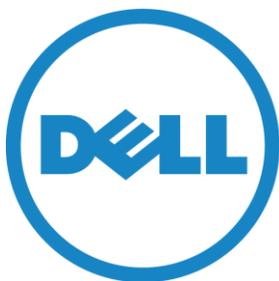


Advanced Reliability for Intel[®] Xeon[®] Processors on Dell[™] PowerEdge[™] Servers

A Dell Technical White Paper

Dell | Intel



THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information. The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order. Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web Site <http://www.intel.com/>.

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel® products, visit Intel Performance Benchmark Limitations.

Copyright © 2010 Intel Corporation and Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell, the *DELL* logo, and the *DELL* badge, *PowerEdge*, and *PowerVault* are trademarks of Dell Inc. Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and other countries. *Microsoft*, *Windows*, *Windows Server*, and *Active Directory* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

Contents

Executive Summary	2
Injecting Higher Value into Mission-Critical Computing	2
High-End Support for Today’s Most Demanding Workloads	5
Performance and Scalability	5
Reliability, Availability, and Serviceability (RAS)	6
Complete High Availability Solutions	6
The Importance of Memory Error Correction	8
Detailed RAS Features	9
Better Data Integrity	9
Error Prevention	9
Error Detection, Correction, and Containment	9
Improved System Availability	10
Automatic Recovery from Uncorrectable Errors	10
Interconnect Error Checking and Repair	11
Redundancy and Automatic Failover	12
Complete High Availability Solutions	14
Recommendations for Affordable High Availability	14
Conclusion	15
APPENDIX A: Glossary of Useful RAS Terms	16
About the Authors	17

Tables

Table 1. PowerEdge Server Processor Scalability	3
Table 2. Mission-Critical Features on Intel Xeon 7500 Series-Based Servers	8
Table 3. Protecting Data	10
Table 4. Increasing System Availability	13

Figures

Figure 1. R810 and M910 2 Processor to 4 Processor Configuration	4
Figure 2. R910 Expandable System Design	5
Figure 3. Complete High Availability Solutions	6
Figure 4. Advanced Redundancy and Failover Throughout	7
Figure 5. New MCA Recovery Interfaces	11

Executive Summary

Over the past decade, successful companies have increasingly turned to Dell™ PowerEdge™ Servers featuring Intel® Xeon® processors to support their mission-critical workloads, including database, enterprise resource planning (ERP), customer resource management (CRM), virtualization, and business intelligence applications. The proven capability and reliability of Dell PowerEdge Servers, combined with their unparalleled flexibility, have helped these companies stay at the forefront of technology, reduce total costs, and compete more effectively in today's challenging business environment.

The latest Intel® Xeon® processor 7500 and 6500 series can be expected to dramatically accelerate this trend. Four-socket servers based on this new processor family deliver the greatest performance and scalability increase ever for a new Intel Xeon processor generation. They also provide advanced support for 24/7 computing environments with abundant new reliability, availability, and serviceability (RAS) features. (The Intel® Xeon® processor 6500 series offers similar robust capabilities for the PowerEdge R810 rack server and the PowerEdge M910 blade server in two-socket configurations.)

Operating system (OS) support is available for these advanced RAS features today with increasing support over time, and Dell is delivering scalable and highly resilient servers that are ideal for critical business applications and large-scale consolidation. These systems provide:

- Superior data integrity with advanced support for error detection, correction, and containment across all major components and communication pathways.
- Improved system availability with multiple levels of redundancy, plus OS-assisted system recovery from certain uncorrectable errors that would have brought down previous-generation servers.
- Enhanced serviceability with predictive failure analysis that enables problematic components to be identified and replaced before they fail.

For anyone looking to optimize performance, scalability, reliability, and value for core business solutions, Intel Xeon processor 7500 series-based servers offer an essential new resource.

Injecting Higher Value into Mission-Critical Computing

Dell PowerEdge servers with Intel® processors have been established in the data center for over a decade. In recent years, they have become a high-value alternative to proprietary mainframe and high-end RISC systems as companies seek to drive down total cost of ownership (TCO) and provide a more flexible foundation for growth. They have been widely used for core business applications, including ERP, CRM, database, virtualization, and business intelligence solutions. They are also being used as consolidation platforms, since Intel® Virtualization Technology enables near-native application performance in virtual machines, even for I/O-intensive applications, such as enterprise databases and high-volume transactional applications.

With high-end scalability and advanced RAS support, the latest Intel® Xeon® processor 7500 series provides a dramatic increase in capability and value for mission-critical computing solutions. Servers based on this processor family support the highest levels of scalability, availability, and data integrity, but at a fraction of the cost of proprietary mainframe and RISC architectures and with far greater choice of hardware, software, and vendor support. In short, the reasons for using Dell Intel® Xeon® processor-based servers for core business applications have become more compelling than ever before.

Dell PowerEdge Server designs based on the Intel® Xeon® processor 7500 series scale up to four sockets (32 cores) as shown in Table 1.

