WHITE PAPER

iCCS IDEAS Custom Consulting Services

PRIMEQUEST System Architecture

July 2006

PREPARED FOR

Fujitsu

TABLE OF CONTENTS

TO DIG DEEPER

This white paper prepared by Ideas International (IDEAS) offers an in-depth examination of the system architecture of Fujitsu's PRIMEQUEST servers. Three companion white papers provide similar detail on: high availability, operating systems, and system management. An overview paper provides a summary of the capabilities of the PRIMEQUEST offerings. For additional details, visit the Fujitsu website at http://www.fujitsu.com/global/.

Executive Summary

Observing that a growing number of customers were attracted to Linux and Windows operating environments and to servers employing Intel processors, in 2005 Fujitsu launched its PRIMEQUEST 400 Series employing Intel's Itanium 2 processors. With the availability of Dual-Core Intel Itanium 2 processors (previously known by the codename "Montecito"), the PRIMEQUEST line has been refreshed. The new PRIMEQUEST 500 Series models deliver the enhanced scalability enabled by the dual-core processors as well as new features uniquely offered by Fujitsu. Designed to meet the mission-critical needs of enterprise customers, PRIMEQUEST offers those customers a family of highly reliable servers that benefit from the economies of industry-standard hardware and software.

By its definition, the term "industry standard" implies that other systems exist that also employ Itanium Processor Family (IPF) chips running Linux or Windows. What differentiates Fujitsu's servers from other IPF-based platforms is the careful attention paid to mission-critical design. Fujitsu has leveraged its proven mainframe and PRIMEPOWER design experience to create an Itanium-based family that satisfies customer desires for running mission-critical workloads on industry-standard servers using broadly deployed operating environments.

In addition to extensive use of error checking and correction (ECC), PRIMEQUEST servers offer various options for mirroring critical components to insure high availability. The address and data crossbar switches that form the system interconnect are implemented redundantly so that they can be configured to carry duplicate signals. Using the mirrored redundant paths, the system can continue operating despite a failure in one section of the crossbar switches. Memory modules can also be redundantly configured.

Scaling up to 32 processor chips with 64 cores, PRIMEQUEST can handle the largest application workloads. For customers who do not need to devote all processors to a single application environment, PRIMEQUEST servers offer flexible partitioning. Partitioning takes advantage of the underlying modular building block approach consisting of system boards with processors and memory, and I/O units providing PCI slots and Gigabit Ethernet access. Up to 16 isolated partitions can be configured. Each partition runs its own operating system instance: various releases of Linux or Windows, each with custom software stacks. Each independent partition is electrically isolated from hardware or software failures in other partitions.

An important feature of Fujitsu's implementation, PRIMEQUEST partitioning permits I/O units to be assigned independently of the system boards. PRIMEQUEST

Flexible I/O allows I/O resources to be associated with any system boards, enabling partitions to be defined with the appropriate mix of processing power and I/O connectivity.

Overall, Fujitsu has leveraged its long experience in designing enterprise-class servers to create its Mission-Critical IPF Server. Even customers not familiar with prior Fujitsu high-end systems will recognize that the PRIMEQUEST platform offers unique features that make it attractive to those seeking industry-standard servers for critical business applications. The body of this paper describes, in more detail, some of the key architectural characteristics of the PRIMEQUEST server design. Companion papers focus on other PRIMEQUEST attributes, such as operating system support, high availability features, and system management capabilities. A "PRIMEQUEST Overview" white paper is also available.

Introduction

Customer Desire for Industry-Standard Mission-Critical Servers

Many customers are increasingly attracted to the open and thoroughly examined Linux operating environment as a secure and robust solution for their business-critical computing. Other enterprise customers value the broad application base available for the widely used Windows environment. These Linux and Windows users are similarly attracted to servers powered by industry-standard processors as a way to avoid lock-in, thanks to wide selection of platforms offered by a variety of vendors.

