

Hot plug RAID memory technology for fault tolerance and scalability

technology brief



Abstract.....	2
Introduction.....	2
Memory reliability.....	2
Error detection and correction	3
Parity checking	3
Error checking and correcting.....	3
Potential for system failures.....	3
Hot plug RAID memory	4
Performance	5
Basic operation.....	5
Hot-plug capabilities.....	8
Ease-of-use capabilities	9
Conclusion.....	9
Call to action	10

Abstract

This technology brief describes the Hot Plug RAID Memory technology developed by HP to give enterprise-class servers the level of memory fault tolerance today's 7x24 applications demand. It provides background information on memory reliability, reviews current error detection and correction techniques, and explains why the likelihood of memory errors grows as memory capacity increases. It discusses Hot Plug RAID Memory in depth and provides information on less robust, alternative fault-tolerant memory solutions.

Introduction

The 1990s brought fundamental changes in enterprise computing. The proliferation of web browsers and the Internet led to a dynamic, global marketplace that demands instant answers, products, and services. Customer requirements for a high-performance, highly available, and easily managed computing infrastructure have increased exponentially.

As a result, the changes of the 1990s spurred innovation in one of the most critical subsystems of enterprise-class servers: memory. Operating system support for more than 4 gigabytes (GB) of memory and availability of low-cost, high-capacity memory modules have driven requirements to support unprecedented memory capacity in today's industry-standard servers. Recent ProLiant servers support up to 64 GB of memory, and memory capacities will continue to grow in the near future.

Error checking and correcting (ECC) memory, introduced in PC servers in 1992, still offers excellent protection for many servers. As memory capacity grows, however, the level of effectiveness ECC provides actually decreases.

HP developed Hot Plug RAID Memory to extend the effectiveness of ECC and give enterprise-class servers the level of memory fault tolerance today's 7x24 applications demand. Hot Plug RAID Memory provides redundancy and hot-plug capabilities for industry-standard dual inline memory modules (DIMMs) to deliver unprecedented levels of availability, scalability, and fault tolerance.

Memory reliability

A well-designed memory subsystem, such as those employed in ProLiant servers, can be extremely reliable. For example, the memory subsystems in ProLiant servers are designed and extensively tested to ensure the highest quality possible. The memory modules in ProLiant servers undergo extensive qualification through the HP World Class Suppliers Process to ensure compliance with the industry-standard specifications.

Memory system integrity begins with the reliability of the DIMMs. All ProLiant servers use industry-standard DIMMs, but just meeting industry standards is not enough. Rigorous testing also ensures that all DIMMs in ProLiant servers meet exacting electrical standards.

Because memory is an electronic storage device, it has the potential to return information different from what was originally stored. Dynamic random access memory (DRAM) stores ones and zeros as charges on extremely small capacitors that must be frequently refreshed to ensure the data is not lost. Every bit of memory is either a zero or a one, the standard in a digital system. A relatively small electrical disturbance near the memory cell can alter the amount of charge on the capacitor, changing the state of the data bit stored in that memory cell and causing a memory data error.

Two kinds of errors can typically occur in a memory system. The first is called a hard, error and is characterized by the fact that it is repeatable, though it may be very inconsistent. In this situation, a piece of hardware is broken and will continue to exhibit incorrect behavior over time. For example, a

bit may be stuck so that it always returns "0", even when a "1" is written to it. Hard errors indicate physical problems such as memory defects or a broken connection.

Most errors that occur in the memory subsystem are soft errors. A soft error is a randomly occurring event that causes the data stored in a device to be changed. Because a soft error is not caused by a problem with the circuit, once the data is corrected, the error will not recur.

Error detection and correction

The only true protection from memory errors is to use some sort of memory detection or correction protocol. Some protocols can only detect errors, while others can both detect and correct memory problems, seamlessly.

