# HP Advanced Memory Protection technologies

technology brief, 3rd edition

# Abstract

This technology brief explains the causes of system memory errors and describes the methods that HP uses to prevent "hard" memory errors and to detect and correct inevitable "soft" errors. This document then describes the HP Advanced Memory Protection technologies that go beyond error correction to increase the fault tolerance of HP ProLiant ML and DL 300- and 500-series servers.

# Introduction

Businesses are increasingly dependent on industry-standard servers to run memory-intensive and mission-critical applications. This trend is driving operating systems to support more memory and pushing the memory capacity of servers to new levels. System memory has become more reliable over the years because of better manufacturing processes and memory protection technologies like error checking and correcting or error correcting code (ECC), first introduced in industry-standard servers by HP. However, as the density of memory components and server memory capacity continue to increase, there is a higher probability of memory errors occurring. Memory errors can corrupt data and cause servers to crash, resulting in the permanent loss of business data and lost revenue from downtime.

To meet this memory reliability challenge, HP offers three levels of Advanced Memory Protection that provide increased fault tolerance for applications requiring higher levels of availability. HP customers can choose a system with the level of memory protection they prefer—Online Spare Memory, Hot Plug Mirrored Memory, and HP Hot Plug RAID Memory. In reference to memory, *RAID* stands for Redundant Array of Industry-standard DIMMs.

First, this paper explains the causes of memory errors and why the possibility of memory errors is increasing in servers. Then it describes the functions and limitations of established methods to prevent and detect most memory errors. Lastly, it summarizes the fault-tolerant Advanced Memory Protection technologies included in the HP ProLiant ML and DL 300- and 500-series servers. For information about Advanced Memory Protection capabilities of specific ProLiant servers, refer to the server's user guide.

# Memory errors

Memory modules used in servers are electronic storage devices; therefore, they are inherently susceptible to memory errors. Computers use two types of memory devices—static random access memory (SRAM) and dynamic RAM (DRAM). SRAM is used for cache memory because it is very fast and retains its data until the power is turned off. DRAM chips, which are used for main memory, are installed on 168-pin dual inline memory modules (DIMMs). Each DRAM chip stores data in columns and rows of capacitors (memory cells) that must be continuously recharged, or refreshed, to preserve the data. A charged capacitor represents a "1" data bit and an uncharged capacitor represents a "0" data bit. The level of the electrical charge is determined by the operating voltage of the memory device.

When a memory cell is accessed during a read operation, the capacitor's charge determines whether it is read as a "1" or "0" data bit. For example, in a 1.8-volt system, a sensor reads a capacitor with a charge of +1.8 V as a "1" bit and a sensor reads a capacitor discharged to 0 V as a "0" bit. As long as the voltage is closer to +1.8 V than to 0 V, the sensor will read the value correctly. However, if a capacitor's charge is affected by some external event, the data may become incorrect. For servers running business-critical applications, these memory errors can cause servers to crash and can result in the permanent loss of business data. Memory errors are classified by the number of bits that are affected—single-bit or multi-bit—and the cause of error.

## Single-bit and multi-bit errors

The memory bus is a circuit that consists of two parts: the data bus and the address bus. The data bus is a set of traces that carry the actual data to and from SDRAM. Each trace carries one data bit at a given time. Today's computers have a 64-bit wide data bus, which means that the bus transports 64 bits at a time. These 64 bits constitute an ECC data word. An error in one bit of a data word is called a single-bit error. An error in more than one bit of a data word is called a multi-bit error.

## Hard errors and soft errors

Memory errors are referred to as either hard errors or soft errors, depending on their cause. A hard error is caused by a broken or defective piece of hardware, so the device consistently returns incorrect results. For example, a memory cell may be stuck so that it always returns "0" bit, even when a "1" bit is written to it. Hard errors are caused by DRAM defects, bad solder joints, connector issues, etc. Soft errors are more prevalent. They occur randomly when an electrical disturbance near a memory cell alters the charge on the capacitor. A soft error does not indicate a problem with a memory device because once the stored data is corrected (for example, by a write to a memory cell) the same error does not recur.

## Increasing possibility of memory errors

Two trends are resulting in the increased likelihood of memory errors: the memory capacity of servers is expanding and the storage density of memory components is increasing.