However, the choice of multiple platforms afforded by the use of industry-standard processors poses a challenge to vendors – how to differentiate their systems while adhering to those industry standards. For Fujitsu, differentiation comes from its long heritage of mission-critical server designs. This paper, along with the companion papers in this series, highlights the features that distinguish PRIMEQUEST servers from other industry-standard platforms supporting Windows and Linux.

Figure 1. High Level Diagram of PRIMEQUEST Architecture

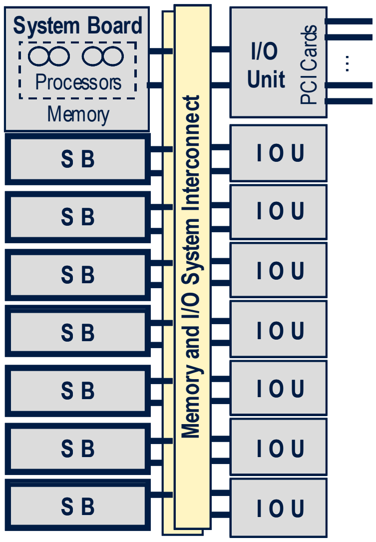

PRIMEQUEST Family

Basically, the PRIMEQUEST systems are scalable SMPs constructed from modular building blocks consisting of system boards, I/O units, crossbars, and System Management Modules. Each of the system boards (SBs) contains 4 processor chips and 32 DIMM slots for memory. The system boards plug into crossbar switches that allow all memory to be accessed by any processor. The crossbar switches also provide the paths that connect processors to the I/O subsystems. Figure 1 (left) depicts a high-level block diagram of the PRIMEQUEST architecture.

In some ways, Figure 1 appears similar to the high-level block diagrams of other modularly scalable SMPs; however, there are some important differences. For example, the interconnect crossbar is a key determinant of overall system performance. Fujitsu has leveraged its supercomputing and mainframe experience to design a crossbar switch that delivers higher bandwidth at lower latencies than competitive systems. As a result, PRIMEQUEST will deliver higher system performance than other implementations that may use the same processor chips.

Another item to note is that system boards and I/O units can be independently assigned to different partitions. Unlike other vendors, Fujitsu's approach does not predefine which I/O units are associated with particular system boards, and it permits partitions to flexibly contain different numbers of I/O units. Further discussion of PRIMEQUEST partitioning features is covered in a subsequent section. Additional unique capabilities of the Fujitsu designs, such as mirroring, are also covered in detail in subsequent sections.

For simplicity, Figure 1 does not show the dedicated System Management Board, which monitors the health of the server and provides configuration controls, including partitioning. Running a real-time Linux kernel on an embedded PowerPC 440 processor, the server management function included in the System Management Board eliminates the need for a system console or dedicated console. The figure also omits the optional Gigabit Ethernet Switch Board that provides an integrated switch for Gbit Ethernet connections to the I/O units.

The high-level diagram of Figure 1 depicts common characteristics of PRIMEQUEST servers, covering both the existing 400 series as well as the new 500 series. Differences between PRIMEQUEST models as well as future plans are summarized in the next few paragraphs.

The PRIMEQUEST 500 Series

Powerful processors form the foundation for high-end systems, and Fujitsu employs Intel's top-end Itanium 2 processors in PRIMEQUEST. The initial models, announced in 2004, used the single-core Itanium 2 processor known as "Madison 9M" that could contain up to 9 MB of on-chip cache. These 400 Series models were available in configurations supporting 8, 16, or 32 Intel Itanium 2 processor chips. With Intel's release of a Dual-Core Itanium 2 processor chip (previously code named "Montecito"), Fujitsu now introduces the PRIMEQUEST 500 Series, also available in configurations with 8, 16, or 32 processor chips, but containing twice as many processor cores compared to prior models. Table 1 summarizes the PRIMEQUEST family.