Parity checking

Parity checking is the most basic form of memory error detection. Although it detects many errors, it does have some drawbacks. Parity checking can only reliably detect a single-bit error. In addition, parity checking cannot locate and correct erroneous data. Even if parity checking detects an error, it has no ability to correct the error, and the server will halt operation.

Error checking and correcting

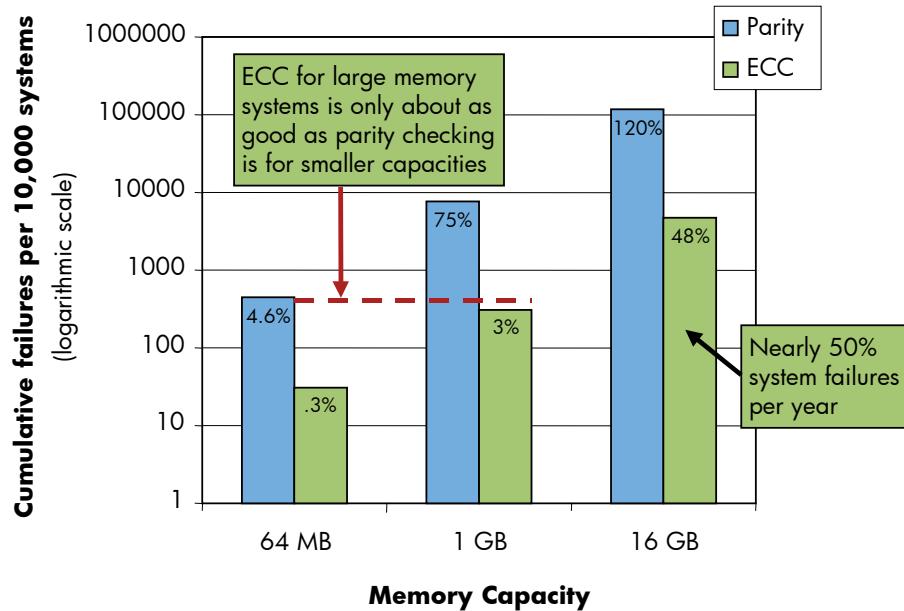
ECC memory is now standard in all ProLiant servers and significantly reduces the probability of fatal memory failures. The ECC commonly used in industry-standard servers is superior to parity checking because this ECC not only detects both single-bit and multibit errors, but it will actually correct single-bit errors.

Moreover, this ECC will detect (but not correct) errors of two, three, or even four bits. ECC protected memory systems handle these multibit errors much as parity checking handles single-bit errors: by generating a nonmaskable interrupt (NMI) that instructs the system to shut down to avoid data corruption.

Potential for system failures

Research has shown that the number of soft errors increases as memory capacity increases. Some percentage of these soft errors will be multibit errors that ECC cannot correct, so the potential for failure in ECC systems also increases as memory capacity increases. In fact, servers with 1 GB of memory using ECC are protected against memory failures only about as well as servers with 64 MB of memory using parity checking (Figure 1). With each new generation of servers, memory capacity increases, and so does the potential for system failures.

Figure 1: Server outages during a one-year period due to memory failures¹



Hot plug RAID memory

To help meet the availability and scalability demands of today's eBusiness world, HP developed a solution that allows customers to take advantage of industry-standard memory technology, increase server fault-tolerance, increase memory capacity, and increase server availability. Hot Plug RAID Memory provides a level of protection far greater than standard ECC-based solutions and allows the detection of otherwise undetectable errors (Table 1).