**Effect of memory capacity**

Software vendors are developing increasingly complex and memory-intensive applications to meet the demands of the financial, telecommunications, and entertainment industries. This is driving operating systems to address more memory, thus causing manufacturers to expand the memory capacity of servers. For example, some HP ProLiant servers now support up to 128 GB of memory (Figure 1). As manufacturers expand the memory capacity of servers, the possibility of memory errors likewise increases.

**Figure 1.** The memory capacity of HP ProLiant servers is growing to meet application demands.

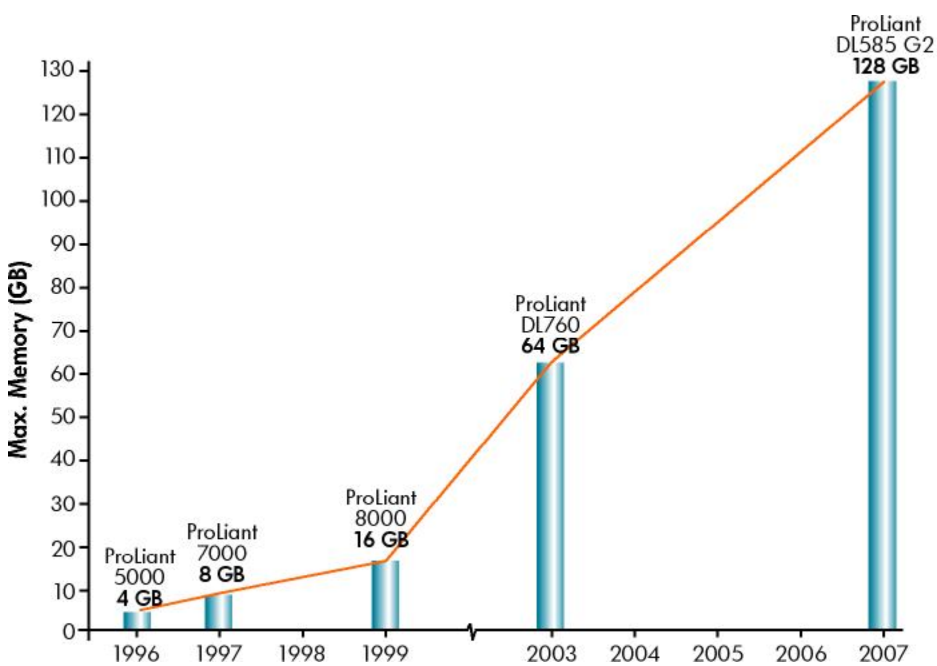

### Effect of SDRAM storage density and operating voltage

Two parameters of DRAM are inextricably tied together—the storage density of the DRAM chips and the operating voltage of the memory system. As the size of memory cells decreases, the DRAM storage density increases along with the memory cells' sensitivity to voltage. Initially, industry-standard DIMMs operated at 5 volts. However, because of improvements in DRAM storage density, the operating voltage was decreased to 3.3 V, 2.5 V, and then 1.8 V to allow memory to run faster and consume less power. Because memory storage density is increasing and operating voltage is shrinking, there is a higher probability that an error may occur. Whenever a data bit is misinterpreted and not corrected, the error can cause an application to crash.

# Methods to prevent memory errors

There are two ways to protect against memory errors: testing and the use of error detection/correction technologies. HP has long established its leadership in the qualification and testing of memory components and backs its testing procedures with a pre-failure warranty.

## Memory testing

As memory chips become faster and more complex, testing them becomes more difficult and expensive. Memory device manufacturers invest heavily in testing systems, and they continually revamp their testing procedures to maintain device quality. Due to the constant changes in manufacturing processes, HP qualifies each memory module design and manufacturing process to minimize the occurrence of hard errors. In addition to the rigorous qualification processes of module manufacturers, HP further tests every memory module in the model of ProLiant server in which it will be installed. This process includes testing each manufacturer's modules on every model of HP ProLiant server currently shipping and re-qualifying every module manufacturer each time HP offers a new processor speed or a new server platform.[1] This testing and re-qualification process results in continuous improvement of memory module reliability.

Superior qualification and testing procedures allow HP to offer a three-year Pre-failure Warranty on HP memory. The HP Pre-failure Warranty allows for replacement of any HP DIMM that exceeds predefined limits for correctable errors. These errors are recorded by the server and can be verified through HP Insight Manager or a diagnostics program.