Table 1. PowerEdge Server Processor Scalability

PowerEdge R910	PowerEdge R810	PowerEdge M910
2 or 4 processors	1, 2, or 4 processors	2 or 4 processors
Up to 32 cores/64 threads	Up to 32 cores/64 threads	Up to 32 cores/64 threads
Up to 64 DIMMS (1 TB)	Up to 32 DIMMS (512 GB)	Up to 32 DIMMS (512 GB)
4U form factor	2U form factor	10U Chassis with 8x M910 will support up to 256 Cores

With its 4U chassis and full memory allotment, the PowerEdge R910 provides the highest scalability and capacity. Through the use of Dell FlexMem Bridge, the PowerEdge R810 and M910 can selectively switch memory connectivity to allow usage of all DIMMs regardless of 2- or 4-processor configurations. Figure 1 shows the high-level connectivity of the system buses for these configurations. Figure 2 shows the high-level architecture of the R910.

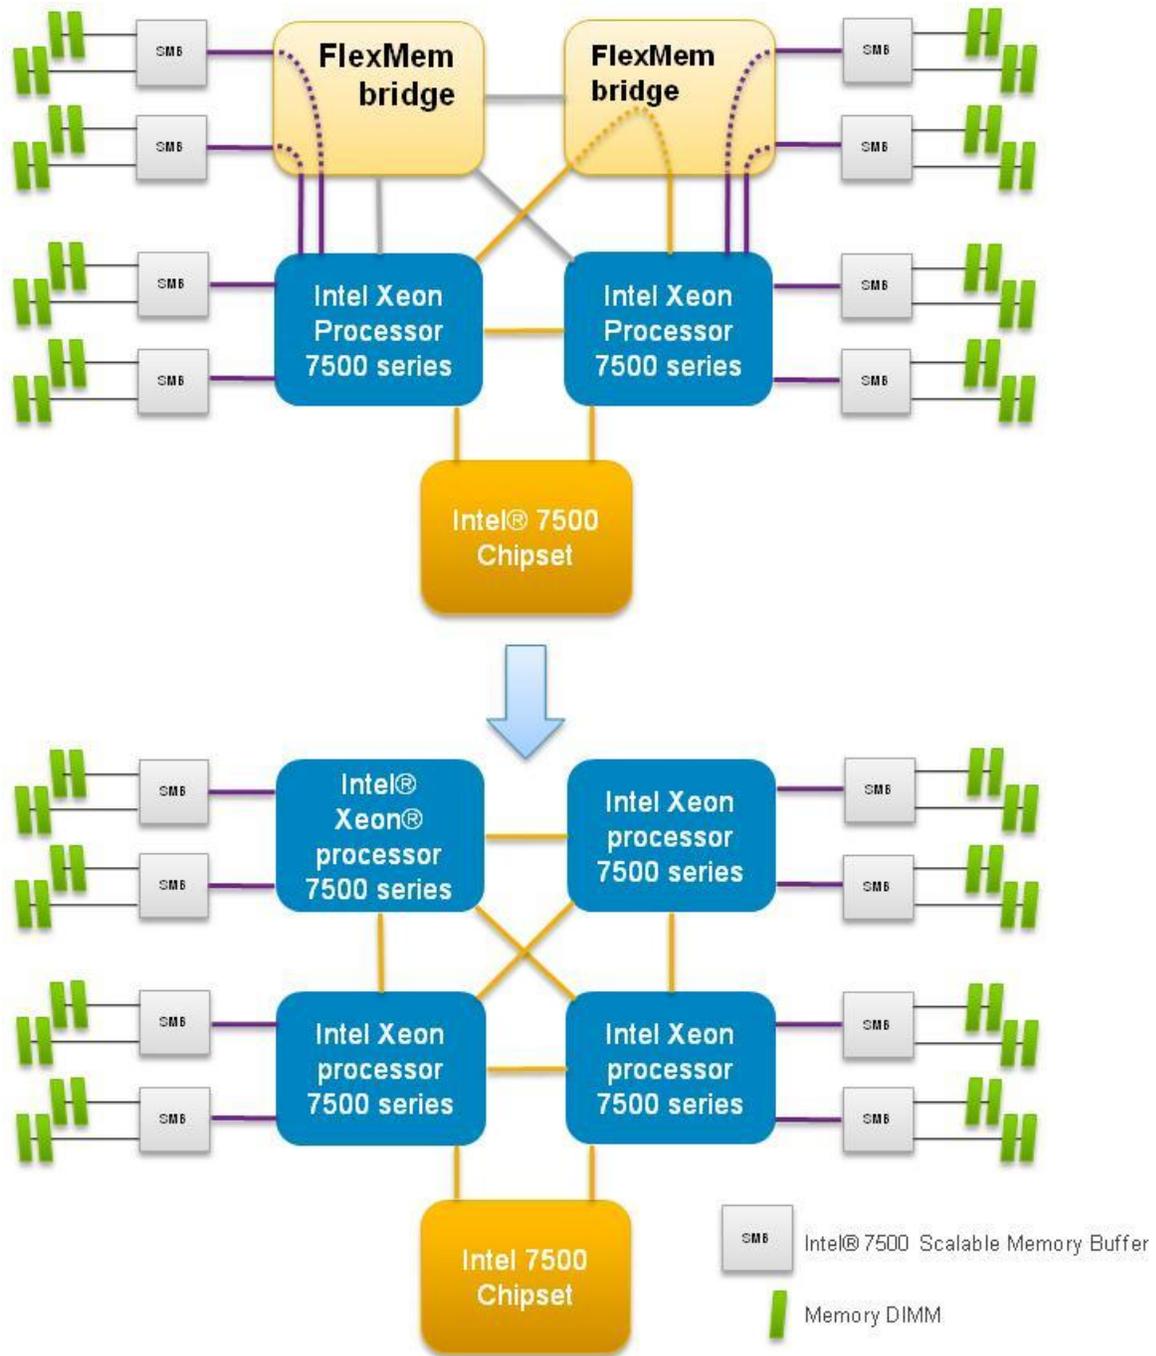


Figure 1. R810 and M910 2 Processor to 4 Processor Configuration

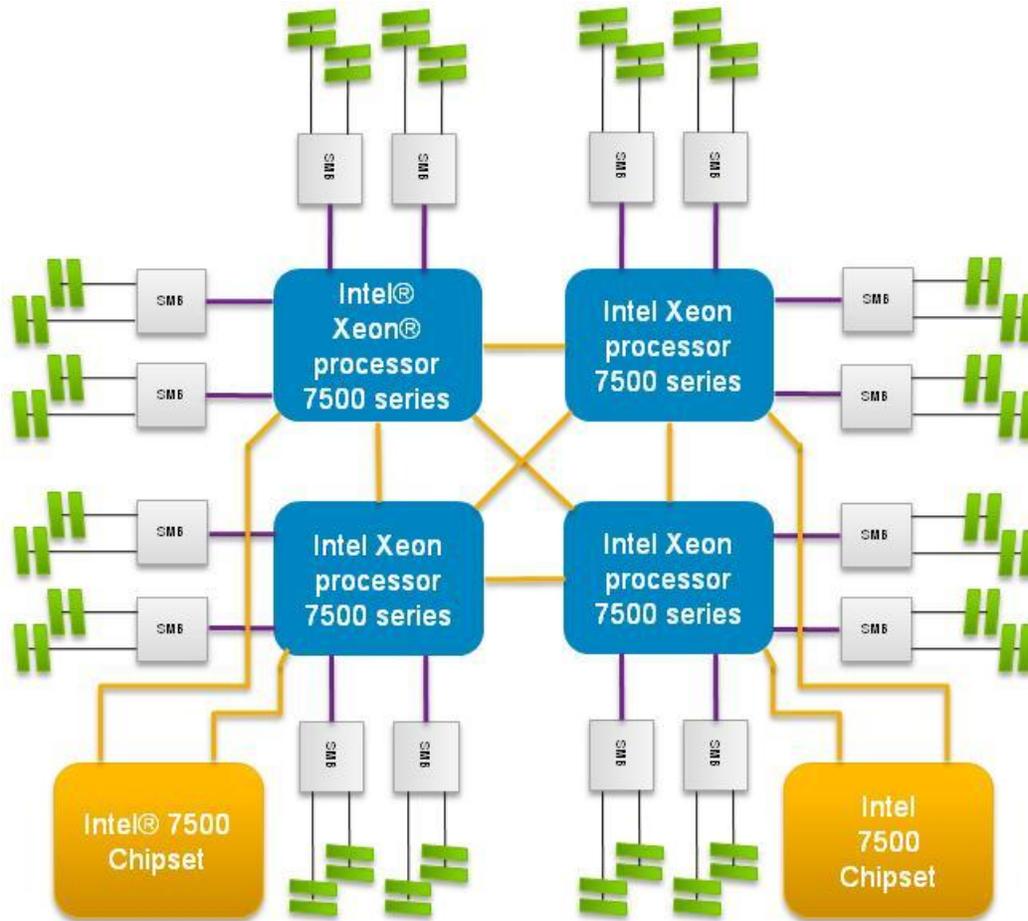


Figure 2. R910 Expandable System Design

High-End Support for Today's Most Demanding Workloads

Performance and Scalability

The Intel® Xeon® processor 7500 series provides the biggest leap ever in performance and scalability for a new Intel Xeon processor generation. Preliminary testing by Dell SPA labs has demonstrated a 3x increase in consolidation performance¹ and a 3x increase in database performance², with comparable gains across many other demanding workloads. This processor series provides up to eight high-performance cores, 16 execution threads, and 24 MB of last level cache, compared with six cores, six threads and 16 MB of cache in previous generation processors. It also delivers increases of up to 6x in both memory bandwidth³ and 4x memory capacity. This dramatic leap in performance and scalability is

¹ Source: Dell internal measurement (April 2010) using the VMmark benchmark.

² Source: Dell internal measurement (March 2010) using TPC-E* benchmarks.

³ Source: Dell internal memory bandwidth measurement (March 2010).

only half the story. The other half is the equally dramatic increase in RAS features integrated into this new processor family.

Reliability, Availability, and Serviceability (RAS)

Companies across many industries are focused on becoming more responsive to their customers and the marketplace by integrating their business applications more extensively and moving toward real-time business processes. As a result of this trend, downtime for a key business application will often impact transactions across many other systems and applications. It can also halt core processes, impede customer and partner transactions, impair productivity for large numbers of employees, and damage credibility in the marketplace. For many businesses, the failure of a core business application can easily cost hundreds of thousands to millions of dollars per hour.