Table 1: Comparison of protection provided by parity checking, ECC, and Hot Plug RAID Memory

Error Condition	Parity	Standard ECC	RAID Memory
Single-bit	Detect	Correct	Correct
Double-bit	Unreliable	Detect	Correct
4-bit DRAM	Unreliable	Detect	Correct
8-bit DRAM	Unreliable	Unreliable	Correct
Greater than DRAM	Unreliable	Unreliable	Detect

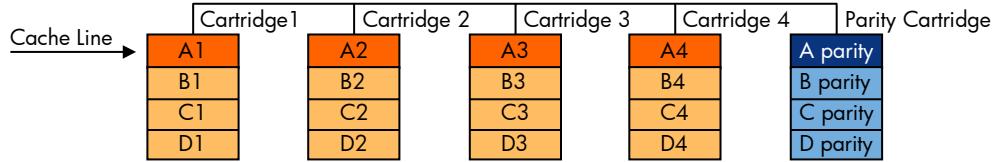
For years, the computer industry has used redundant array of independent disk (RAID) technology to provide fault tolerance and high availability for disk drive subsystems in servers. The technology used in Hot Plug RAID Memory is conceptually similar to RAID storage technology. However, in the context of the memory solution, RAID stands for redundant array of industry-standard DIMMs.

ProLiant servers with Hot Plug RAID Memory technology use five memory controllers to control five cartridges of industry-standard synchronous DRAM (SDRAM). When a memory controller needs to write data to memory, it splits a cache line of data into four blocks (shown as A, B, C, and D in

¹ Source: Timothy J. Dell, "A White Paper on the Benefits of Chipkill-Correct ECC for PC Server Main Memory," IBM Microelectronics Division – Rev. 11/19/97

Figure 2). Then each block is written, or striped, across four of the memory cartridges. RAID logic calculates parity information, which is stored on the fifth cartridge. With the four data cartridges and the parity cartridge, the data subsystem is redundant such that if the data from any DIMM is incorrect or if any cartridge is removed, the data can be recreated from the remaining four cartridges.

Figure 2: Data striping in Hot Plug RAID Memory



Hot Plug RAID Memory technology is implemented in ProLiant servers as part of a next-generation chipset designed by HP that includes four application-specific integrated circuits (ASICs). The ASICs enable the chipset to provide exceptional memory performance, a high-level of fault tolerance, and hot-plug memory capabilities. Hot Plug RAID Memory provides the ability for the memory subsystem to withstand a complete memory device failure and to continue operating normally.

Performance

Although Hot Plug RAID memory is conceptually similar to RAID technology in disk drive subsystems, there are some key performance and implementation differences between Hot Plug RAID Memory and typical storage subsystem RAID.

Hot Plug RAID Memory does not have the mechanical delays of seek time and rotational latency associated with hard disk drive arrays. Storage subsystem arrays use a single bus to write the stripes sequentially across multiple drives. In contrast, Hot Plug RAID Memory uses parallel, point-to-point connections to write data simultaneously across multiple memory cartridges.

Also, Hot Plug RAID Memory eliminates the write bottleneck associated with typical storage subsystem RAID implementations. In a storage array, the RAID controller generally performs a read operation of existing parity before a write operation can be completed. If a dedicated parity drive is being used, a bottleneck occurs. However, because Hot Plug RAID Memory almost always operates on an entire cache line of data, there is no need to read existing parity before a write operation. Therefore, no performance bottleneck occurs.

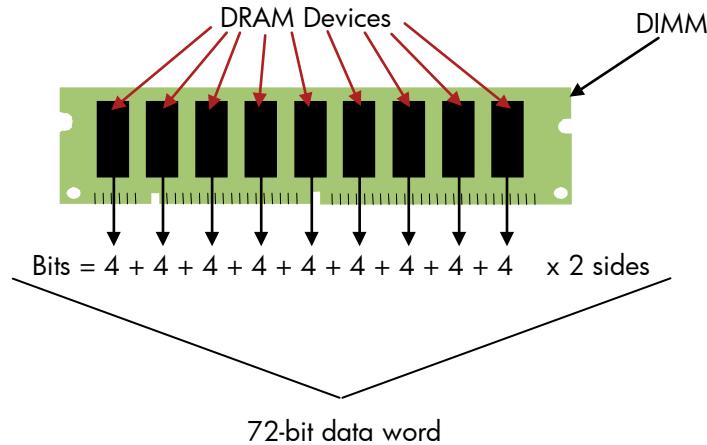
When a traditional striped RAID storage subsystem rebuilds data, data is not protected should another drive fail. However, Hot Plug RAID Memory operates in a typical (nonredundant) ECC mode while data is being rebuilt. As a result, even if a secondary memory failure occurs during a rebuild operation, the data is protected by ECC.