## HP Pre-failure Warranty

The Pre-failure Warranty, which is standard on all HP ProLiant servers, extends the advantage of an HP three-year, limited warranty on critical components, such as memory, before they actually fail. Specifically, the Pre-failure Warranty ensures that when customers receive notification from HP Systems Insight Manager that a critical server component may fail, the component is replaced free of charge under the warranty. With the Pre-failure Warranty, system administrators can proactively schedule downtime for maintenance and not interrupt critical business operations that rely on these enterprise servers.

During the warranty period, the Pre-failure Warranty covers the replacement of DIMMs used in a server’s main memory when the predefined thresholds for correctable errors have been exceeded. The predefined thresholds can differ among system architectures.

Non-repeating, correctable soft errors are not covered under warranty since their occurrence requires no action.

[1] For more information, download the white paper *Why Buy HP Memory?*, at http://h71028.www7.hp.com/ERC/downloads/4AA0-4216ENW.pdf.

## Error detection/correction technologies

HP minimizes the occurrence of most manufacturing-related memory errors by certifying memory manufacturers and continuously testing their memory products. However, soft errors cannot be prevented by these means. As memory manufacturers increase the capacity of memory devices, the likelihood of soft memory errors likewise increases. 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.

### ECC memory

In 1993, HP was the first company to introduce ECC memory in industry-standard servers. This significantly reduced the probability of fatal memory failures. ECC memory is now standard in all HP ProLiant servers. ECC provides added protection over parity checking. Parity checking provides single-bit error detection, but it does not correct memory errors or handle multi-bit errors. ECC detects both single-bit and multi-bit errors in a 64-bit data word, and it corrects single-bit errors.

ECC encodes information in a block of 8 bits to permit the recovery of a single-bit error. Every time data is written to memory, ECC uses a special algorithm to generate values called check bits. The algorithm adds the check bits together to calculate a checksum, which it stores with the data. When data is read from memory, the algorithm recalculates the checksum and compares it with the checksum of the written data. If the checksums are equal, then the data is valid and operation continues. If they are different, the data has an error and the ECC memory logic isolates the error and reports it to the system. In the case of a single-bit error, the ECC memory logic can correct the error and output the corrected data so that the system continues to operate (Figure 2).

**Figure 2.** Diagram of ECC logic locating and correcting a single-bit error

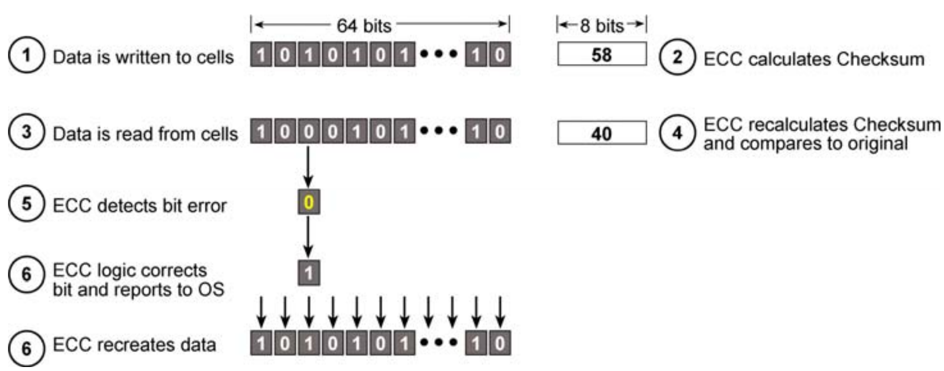


In addition to detecting and correcting single-bit errors, ECC will detect (but not correct) errors of two random bits and up to four bits within a single DRAM chip. ECC memory handles these multi-bit errors by generating a non-maskable interrupt (NMI) that instructs the system to halt to avoid data corruption.

ECC technology has provided adequate protection for many applications. However, the effectiveness of ECC protection decreases as memory capacity rises. This fact is significant because of the following factors driving industry-standard servers to support more memory capacity:

- operating systems now support greater amounts of memory;
- low-cost, high-capacity memory modules are more available;
- and server virtualization.

### Advanced ECC technology

To improve memory protection beyond standard ECC, HP introduced Advanced ECC technology in 1996. HP and most other server manufacturers continue to use this solution in industry-standard products. Advanced ECC can correct a multi-bit error that occurs within one DRAM chip; thus, it can correct a complete DRAM chip failure. In Advanced ECC with 4-bit (x4) memory devices, each chip contributes four bits of data to the data word. The four bits from each chip are distributed across four ECC devices (one bit per ECC device), so that an error in one chip could produce up to four separate single-bit errors. Figure 4 shows how one ECC device receives four data bits from four DRAM chips.