Server clustering and live virtual machine migration is helping businesses meet rising high-availability requirements. However, many IT organizations are not comfortable clustering and virtualizing certain applications, such as large databases and high-end transactional applications. In addition, a fatal system error can potentially bring down a server before failover can complete. In many scenarios, this is not an acceptable risk.

Intel Xeon processor 7500 series-based servers help to mitigate this risk by providing silicon-level support for an abundance of new advanced RAS features. These features provide the foundation for advanced system-level RAS, and Intel has been working extensively with server, OS, virtual machine monitor (VMM), and application vendors to ensure tight integration and broad support across the entire hardware and software solution stack. These efforts are enabling the delivery of fully realized solutions targeting the needs of mission-critical applications (Figure 3).

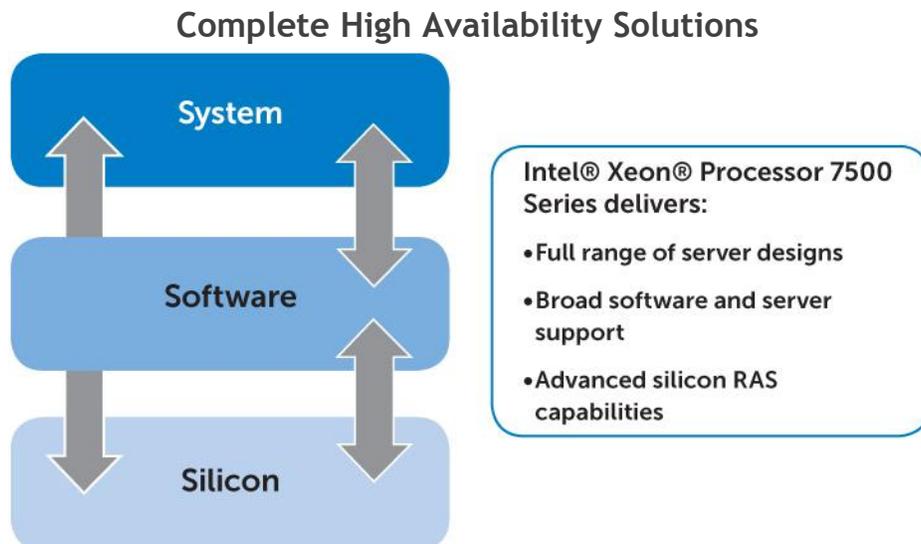


Figure 3. Complete High Availability Solutions

Key benefits include:

- **Robust Data Integrity.** Data errors are prevented, detected, corrected, and contained more comprehensively and effectively to preserve data integrity. If an uncorrectable error does occur, it is tagged⁴ and contained to help prevent propagation to other systems and applications.
- **Improved System Availability.** Machine Check Architecture Recovery (MCA Recovery) enables OS-assisted system recovery from certain uncorrectable errors that would have brought down previous-generation servers. Microsoft, Novell, Oracle, Red Hat, and VMware have announced support for MCA Recovery. In each case, these capabilities are either supported in currently available versions or will be supported in a future scheduled release. Redundancy and automated failover capabilities are also integrated extensively at both the silicon level (Figure 4) and the system level to increase uptime. With the release of new PowerEdge servers with Intel Xeon processor 7500 series, Dell has also added a redundant embedded hypervisor option for supporting VMware ESXi.
- **Enhanced serviceability.** Enhanced error logging and reporting enables predictive failure analysis to identify problematic components before they cause downtime or uncorrectable errors.

Advanced redundancy and failover features are integrated throughout the Intel® Xeon® processor 7500 series to deliver superior resilience for mission-critical computing environments as shown in Figure 4.