Table 1. 400 Series and 500 Series Models

PRIMEQUEST Model	Processor Cores Processor Type	Maximum Memory
580	64 cores 1.6 GHz Dual-Core Intel Itanium 2 processor	2 TB
540	32 cores 1.6 GHz Dual-Core Intel Itanium 2 processor	1 TB
520	16 cores 1.6 GHz Dual-Core Intel Itanium 2 processor	256 GB
480	32 cores 1.6 GHz single-core Itanium 2	1 TB
440	16 cores 1.6 GHz single-core Itanium 2	512 GB
420	8 cores 1.6 GHz single-core Itanium 2	256 GB

More cores can deliver more throughput; however, core count is but one contributor to system performance. Table 2 compares the top-of-the-line models 480 and 580 to highlight differences, and similarities, between the two models.

Table 2. Comparison of PRIMEQUEST 480 and 580 Models

	Model 480	Model 580
Processors	32 single-core Itanium 2 processors	32 Dual-Core Itanium 2 processors with Hyper-Threading
Clock / Cache	1.6 GHz / 9 MB, 1.5 GHz / 4 MB	1.6 GHz / 24 MB, 1.42 GHz / 12 MB
Front Side Bus	400 MHz	533 MHz
Processor Cores per System Board	4 single-core chips	4 chips with 8 cores
Memory	1 TB DDR2 (4 GB DIMMs)	2 TB DDR2 (8 GB DIMMs)
Crossbar Bandwidth (peak)	103 GB/sec.	137 GB/sec.
Internal Hard Disk	4.7 TB (32 drives @ 147 GB)	4.7 TB (32 drives @ 147 GB)
Optional I/O Units	8 each with 12 PCI-X slots	8 each with 12 PCI-X slots
PCI Slots	128 (32 internal, 96 in I/O boxes)	128 (32 internal, 96 in I/O boxes)
Partitions	8 4 socket granularity	16 2 socket granularity

Note that the dual-core processor chip contains a large on-chip cache that will help accelerate performance. Dual-Core Intel Itanium 2 processor "Montecito" implementation moves to a 90 nm fabrication technology from the 130 nm process of the single-core "Madison 9M." The smaller transistor sizes allow this dual-core chip to have 1.72 billion transistors compared to the 592 million transistors of the prior single-core implementation. As a result, Intel can place two full Itanium 2 processor cores on each chip as well as increase the L3 cache size to 12 MB per core. Furthermore, the dual-core Itanium 2 offers a form of multithreading, called Hyper-Threading that allows each core to execute instructions from two different code threads.

The increased L3 cache should help offset contention between the two cores for the common front-side bus. In addition, multithreading should keep the execution units busy even while one thread is waiting for data from memory. The benefits of multithreading depend greatly on the code being run, but could provide a boost of 10 to 20%. Overall, the combination of dual cores, large caches, and Hyper-Threading should enable the dual-core "Montecito" Itanium 2 chips to deliver up to 2.5 times the throughput compared to the single-core "Madison 9M" chips, across a range of applications including transaction processing workloads typically found on large SMP servers.[1] Furthermore, the increased throughput capability of the dual-core Intel Itanium 2 chips does not require additional electrical power. In fact, Intel has designed the "Montecito" dual-core chip to use less electrical power (and therefore require less cooling) than the single-core "Madison 9M" implementation.

Table 2 also indicates that the front-side bus runs faster on the dual-core chips, providing greater bandwidth to memory, I/O, and the PRIMEQUEST system interconnect. Fujitsu redesigned its supporting chip set to take advantage of the faster transfer rate, resulting in an increase of crossbar bandwidth as shown in the table. PRIMEQUEST's high-performance crossbar system interconnect is one of its differentiating advantages compared to other IPF-based servers. The PRIMEQUEST crossbar forms the interconnect that joins the Itanium 2 processors together to create a shared-memory multiprocessing server. Fujitsu engineers leveraged their mainframe and supercomputer expertise to devise a high-bandwidth, low-latency crossbar, surpassing competition by offering up to 137 GB/sec. (peak) bus bandwidth.