**Figure 4.** In Advanced ECC, each DRAM chip sends four data bits, which are distributed across four ECC devices.

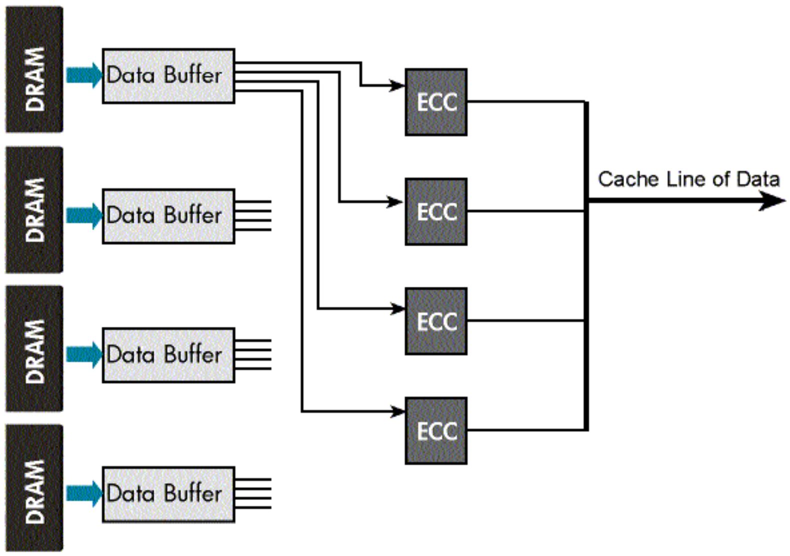


Since each ECC device can correct single-bit errors, Advanced ECC can actually correct a multi-bit error that occurs within one DRAM chip. As a result, Advanced ECC provides device failure protection (Table 1).

**Table 1.** Comparison of error protection with ECC and Advanced ECC

| Error Condition | ECC Outcome | Advanced ECC Outcome |
|---|---|---|
| Single-bit | Correct | Correct |
| Double-bit | Detect | Correct or detect |
| DRAM failure | Detect | Correct |
| ECC detection fault | — | — |

Although Advanced ECC provides failure protection, it can reliably correct multi-bit errors only when they occur within a single DRAM chip. Advanced ECC does not provide failover capability. As a result, if there is a memory failure, the system must be shut down before the memory can be replaced. The latest generation of HP ProLiant servers offers three levels of Advanced Memory Protection that provide increased fault tolerance for applications requiring higher levels of availability.

# HP Advanced Memory Protection technologies

In addition to ECC and Advanced ECC, HP ProLiant 300- and 500-series servers feature one or more Advanced Memory Protection technologies that deliver increased fault tolerance for applications requiring higher levels of availability. These Advanced Memory Protection technologies—Online Spare Memory, Hot Plug Mirrored Memory, and Hot Plug RAID Memory—are optimized for the features and applications of each server series (see Figure 5).

**Figure 5.** Advanced ECC and HP Advanced Memory Protection technologies in the latest generation of HP ProLiant servers