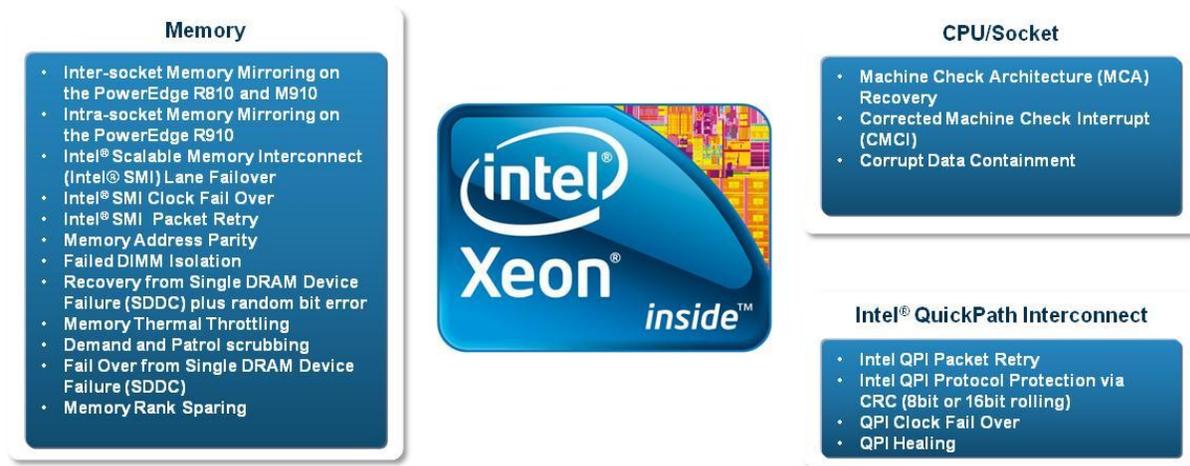


Figure 4. Advanced Redundancy and Failover Throughout

With this combination of advanced capabilities and broad vendor support, the launch of the Intel Xeon processor 7500 series marks a major advance in the overall value of Intel Xeon processor-based servers for mission-critical computing solutions. IT organizations can now take advantage of unprecedented performance, scalability, and RAS on the world’s most flexible and widely supported computing architecture.

⁴ The process of tagging data to prevent propagation is traditionally known as “data poisoning.”

Table 2. Mission-Critical Features on Intel Xeon 7500 Series-Based Servers

Benefits for IT	Silicon Features
Protects Data	
<ul style="list-style-type: none"> • Reduces circuit-level errors • Detects data errors across the system • Limits the impact of errors 	Parity checking and Error Correction Code (ECC)
	Memory thermal throttling
	Memory demand and patrol scrubbing
	Corrupt data containment mode ⁵
Intel® QuickPath Interconnect (Intel® QPI) protocol protection via Cyclic Redundancy Checking (CRC)	
Increases Availability	
<ul style="list-style-type: none"> • Heals failing data connections • Supports redundancy and failover for key system components • Recovers from uncorrected data errors 	Machine Check Architecture Recovery (MCA Recovery) with operating system support ^{5,6}
	Intel® Scalable Memory Interconnect (Intel® SMI) lane failover
	Intel SMI clock failover
	Intel SMI and Intel QPI packet retry
	Intel QPI clock failover
	Intel QPI self-healing
	Single Device DRAM Correction (SDDC) plus random bit error recovery
	Memory mirroring
Memory rank sparing	
Minimizes Planned Downtime	
Helps IT predict failures before they happen	MCA error logging (CMCI) with operating system predictive failure analysis ⁶

The Importance of Memory Error Correction

Google engineers and the University of Toronto worked together to study real-world memory errors across hundreds of thousands of active servers in Google data centers. Among the results from the two-and-a-half year study:

- Memory errors are an order of magnitude more common than previously thought (more than eight percent of DIMMs were affected by errors per year in Google’s computing environment).
- Hard errors, the kind that cause system failure, account for a much higher percentage of total errors than anticipated.

⁵ Currently only supported on the PowerEdge R910. Support planned for the PowerEdge R810 and M910 later in 2010.

⁶ Requires OS support. Check with your operating system vendor for current or planned support.

- Current errors in a memory DIMM provide a strong indication of future errors in the same component.

Based on these results, the error and memory management features provided by the Intel Xeon processor 7500 series, such as automated error detection and correction, predictive failure analysis, DIMM failover, and operating system-assisted system recovery (through Machine Check Architecture Recovery) are essential technologies in data-intensive, mission-critical computing environments.

Download the complete study, DRAM Errors in the Wild: A Large-Scale Field Study, by Bianca Schroeder of the University of Toronto and Google engineers Eduardo Pinheiro and Wolf-Dietrich Weber, at:

www.cs.toronto.edu/~bianca/papers/sigmetrics09.pdf

Detailed RAS Features

Four socket and larger servers based on the previous-generation Intel Xeon processor 7400 series already offer a more extensive array of advanced RAS capabilities than other servers in their class. The new Intel Xeon processor 7500 series takes RAS support to the next level, delivering high-end capabilities that are unprecedented in open system server architectures.

Better Data Integrity

In any computing system, errors in data or processing can affect data reliability, system availability, or both. There are basically two kinds of errors:

- Soft errors are typically caused by an alpha particle or other micro event that changes the logic state of one or more silicon gates. Most can be fixed simply by correcting the logic state.
- Hard errors are more persistent. They imply a failure at the hardware level and require a permanent fix to ensure correct and uninterrupted operation.

The Intel Xeon processor 7500 series incorporates extensive features for preventing, detecting, correcting, containing, and reporting both kinds of errors on the processor die, in attached components and along the pathways that connect those components. These mechanisms provide robust support for maintaining data integrity and keeping mission-critical services online.

Error Prevention

Overheating of memory components can cause errors and accelerate component failure. To address this risk, the Intel Xeon processor 7500 series monitors memory temperatures and can either reduce the rate of commands issued to the memory controllers or, in conjunction with Dell's systems management solution, increase fan speeds as needed to keep memory components operating within acceptable thermal limits.

Error Detection, Correction, and Containment

Though thermal throttling can help reduce the frequency of errors, errors can and do happen in all computing systems. All Intel Xeon processors incorporate advanced mechanisms for automatically detecting and correcting errors. The Intel Xeon processor 7500 series extends these capabilities to provide more comprehensive coverage in both the processor and the memory subsystem (Table 3). Many of these mechanisms are performed automatically and transparently to detect and correct errors. When errors cannot be corrected automatically, they are tagged and reported to the OS for remediation if the O/S does support.