Beyond "Montecito," Intel plans a more powerful dual-core implementation (with the development code name "Montvale"), currently expected in the last half of the 2007 calendar year. Although specific details regarding future Intel chips are not available, Fujitsu has designed its PRIMEQUEST systems to be able to accommodate future generations of multicore Intel Itanium 2 processors, satisfying customer desires for investment protection.

Leadership through Mission-Critical Design

Advantages of IPF Processors

Intel has designed its Itanium processors to vigorously compete with the best of the RISC chips, not just in performance but also with built-in resilience to failures. From the beginning, Itanium's large L3 caches have been protected by ECC that can repair failures caused by a bad bit. The dual-core Intel Itanium 2 adds circuitry to allow malfunctioning sections of cache to be disabled. This dual-core chip implementation also includes a form of SpeedStep technology, first deployed in notebook chips, to slow down clock rates during lull periods to conserve power. In addition to saving power costs, SpeedStep will help the dual-core processor chip stay cooler to improve circuit reliability. Overall, Intel indicates that it has added a number of additional error-checking features as part of ongoing enhancements to Itanium's Machine Check Architecture (MCA). Although the specific details of MCA enhancements are beyond the scope of this paper, Intel has reported that it had subjected this dual-core Itanium 2 chip to a bombardment of gamma rays and alpha particles during testing of soft error recovery at a national laboratory. That accelerated stress testing indicated that individual chips possessed an average time between failures in the hundreds of years.

The 1.72 billion transistors afforded by the 90 nm technology also enable Intel to add functionality such as its Virtualization Technology. Intel Virtualization Technology is a hardware assist that enables systems to better run multiple operating systems and applications in different virtual partitions within the same processor. Note that Intel Virtualization Technology within Itanium chips complements, and does not replace, the physical partitioning that PRIMEQUEST servers offer (discussed later).

Error Checking and Correction in Fujitsu Chips

In addition to the ECC circuitry incorporated within Intel's Itanium 2 processors, Fujitsu has included ECC protection in the PRIMEQUEST chipset that its engineers designed. Two main areas that Fujitsu has protected by ECC are the system interconnects in the crossbar switch and the interface to memory.

The system interconnects are actually comprised of six separate crossbar switches: four handling data and two for addresses. Each of those interconnects is protected with ECC circuitry that allows single-bit errors to be fully corrected on the fly, without affecting system operation. Multiple-bit errors are detected, but those result in a machine check since they cannot be automatically corrected. Should the System Management Board processor determine that an interconnect bus is permanently defective, that bus can be deactivated.

Memory suffers not only from single-bit errors, but also from failures that disable an entire chip. PRIMEQUEST memory is protected by an advanced ECC code known as S4EC-D4ED.[2] Commonly described as “chip kill” coverage, this 16-bit ECC code can not only correct single errors, but it can also reconstruct all four bits from any failing memory chip.

Overall, Fujitsu has employed state-of-the-art ECC to protect those portions of the system that are most vulnerable to errors.

Mirrored Buses and Memory

Even though Fujitsu has incorporated traditional ECC on top of an already highly reliable Intel processor, there are mission-critical customer environments that warrant even higher levels of availability. For such situations, Fujitsu offers a spectrum of mirroring choices that are unmatched by any other high-end server. In particular, PRIMEQUEST systems offer various options for mirroring buses or memory to ensure continued operation in the face of failure. For PRIMEQUEST servers, two modes of mirroring are offered – Standard and Extended – in addition to the option to run in a non-mirrored configuration.