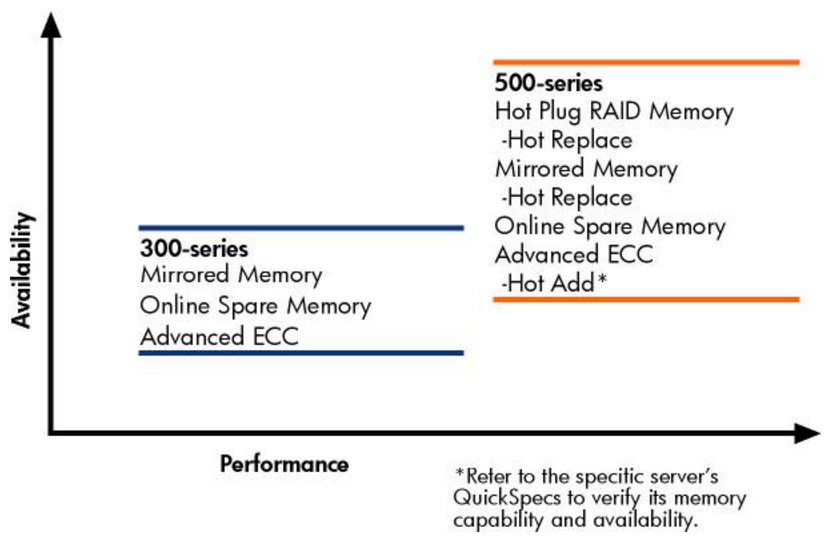


The implementation of each hot-plug operation—hot-add, hot-replace, and hot-upgrade—varies depending on the server series. For ProLiant 300- and 500-series servers, hot-add refers to adding memory boards or cartridges to the system while it is running. The hot-add feature allows customers to increase the amount of memory available to the OS without rebooting the system. Hot-add must be enabled in the ROM-Based Setup Utility (RBSU) and must be supported by the OS. Hot-replace allows customers to remove a memory board (or cartridge), replace a failed or degraded DIMM, and re-install the memory board; all while the server is running. Hot-replace is available without any OS support and can be used with either memory mirroring or RAID techniques. Hot-upgrade allows the customer to replace smaller capacity DIMMs with larger capacity DIMMs while the system is running.

## Online Spare memory

Online Spare memory mode is a higher level of memory protection that complements Standard Memory mode with Advanced ECC. With Online Spare mode, a DIMM with a rank of memory at least as large as the other ranks in the system, or memory board, is designated as the Online Spare rank.[2] If one of the other DIMMs exceeds a threshold rate of correctable memory errors, the affected rank of memory within that DIMM is taken offline and the data is copied to the Online Spare rank. This capability maintains server availability and memory reliability without service intervention or server interruption. The DIMM that exceeded the error threshold can be replaced at the customer's convenience during a scheduled shutdown. Online Spare reduces the chance of an uncorrectable error bringing down the system; however, it does not fully protect the system against uncorrectable memory errors.

Online Spare mode is beneficial to businesses with sites that do not have IT staff available to service a failure, do not always have replacement memory readily available, or cannot bring down the server until a scheduled shutdown. Online Spare mode does not support any hot-plug operations; therefore, the server must be powered down to replace a bad memory module. The cost of implementation for Online Spare over Advanced ECC is the hardware cost of the extra DIMMs for the spare memory bank.

The user configures the system for Online Spare mode via the RBSU[3]. Online Spare configuration requirements vary for each server series. Online Spare mode does not require operating system support or special software beyond the System BIOS. However, if messaging and logging are desired at the console along with messages in HP Systems Insight Manager, an operating system must be used that has system management and agent support for Advanced Memory Protection.

For information about whether a specific ProLiant server supports Online Spare memory, refer to the QuickSpecs for the server.

---

[2] For more information about memory ranks, refer to the technology brief "Memory technology evolution: an overview of system memory technologies" edition at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00256987/c00256987.pdf .

[3] Please read the user guide for each server for instructions to run RBSU.

## Mirrored Memory

Mirrored Memory mode is a fault-tolerant memory option that provides a higher level of availability than Online Spare mode. While Online Spare mode protects against single-bit errors and entire DRAM failure, Mirrored Memory mode provides full protection against single-bit and multi-bit errors.

There are two types of Mirrored Memory mode: non-hot plug and hot plug. Non-hot plug Mirrored Memory mode is beneficial to sites that cannot afford unscheduled downtime. Hot-plug Mirrored Memory is beneficial to sites that cannot risk waiting until scheduled downtime to replace degraded memory modules. No operating system support is required for this option and all software and drivers are in the system BIOS.

For information about whether a specific ProLiant server supports non-hot plug or hot plug Mirrored Memory, refer to the QuickSpecs for the server.

### Non-hot plug Mirrored Memory mode

ProLiant servers that support non-hot plug Mirrored Memory do not offer hot-replace capability. This option provides customers with full protection against single-bit and multi-bit errors. To enable non-hot plug Mirrored Memory mode, customers designate half of the memory banks as system memory and the remaining banks as mirrored memory. All banks must be configured identically.

In non-hot plug Mirrored Memory mode, the same data is written to both the system memory and mirrored memory banks, but data is read only from the system memory banks. If a DIMM in the system memory banks experiences a multi-bit error or reaches the pre-defined error threshold for single-bit errors, the system will still write data to both the system and mirrored memory banks. However, the system will only read data from the mirrored memory banks. This allows continuous operation and maintains the level of server availability except in the highly unlikely case of a simultaneous error in exactly the same location on a DIMM and its mirrored DIMM.

### Hot Plug Mirrored Memory mode

Hot Plug Mirrored Memory mode provides hot-replace capability to increase server availability. Hot-replace allows customers to remove a memory board, replace a failed or degraded DIMM, and re-install the memory board; all while the server is running. This capability is especially useful for businesses that cannot afford downtime and cannot risk waiting until scheduled downtime to replace degraded memory modules.

