also avoids system failures caused by circuit failure, thus providing the robust data integrity that may be needed for extreme mission-critical applications.

Of course mirroring memory does reduce the total memory capacity in half. However, with the extensive amount of memory that can be configured on PRIMEQUEST servers, the reduction in total memory capacity should not be a concern. Admittedly, having data stored in duplicate memory locations does double

FIGURE 2

PRIMEQUEST Flexible Partitioning

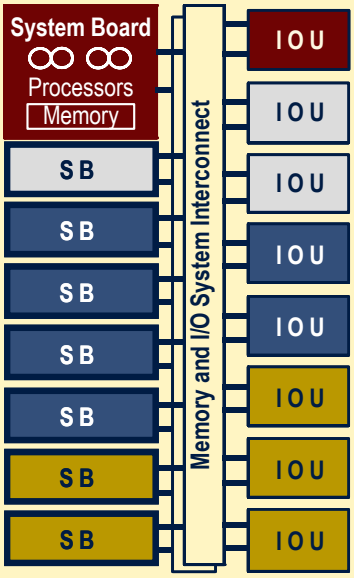

the cost for the memory. Thus, Extended Mirror Mode would be most appropriate for environments that place great importance on avoiding system disruption.

Enterprise servers from other vendors do not offer redundancy features comparable to PRIMEQUEST Extended Mirror Mode. Those who have been accustomed to mainframe levels of high reliability and availability will especially value PRIMEQUEST's mirroring capabilities.

Flexible Partitioning

PRIMEQUEST's scalability to 32 Itanium 2 processor chips may be beyond the day-to-day requirements of many individual applications, particularly those running under Windows or Linux. Nonetheless, customers often procure large servers that can be flexibly configured to meet peak requirements and can also be divided into independent, isolated partitions for normal operations. PRIMEQUEST's partitioning capabilities hold an edge over its primary competitors. And, large systems can be used to consolidate the workloads of multiple servers into an easy-to-manage, central server. For example, different workloads may have conflicting software stacks, such as requiring different versions of middleware. The error isolation afforded by partitioning is also very valuable to assure that failures in software under development or test do not affect production workloads.

Recall that PRIMEQUEST systems are constructed from modular building blocks. Each of the system boards contains up to four processor chips, up to 128 GB memory, and connections to I/O by way of the system interconnect crossbar. Independent partitions can be created within PRIMEQUEST by instructing the crossbar switch to isolate sets of system boards from each other and by assigning I/O units access to those system boards. Each independent partition runs a different instance of an operating system and is electrically isolated from hardware or software failures in other partitions.

Different vendors have different approaches for system partitioning. IBM, for example, focuses on virtual partitioning that depends upon software to keep the partitions isolated. But software partitioning, such as IBM's hypervisor, is susceptible to single points of software failure. And, software overheads may reduce performance compared to running that workload on the native hardware. Although many customers do acknowledge a role for software partitioning, they prefer the electrical isolation of hardware-enforced partitioning to assure them that the partitions are fully isolated.

HP offers hardware partitioning that also divides its large servers along hardware boundaries, called "cell boards" by HP. However, HP's cell board partitioning unit contains not just processors and memory, but also includes the paths to I/O. Thus, each partition has a fixed ratio of I/O to processor/memory.

PRIMEQUEST avoids that fixed relationship by allowing I/O units (IOU) to be flexibly assigned to hardware partitions. Via the crossbar switches, Fujitsu's Flexible I/O allows I/O units to be associated with any system boards, enabling partitions to be defined with the appropriate mix of processing power and I/O connectivity.

Figure 2 (previous page) illustrates dividing an eight-system-board PRIMEQUEST 480 server into four partitions. Each partition runs its own operating system, Linux or Windows, including different software releases or different software stacks. Figure 2 also illustrates that I/O units can be flexibly assigned: One partition has two system boards and three I/O units whereas another partition has four system boards and two I/O units, etc.

Note that Mirror Mode options do not have to be uniform across all partitions. That is, some partitions can be configured as Standard Mirror Mode while other partitions requiring additional redundancy can operate in Extended Mirror Mode.

Partitioning a server is also useful when performing concurrent maintenance activities or staged upgrades. System boards or I/O units can be electrically isolated from active partitions and then removed for repair or component upgrade. Once again, PRIMEQUEST's Flexible I/O permits system boards or I/O units to be isolated independently for service or upgrade. HP's hard partitioning does not have that flexibility. And, IBM's lack of hardware partitioning means its p5 servers cannot undergo concurrent maintenance or staged upgrades.

The IDEAS Bottom Line

Not many years ago, mission-critical computing was relegated to mainframes or a few select UNIX/RISC platforms. Back then, Linux and Windows focused on relatively modest configurations. Today, customers seek industry-standard hardware and software solutions for even their most important workloads. PRIMEQUEST addresses that need with solidly designed systems engineered for mission-critical environments.

From the midrange model 420 through the 32-Itanium-processor-chip model 480, PRIMEQUEST delivers a balance of performance, memory, and I/O capacity that positions it very favorably among competitive servers targeting enterprise environments. PRIMEQUEST's hardware-based partitioning offers configuration flexibility advantages over other hardware-based implementations and permits concurrent maintenance and upgrades that are not possible with software-based partitioning solutions.

Where Fujitsu's design clearly stands apart from other Itanium-based SMPs is in its focus on hardware high availability. Beyond extensive use of ECC and other traditional availability techniques, the unique mirrored crossbars and mirrored memory assures the most demanding customers that PRIMEQUEST platforms can be used to deliver mission-critical business computing.

Customers with enterprise computing needs have a variety of systems to consider, including models offered by traditional suppliers such as IBM, HP, Unisys, and others. Customers with mission-critical workloads should make sure that Fujitsu is on their short list and carefully examine PRIMEQUEST's advantages.

Americas
Ideas International, Inc.
800 Westchester Avenue
Suite S620
Rye Brook, NY 10573-1330
USA
Tel + 1 914 937 4302
Fax +1 914 937 2485

Asia/Pacific and Worldwide Headquarters
Ideas International Limited
Level 3
20 George Street
Hornsby, NSW, 2077
Australia
Tel +61 2 9472 7777
Fax +61 2 9472 7788

Europe, Middle East, Africa
Ideas International Europe
1 Deanes Close
Steventon
Oxon OX13 6SZ
United Kingdom
Tel +44 (0) 1235 437 850
Fax +44 (0) 1235 437 851

www.ideasinternational.com