When configured as mirrored, resources are paired and perform the identical function in parallel. Should one of the pair fail, the other resource will still provide a valid result, allowing the system to continue without interruption. As shown in Table 3, Standard Mirroring starts by mirroring the global address crossbars that interconnect all system boards and I/O units. Extended Mirroring arranges the four global data crossbars as two duplicate pairs and configures memory so that all data is replicated in two separate memory modules.

Table 3. System Mirror Options Offered by PRIMEQUEST Servers

	Standard Mirror	Extended Mirror
Address Bus	Mirrored	Mirrored
Data Bus	Independent	Mirrored
SB Bandwidth	12.8 GB/sec. (400 Series) 17.06 GB/sec. (500 Series)	6.4 GB/sec. (400 Series) 8.53 GB/sec. (500 Series)
IOU Bandwidth	3.2 GB/sec. (400 Series) 4.26 GB/sec. (500 Series)	1.6 GB/sec. (400 Series) 2.13 GB/sec. (500 Series)
Memory	Independent	Mirrored
Memory Capacity	100%	50%
Performance	>99%	>95%

The global address crossbar broadcasts the address of requested data from memory or I/O to all 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. The PRIMEQUEST address crossbar is actually comprised of two complete address buses. Each system board has separate chips to interface to each of the address buses. In non-mirrored operation, both address buses are used. But, the total anticipated load on the two buses is low enough 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 is the default mirroring on PRIMEQUEST servers.

Extended Mirror goes a couple steps further by duplicating the global data crossbars and main memory. The global data crossbar consists of four separate data paths. Each is protected by single-bit error correction ECC, and has dedicated interface chips on the system and I/O boards. But for customers who want an extra degree of protection, the four data buses can be configured 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 duplicates data stored in main memory. As mentioned previously, memory is covered by an ECC code that not only corrects single-bit errors but also reconstructs all four bits from a defective memory chip. However, if additional multiple errors were to 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, mirrored memory provides 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. With the extensive amount of memory that can be configured on PRIMEQUEST servers, the reduction in total memory capacity should not be a concern. Nonetheless, having data stored in duplicate memory locations does double the cost for the memory. Therefore, Extended Mirroring does carry a financial cost related to doubling the amount of memory installed. As far as performance penalties, Fujitsu projects that mirroring the data buses does decrease available crossbar bandwidth but it extracts a performance penalty of less than 5%.

Cable Reduction Improves Reliability

Servers that scale to large numbers of processors, such as PRIMEQUEST, are not always run as a single shared-memory configuration under the control of a single instance of an operating system. Often, large servers are partitioned into smaller subsets to act as server consolidation platforms. Each of the subsets runs its own software and acts as if it were an independent server. That independence requires that each partition have independent connections to data LANs as well as connections to the separate management system LANs. Additionally, each of the partitions usually requires keyboard/video/mouse (KVM) connections.

For the 16 partitions available in high-end PRIMEQUEST systems, Fujitsu calculates that over 100 cables would be needed for all connections. In addition, external Ethernet connections consume PCI slots. The external LAN and KVM switches add additional cost, and the reliability of so many physical cables is likely to present problems.

The PRIMEQUEST solution provides prewired internal paths that connect to integrated and redundant switches for Gigabit Data LAN, System Management Board switches, and KVM. No rewiring is required to reconfigure partitions. This integrated approach presents an overall cost savings compared to the hardware cost, installation cost, and maintenance costs of physical cables and external switches. In addition, internal connections have higher reliability than external cabling.

Configuration Flexibility

Not many individual applications, particularly those running under Linux or Windows, need a system as large as the 64-core PRIMEQUEST 580 SMP. At the same time, customers value large systems as vehicles for consolidating the workloads of multiple servers into an easy-to-manage, central server. Some applications do not “play together well” as they may have conflicting software stacks; for example, they may depend upon different versions of the middleware software. In such cases, the large consolidation server needs to be partitioned into smaller operating system images. 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. Fujitsu builds upon its experience with PRIMEPOWER partitioning to provide unique capabilities for its PRIMEQUEST Itanium server.