Table 3. Protecting Data

ADVANCED DATA INTEGRITY FEATURES (Supported in the processor or the associated Intel® chipset)	Intel® Xeon® processor 7400 series	Intel® Xeon® processor 6500/7500 series
Processor/Socket		
Corrupt data containment mode: Memory locations with corrupted data can be tagged (sometimes called “data poisoning”). This limits the impact to the currently running program and greatly reduces the need to reset the system. ⁷		✓
Memory		
Parity checking and Error Correcting Code (ECC): Well established algorithms are used to detect and correct soft errors.	✓	✓
Thermal throttling: Memory command rates can be reduced or fan speeds increased to keep memory components operating within acceptable thermal limits.	✓	✓
Demand and patrol scrubbing: Memory is constantly monitored to detect and correct errors so they don’t accumulate and become uncorrectable. Performed proactively (patrol) and when a correctable error is detected during a read transaction (demand).	✓	✓
Intel® QuickPath Interconnect (Intel® QPI)		
Intel QPI protocol protection via Cyclic Redundancy Checking (CRC)		✓

Improved System Availability

The Intel Xeon processor 7500 and 6500 series establishes a strong foundation for highly available servers with a wide variety of micro-architectural enhancements that help to increase resilience both for individual components and for the system as a whole (Figure 4).

Automatic Recovery from Uncorrectable Errors

The Intel Xeon processor 7500 and 6500 series support MCA Recovery in Dell PowerEdge servers. This silicon-level technology interfaces directly with the host OS or VMM to enable automatic recovery from certain uncorrectable errors that would have caused crashes in previous-generation servers (Figure 5). When an uncorrectable error is detected, the silicon interrupts the OS or VMM and passes along the address of the memory error. The OS or VMM then resets the error condition, marks the defective memory location so it will not be used again, and continues operation. (If the memory location is being used for certain critical kernel operations, the system or application may not be able to continue and will be shut down by the OS.)

Machine Check Architecture Recovery (MCA Recovery) is a mechanism where the silicon works with CMCI-enabled operating systems in an attempt to prevent the server from utilizing areas of main

⁷ Currently only supported on the PowerEdge R910. Supported planned for the PowerEdge R810 and M910 later in 2010.

memory susceptible to uncorrectable errors which would have otherwise caused a system crash in prior generations, as illustrated in Figure 5.

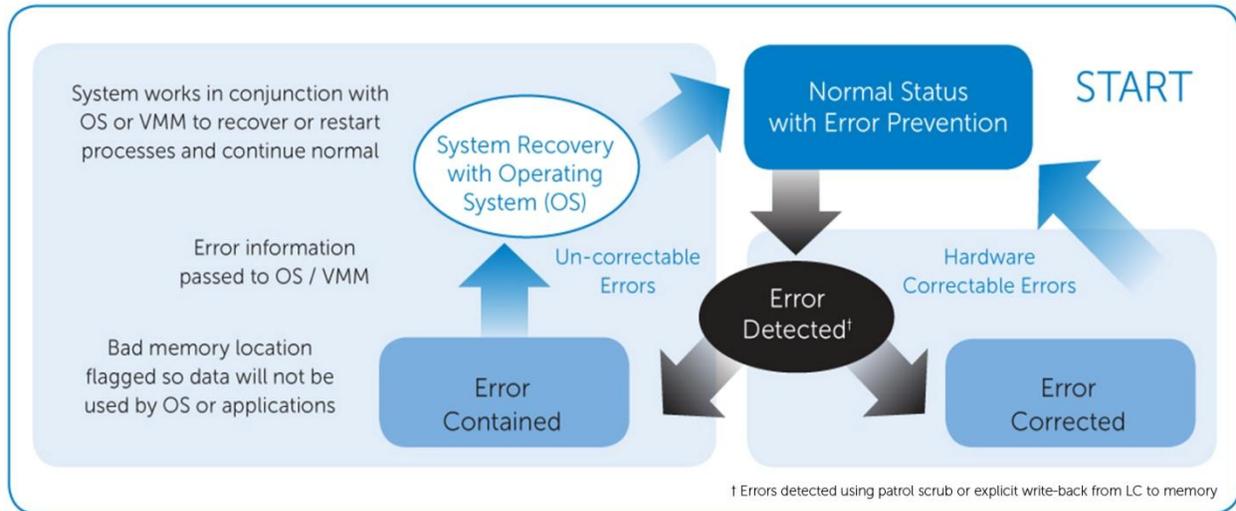


Figure 5. New MCA Recovery Interfaces

Intel has added enhancements to its processor’s Machine Check Architecture to support MCA Recovery from two specific types of memory errors—starting with Intel Xeon processor 7500 and 6500 series. In previous generations, these two specific MCA errors would result in Uncorrectable Memory Errors (failure):

- Uncorrectable error encountered during Patrol Scrub
- Triple Bit ECC error on Last Level Cache (LLC) read when generating an Explicit Write Back Operation from LLC to memory

Through the use of MCA recovery, system failures stemming from these scenarios can now be avoided by preventing the incorrect data from being consumed by an application or end target device.