It is also important to note that like ECC memory protection, Hot Plug RAID Memory protection creates only minimal performance overhead. In Hot Plug RAID Memory, a RAID logic circuit calculates parity in parallel to the data flow, so error correction creates almost no additional data latency.

Basic operation

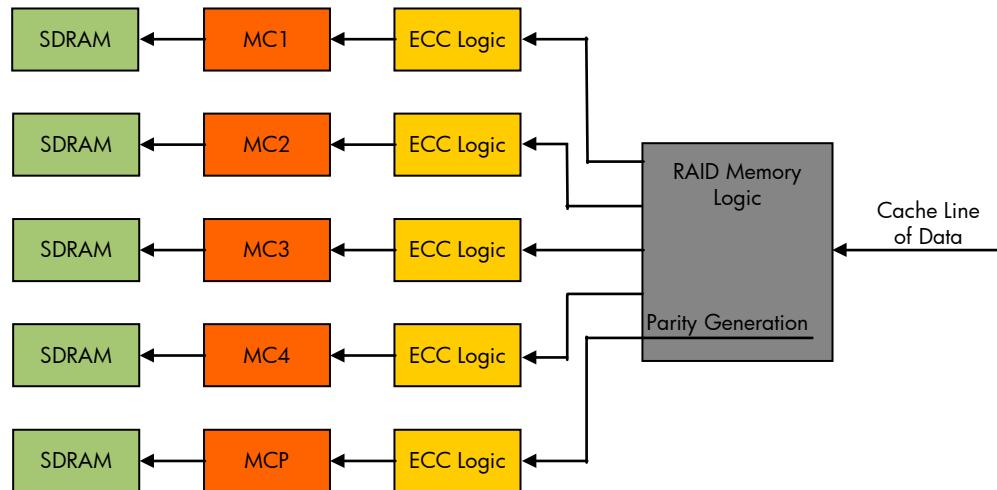
The operation of Hot Plug RAID Memory is dependant on the operation of processors, which use cache lines of data. A cache line of data is formed using data words from a group of DIMMs. In a memory transaction, a single access to a DIMM will access a number of bits from each DRAM device to create two 72-bit data words. For example, each of 18 devices provides 4 bits of data for each data word (Figure 3). Eight data words combine to form one cache line of data.

Figure 3: DIMM-level memory organization



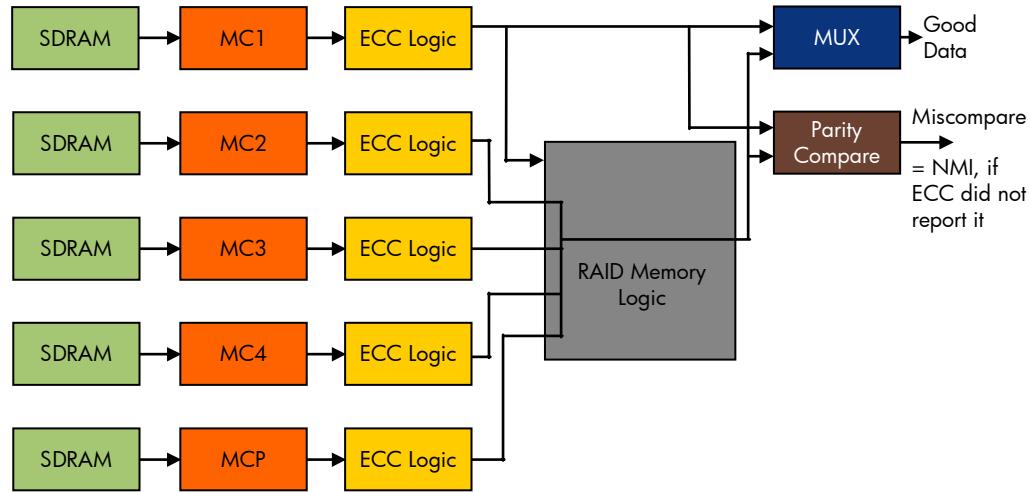
In a memory write transaction, parity is generated from the cache line of data. Simultaneously, the cache line of data is striped across four memory cartridges and the parity information is written to the fifth cartridge (Figure 4).