Figure 2. PRIMEQUEST PPAR Flexible Partitioning

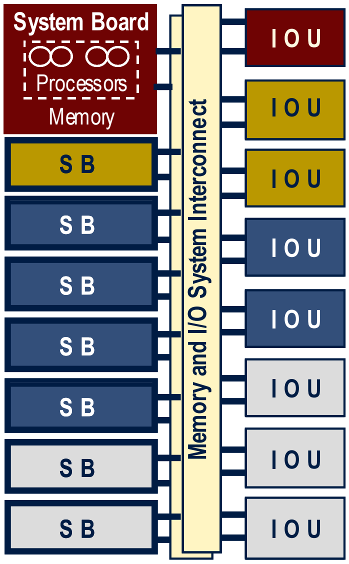

Partitioning Granularity

As described earlier, PRIMEQUEST systems are constructed from modular building blocks. Each of the system boards contains up to four processor chips, up to 256 GB of 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.

Reconfigurable system board partitioning is not exclusive to PRIMEQUEST systems. Similar capabilities can be found on other high-end server platforms. However, Fujitsu’s implementation on its PRIMEQUEST servers assigns I/O units independently of the system boards. Other designs sometimes include I/O connections on the system board or package system boards and I/O together in a way that limits the flexibility of assigning different combinations of system and I/O boards. The PRIMEQUEST Flexible I/O design allows I/O units to be associated with any system boards, through the crossbar switches, enabling partitions to be defined with the appropriate mix of processing power and I/O connectivity.

Figure 2 depicts an eight-system-board PRIMEQUEST server that has been divided into four partitions using PPAR physical partitioning. 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

Figure 3. Partitioning with XPARs

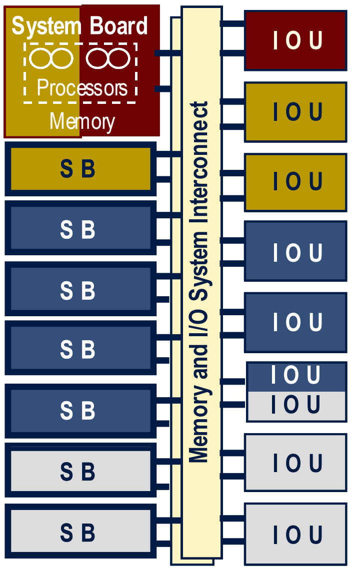

assigned: One partition has two system boards and three I/O units whereas another partition has four system boards and two I/O units.

Figure 3 illustrates the extended partitioning offered by PRIMEQUEST XPARs. The XPAR capability allows a system board to be split into two partitions, each with two dual-core processor chips, or four cores altogether. Similarly, I/O can be split and assigned to two different partitions. Using XPARs, the partition on the top of Figure 3 consists of half a system board (two chips, four cores) with one I/O unit. The next partition uses one and a half system boards (with a total of six chips and twelve cores) and two I/O units. The other two partitions share one of the I/O units. The smaller granularity afforded by XPARs offers a level of partitioning flexibility not available on other IPF-based servers.

Partitioning is available for non-mirrored operation as well as for the various mirrored alternatives. However there are some restrictions in mixing various modes due to engineering considerations. If the system is in non-mirrored mode, then all partitions must be non-mirrored. If the system is configured as mirrored, then each of the partitions must run in one of the mirrored modes: Standard Mirror or Extended Mirror. Since there are no performance or cost penalties for Standard Mirror, it is assumed that customers will choose to run with mirroring enabled. Applications requiring the additional availability of Extended Mirror can set up partitions to run in that mode while other partitions remain in Standard Mirror.

Concurrent Maintenance Capabilities

The electrical isolation fundamental to hardware partitioning ensures that faults in one partition do not affect other partitions. If a processor or memory module has failed and has been taken offline, the remaining processors and memory can continue to operate. However, to replace the failed chips requires that the full system board be removed. Similarly, individual PCI cards can be swapped if they have failed, but if the failure lies within the I/O unit circuitry, that I/O unit must be removed from service to repair.