Interconnect Error Checking and Repair

Errors can happen to data and instructions as they are transmitted among processors, or from processors to memory or I/O devices. The Intel Xeon processor 7500 series includes new, high-bandwidth interconnect systems for these communication tasks: the Intel® Scalable Memory Interconnect (Intel® SMI) and the Intel® QuickPath Interconnect (Intel® QPI). Both include extensive error detection and recovery mechanisms. They also include mechanisms that enable uninterrupted operation in the event of a hard failure, such as a failed lane or clock.

- Intel SMI lane failover. Intel SMI provides a spare lane for traffic in both directions. If a single lane fails or shows persistent errors, the processor can failover to the spare lane to provide uninterrupted operation with no loss of bandwidth.
- Intel QPI self-healing. A port that shows persistent errors can automatically map out the failed transmission path. In this case a full-width port is reduced to a half-width port or a half-width port to a quarter-width port. There is some loss of bandwidth, but operation can continue.
- Dell’s PowerEdge R810 and M910 servers, when configured with two processors and two FlexMem bridges, utilize an additional QPI link between the processors. This increases

available bandwidth for inter-processor traffic, as well as reduces the impact to performance in the event of a QPI width reduction on one of those links.

Redundancy and Automatic Failover

Intel Xeon processor 7500 series-based servers can automatically detect failing memory and failover to spare components to continue uninterrupted operation. In the case of memory, the functional state of the failing component (for example, register states data for a DIMM) can be migrated to the spare component so operations can be continued without interruption.

The Intel Xeon processor 7500 series employs a variety of advanced mechanisms to enable uninterrupted operation in the event of uncorrectable hard or soft errors as shown in Table 4.

Table 4. Increasing System Availability

Advanced High-Availability Features to Increase Uptime (Supported in the processor or the associated Intel® chipset)	Intel® Xeon® processor 7400 series	Intel® Xeon® processor 6500/7500 series
Uncorrectable Data Error Recovery		
Machine Check Architecture Recovery. ⁸ (MCA Recovery): Interfaces with the silicon-based error detection and correction mechanisms and an enabled operating system (OS) to recover from otherwise fatal system errors, enhancing both data integrity and system uptime. ⁹		✓
Intel® QuickPath Interconnect (Intel® QPI)		
Intel QPI packet retry: Enables retransmission of data or instructions following the detection of an error.		✓
Intel QPI clock failover: In the event of a forwarded clock failure, forwarded clocks can be redirected to one of two failover clock lanes to enable uninterrupted operation.		✓
Intel QPI self-healing: Enables an Intel QPI link to map out a failed lane or lanes and continue operation in the event of persistent errors. Bandwidth is reduced but operation can continue.		✓
Memory		
Intel® Scalable Memory Interconnect (Intel® SMI) packet retry: Enables retransmission of data or instructions following the detection of an error.		✓
Intel SMI lane failover: Enables identification and mapping out of a bad data path to reduce errors and maintain operation.		✓
Intel SMI clock failover: In the event of a forwarded clock failure, forwarded clocks are redirected to the clock failover lane to enable uninterrupted operation.		✓
Single DRAM Device Data Correction (SDDC) plus random bit error recovery: Supports continued operation in the event of a single DRAM device failure. Can additionally correct single bit memory errors on the same DIMM.	SDDC only	SDDC plus random bit error recovery
Memory Mirroring: A backup copy of main memory can be maintained on a second DIMM attached to the same or a different socket. Failover to the mirrored DIMM is automatic. Cannot be enabled concurrently with memory sparing.	Intra-socket	Intra-socket or inter-socket
Memory sparing: Data from a failing rank pair is copied to a spare rank pair behind the same memory controller. The failed component is mapped out to enable uninterrupted operation. The OS is not involved. Cannot be enabled concurrently with memory mirroring.	✓ (DIMM)	✓ (Rank ¹⁰)
Failed DIMM isolation: Enables the identification of a failed DIMM channel pair, so only the identified DIMM pair need be replaced.	✓	✓

⁸ Requires operating system (OS) support. Check with your operating system vendor for current or planned support.

⁹ Currently only supported on R910. R810 and M910 will be supported later in 2010.

¹⁰ Dell PowerEdge Servers based on Intel 7500 Series implement Rank sparing, which incurs a lower penalty to usable memory size.

Complete High Availability Solutions

Dell is taking advantage of the advanced RAS features of the Intel Xeon processor 7500 series to deliver affordable servers that are scalable, flexible, manageable, and resilient. Features and capabilities vary by model and include:

- Redundant components with rapid failover capabilities, memory and internal dual SD module for hypervisor redundancy
- Integrated monitoring, diagnostics, and failure prediction
- Advanced virtualization and workload management for more granular and dynamic allocation of resources in consolidated environments

With the Intel Xeon processor 6500 and 7500 series on PowerEdge servers, MCA error logging with OS predictive failure analysis using Corrected Machine Check Interrupts (CMCI) enables the processor to trigger interrupts on a corrected machine check event to notify the operating system for predictive failure analysis. Never before has such an extensive range of high-availability options been available on Intel Xeon processor-based servers. With this support, IT organizations can meet a wide array of high availability requirements without resorting to costly RISC and mainframe architectures. The potential for cost savings is substantial today and will continue to grow in years to come.