Figure 4: Diagram of a memory write transaction



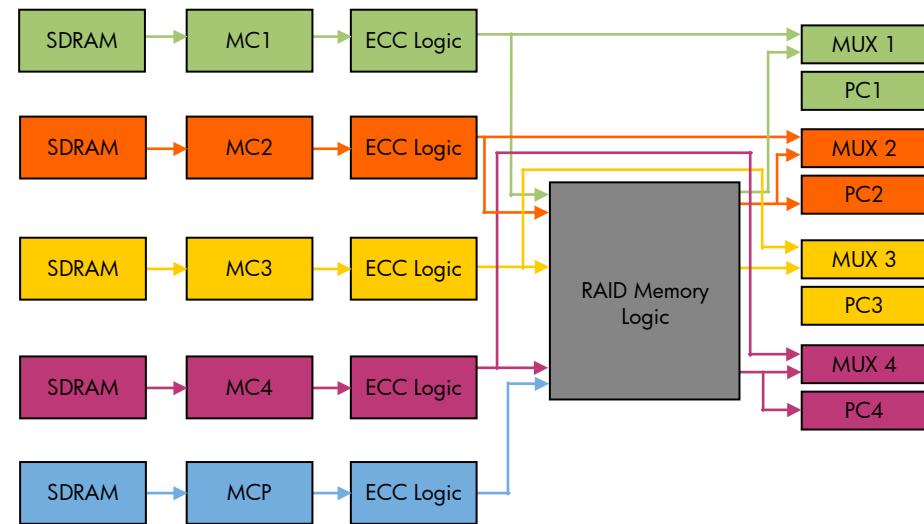
In a memory read transaction (Figure 5), each data word simultaneously travels through a separate memory controller to a separate ECC logic circuit that uses ECC code to detect errors. The ECC logic examines each data word and sends a signal identifying the data as good or bad to another logic device known as a MUX.

Figure 5: Diagram of a memory read transaction for one of the four data paths



During every read transaction, the ECC logic also passes data to a RAID memory logic circuit where a RAID algorithm simultaneously regenerates each data word using the data words from the other three memory controllers and the parity controller. For example, as shown in Figure 5, the RAID memory logic uses the data words from memory controllers 2, 3, 4, and P to regenerate the data word for memory controller 1 (MC1). Each regenerated data word from the RAID memory logic is then passed to a separate MUX (Figure 6).

Figure 6: RAID memory architecture



If the signal from the ECC logic to the MUX indicates the data is good, the MUX sends the original data to the processor. If the signal from the ECC logic to the MUX indicates the data has an error, the MUX sends the regenerated data from the RAID memory logic. At this point, the error detected by the ECC logic has been eliminated and only good data has been transmitted.

If the signal from the ECC logic to the MUX indicates that the data is good, a parity compare logic circuit (for example, PC1, PC2, PC3, or PC4 in Figure 6) compares the data from the ECC logic with

the regenerated data from the RAID memory logic. If all the data words in a read transaction are good, then the original data and the data from the RAID memory logic should be identical. If they are not, a data error undetectable by ECC has occurred. Such an occurrence, although rare, would result in bad data being passed along as if it were good.