A companion paper focusing on High Availability[3] covers the specific details on when failed components can be replaced without affecting other parts of the overall system. In general, if a component is not mirrored and it fails, then all partitions using that component are stopped. The faulty unit can be hot-swapped, with power remaining on, so that unaffected partitions can continue running. If a component that is mirrored by a redundant component fails, then the affected partitions can continue to run and maintenance can be deferred. When convenient, the affected partitions can be stopped and the failed component hot-swapped for repair. Table 4 (below) provides a summary of some of the major components that can be concurrently repaired. Refer to the companion High Availability paper, or to PRIMEQUEST documentation, for additional details.

Table 4. PRIMEQUEST Concurrent Repair Summary

Memory	In Extended Mirror, memory is configured redundantly and will not stop a partition. In Standard Mirror, the partition will fail, but it can be restarted to bypass the failed memory DIMM. Repair requires a system board swap.
System Board	To replace a system board, all partitions using that board must be stopped. Hot swap is supported to allow other partitions to continue operating.
Address Crossbar	In Standard Mirror, the address crossbar is configured redundantly. If one unit fails, it is isolated automatically. A faulty unit can be hot-swapped while isolated. A system restart is required to reactivate the repaired unit.
Data Crossbar	In Extended Mirror, the data crossbar is configured redundantly. If one unit fails, it is isolated automatically. A faulty unit can be hot-swapped while isolated. A system restart is required to reactivate the repaired unit.
I/O Unit	If failure occurs, the partition is stopped. The faulty I/O unit can be hot-swapped when the partition is stopped. Other I/O units can continue operation.
System Management Board	Can be redundantly configured. If the active System Management Board fails, the system automatically switches to the standby board. Hot swap supported.
KVM Switch	Not redundant, but does not crash system. Hot swap supported.
Integrated Gigabit Ethernet Switches	Failure will affect service to connected LANs. If partitions configure redundant GbE links, partitions may continue operation to alternate LANs. Hot swap supported.
Fan Tray	If redundantly configured, system can continue with disabled Fan Tray. Hot swap supported.
Power Supplies	N+1 power supported. System can continue with one power supply out of service. Hot swap supported.

The IDEAS Bottom Line

Fujitsu's PRIMEQUEST Mission-Critical IPF Server has been designed to address customer desires for an enterprise-class server utilizing industry-standard hardware and software. Scaling up to 64 Intel Itanium 2 processor cores and 2 TB of memory, it satisfies customer desires for a high-performance, shared-memory multiprocessor. Of course, not all mission-critical applications require an entire large SMP. Hence, PRIMEQUEST servers offer flexible hardware-based partitioning that creates multiple isolated partitions.

Fujitsu's design stands apart from other IPF-based SMPs through 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 can choose different levels of mirroring, selected on a per-partition basis, to satisfy the requirements of their workloads.

From a hardware perspective, Fujitsu has indeed created an impressive IPF server. Ideas International (IDEAS) believes that customers will recognize Fujitsu's care in designing a machine suitable for mission-critical computing and will welcome PRIMEQUEST servers as a competitive alternative to expensive proprietary systems. Timely support by Fujitsu partners, especially of PRIMEQUEST's unique mission-critical features, will play a key role in assuring the market acceptance of this computing platform.

Endnotes

1 Performance results measured on industry-standard benchmarks were not available when this paper was created. For the latest performance standings comparing servers running all processor types, see: http://www.ideasinternational.com/benchmark/bench.html.

2 Single-4bit block Error Correction – Double-4bit block Error Detection

3 See the companion white paper, "New Fujitsu High-End Dual-Core Intel Itanium 2-Based PRIMEQUEST Servers Offer the Utmost in High Availability," Ideas International, July 2006.

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