Recommendations for Affordable High Availability

Resilient servers are an important component of any high-availability environment, but they are only one component. Delivering true high availability takes a combination of people, processes, and technology including:

- Highly reliable platforms
- Extensive hardware and software testing
- Rigorous change management
- Redundant architectures
- Highly trained staff using effective management tools
- Well-established emergency procedures

With a sufficient investment, level of availability can be achieved, yet costs can be prohibitive as IT organizations strive to increase availability guarantees from 3-nines (99.9% uptime), to 4-nines (99.99%), to 5-nines (99.999%), and higher. The following high-level strategies can be particularly helpful for containing costs while addressing increasing availability requirements:

- **Standardize your infrastructure and operations.** One-of-a-kind solutions and nonstandard technologies increase design, deployment, and support costs. They require specialized expertise and nonstandard processes that can lead to inefficiency, lower quality, and reduced availability. Enterprise standards based on industry-standard systems and technologies help to overcome these limitations by establishing a consistent computing environment across hardware, software, and processes. A standardized environment also reduces training requirements and lets IT organizations leverage the skills of its most competent personnel more effectively.
- **Focus on service delivery.** Server uptime, alone, is not a meaningful measure of high availability. The business value of an application depends on the ability of end-users (and other applications and services) to access and use that application, which may depend on multiple

applications, servers, networks, and so on. Focusing on the availability of the service, rather than the system, can help IT organizations deliver better overall value to the business.

- **Measure the business value of high availability.** Understanding the cost and risk of service downtime, from both business and IT perspectives, is valuable. It allows standard return-on-investment metrics to be established, so decision makers can align high-availability investments with the actual business value they deliver.

Conclusion

Dell PowerEdge servers based on the latest Intel Xeon processor 6500 and 7500 series offer exceptional new capabilities and value for mission-critical applications and large-scale consolidation. These servers deliver the biggest ever generation-to-generation performance gain for an Intel® processor. IT organizations will have wide access to systems with up to 32 processor cores, 64 execution threads, and 1 TB of memory.

This dramatic increase in performance and scalability is combined with an abundance of new RAS features to improve data integrity, system availability, resource management, and serviceability. Advanced error detection, correction, and containment mechanisms are combined with OS-assisted capabilities such as automatic system recovery and predictive failure analysis. Together, these features provide a robust foundation for high-end availability using open systems that are far more flexible and affordable than proprietary RISC and mainframe architectures.

With these advancements, Intel Xeon processor 7500 series-based servers are ushering in a new era of flexibility for mission-critical computing and large-scale consolidation. For anyone looking to inject higher value into a data center, these servers offer an essential new resource.

APPENDIX A: Glossary of Useful RAS Terms

RAS	Reliability, Availability, and Serviceability
Reliability	Assurance that computational results are correct. Errors are detected and corrected when possible and reported if they cannot be corrected.
Availability	Assurance that the system is up and running to support an organization's computing needs.
Serviceability	Assurance that errors are reported and faulty components can be identified and replaced.
Error Detection/Correction	Ability to detect and correct hard and soft errors to increase reliability and availability.
Soft Errors	A transient error that can be corrected by overwriting with the correct data.
Hard Errors	A persistent error that cannot be fixed by overwriting with the correct data (for example, a faulty logic gate).
Parity, Error Correction Code (ECC), Cyclic Redundancy Check (CRC), Polarity	Widely used mechanisms for identifying hard and soft errors.
Field Replaceable Unit (FRU)	A system component that can be physically added, removed, or replaced in the field, such as a processor board, an I/O Hub, a DIMM, or a PCI card. An FRU may or may not be hot pluggable.

About the Authors

Mohan Kumar, Senior Principal Engineer, Intel Architecture Group, Intel Corporation

Mohan is responsible for Server Reliability, Availability and Serviceability (RAS) in the Architecture and Planning organization of the Intel Digital Enterprise Group. He is key lead for the engineering working group that developed the RAS flows for the high-end RAS features implemented in the Intel Xeon processor 7500 series and associated Intel components.

Michael Demshki, Marketing Segment Manager for Mission Critical Servers, Data Center Group, Intel Corporation

Michael has been with Intel for over 13 years focusing on the application of Intel processor technology to ever-growing IT server computing requirements. Prior to joining Intel, Michael spent 8 years marketing high-end servers with Sequent Computer Systems.

Robert Shiveley, Mission Critical Server Product Manager, Data Center Group, Intel Corporation

During his 12 years at Intel, Robert has worked in enterprise product and solutions marketing involving enterprise manageability, security, IT infrastructure optimization, IT business value and enterprise regulatory compliance. Prior to Intel, he worked in strategic and product marketing, solution sales and business development at enterprise ISVs.

Rob Bassman, Senior Systems Engineer, Enterprise Product Group, Dell Inc.

Rob has been with Dell for more than 14 years, designing and developing mainstream and high-end PowerEdge products since the Server Groups' inception. As a platform team lead Rob is responsible for ensuring that quality, functionality and customer experience requirements are met across all disciplines.

Learn more: www.dell.com/poweredge