However, with Hot Plug RAID Memory, the parity compare fails in such a situation and initiates an NMI, preventing the transmission of corrupt data. This feature makes Hot Plug RAID Memory virtually immune to data corruption.

Hot-plug capabilities

The redundancy in Hot Plug RAID Memory provides the ability to hot plug memory cartridges without bringing down the server. This gives unprecedented levels of memory availability and scalability within industry-standard servers. Hot Plug RAID Memory enables the following abilities while the system is running:

- Hot replace: replacing a failed DIMM
- Hot add: adding a DIMM to a memory cartridge
- Hot upgrade: replacing a set of DIMMs with different (higher capacity) ones

Hot-replace capability is offered in a driverless implementation that requires no support from the operating system. ProLiant servers with Hot Plug RAID Memory have hot-replace capability directly out of the box, regardless of the operating system used. This operating system independence was achieved using System Management Mode (SMM), a mode of Intel processors. Use of SMM eliminated the need for HP engineers to develop driver software for every OS and removed the maintenance associated with those drivers.

When an administrator initiates a hot-replace operation, the memory controller tells the server to ignore the cartridge of memory where the hot-replace operation will be performed. Until the hot-replace operation is completed, memory transactions use the other four memory cartridges protected by ECC. Thus, the memory subsystem operates in a nonredundant mode like today's ECC memory subsystems. At this point the cartridge containing the DIMM to be replaced can be removed from the system. The failed DIMM can then be replaced in that cartridge and the cartridge can be inserted into the system. Once the memory cartridge is back online, full redundancy is restored.

When a cartridge is inserted back into the system, Hot Plug RAID Memory automatically rebuilds the data across all the memory cartridges. Rebuilding data can degrade memory performance briefly, but a rebuild for 4 GB of memory takes about 30 seconds—a small price to pay to avoid downtime while increasing fault tolerance.

After the RAID logic rebuilds the data, a verify procedure confirms that the rebuild operation was successful. During a verify procedure, every address location in memory is read. Errors found are reported to the system. If the verify procedure does not confirm that the rebuild operation was successful, the memory will not be brought online until the problem is corrected. The verify command can also be initiated independently of a hot-plug procedure. For example, an administrator can set up a routine that will run the verify procedure periodically and report any errors before they cause problems. This type of proactive monitoring program further reduces downtime.

Hot-add capability allows a user to scale up a computer system as needed by adding extra DIMMs. Hot-add capability requires support from the operating system to recognize the additional memory. HP worked with operating system vendors to ensure that this capability is supported in current and future releases.

Ease-of-use capabilities

Hot Plug RAID Memory also enables several ease-of-use features. The registers and logic in Hot Plug RAID Memory permit software to take action when certain situations arise. For example, a register can collect information on memory errors, and software can be programmed to direct the system to issue warnings and initiate changes in status indicators. Light-emitting diodes (LEDs), locks, and alarms can be used to indicate good or bad DIMMs and to make management of Hot Plug RAID Memory quite easy and intuitive.

Conclusion

Memory error detection and correction technology has not evolved as rapidly as other technologies used in today's enterprise servers. While ECC provides good detection and single-bit correction capabilities, today's systems with more than 1 GB of memory require additional fault-tolerant memory technology to provide a consistent level of protection. Hot Plug RAID Memory technology answers the need for additional data protection. Using traditional RAID technology implemented at the chipset level, Hot Plug RAID Memory provides unprecedented levels of protection while increasing the availability and scalability of the memory subsystem. Because the Hot Plug RAID Memory solution uses industry-standard DIMMs, it provides a fault-tolerant memory form factor that is easily obtainable at competitive prices.

Call to action

To help us better understand and meet your needs for ISS technology information, please send comments about this paper to: TechCom@HP.com.

© 2001, 2003 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.