Hot Plug Mirrored Memory mode can be enabled using two or four memory boards. The memory boards must be configured identically. The customer designates half of the memory boards as system memory and the remaining boards as mirrored memory. If a DIMM on one of the system memory boards experiences a multi-bit error or reaches the error threshold for single-bit errors, the data is read from the corresponding mirrored DIMM. This enables the customer to hot-replace the failed DIMMs without shutting down the server. Hot Plug Mirrored Memory along with Advanced ECC provides protection against all memory errors except in the highly unlikely case of a simultaneous error in exactly the same location on a DIMM and its mirrored DIMM.

## Hot Plug RAID Memory

Hot Plug RAID Memory protects the server against uncorrectable memory errors that would otherwise result in a server failure. Hot Plug RAID Memory allows the memory subsystem to operate continuously, even in the event of a complete memory device failure. For information about whether a specific ProLiant server supports Hot Plug RAID memory, refer to the QuickSpecs for the server.

### Operation

The Hot Plug RAID memory implementation in ProLiant servers requires four memory boards. The implementation is conceptually similar to Redundant Array of Independent Disks (RAID) Level 5 in that the north bridge uses an exclusive-OR engine to generate a parity check line for every three cache lines. The north bridge interleaves the cache lines and the parity check line across all four memory boards (Figure 6). Because the cache lines are striped across the memory boards, all four boards must have the same total amount of memory. If an uncorrectable memory error is encountered, the server can re-create the proper data using the parity information and the information from the other memory boards that contain no failed DIMMs.

**Figure 6.** Hot Plug RAID memory stripes three cache lines and the associated parity information across the memory boards.

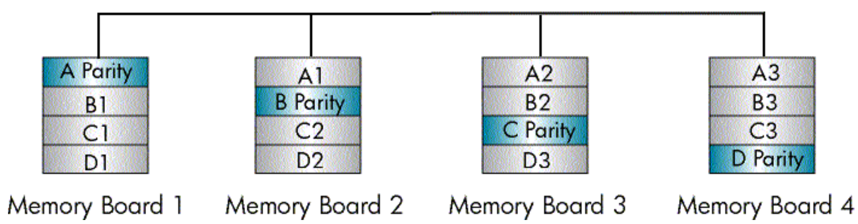


Hot Plug RAID Memory should not be confused with the schemes used for RAID storage disk drive arrays. Hot Plug RAID Memory does not have the mechanical delays of seek time and rotational latency associated with disk drive arrays. Storage subsystem arrays use a single bus to write the stripes sequentially 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 usually 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.

### Hot-plug capabilities

The redundancy in Hot Plug RAID Memory allows customers to hot-replace, hot-add, and hot-upgrade DIMMs without shutting down the server. Hot replace involves replacing a failed DIMM while the system continues to operate. Hot Plug RAID memory offers hot-replace capability in a driverless implementation.

When a hot-plug operation is completed, HP Hot Plug RAID Memory automatically rebuilds the data across all the memory cartridges. The process required to rebuild 4 GB of memory takes less than 30 seconds, which is a small price to pay for increased fault tolerance and no downtime due to system memory failure.

## Interline comparison

Table 2 provides an interline comparison and competitive advantage of HP Advanced Memory Protection technologies to Advanced ECC. All technologies shown provide device failure protection. HP Hot Plug Mirrored Memory will provide hot-replace capability, while HP Hot Plug RAID Memory will provide hot-replace, hot-add, and hot-upgrade capabilities.

**Table 2.** Advanced Memory Protection interline comparison and competitive advantage

| | **HP Advanced ECC Technology** | **HP Online Spare Memory** | **HP Non-Hot Plug Mirrored Memory** | **HP Hot Plug Mirrored Memory** | **HP Hot Plug RAID Memory** |
|---|---|---|---|---|---|
| Device failure protection | Yes | Yes | Yes | Yes | Yes |
| Industry-standard DIMMs | Yes | Yes | Yes | Yes | Yes |
| Hot Plug | No | No | No | Yes | Yes |
| Failed DIMM replacement | Offline | Offline | Offline | Online | Online |
| Additional memory expense | 0% | 10% – 50% | 100% | 100% | 25% |

## Conclusion

The demand for servers with more memory capacity is unrelenting—it is driven by increasingly complex and memory-intensive applications and more powerful processors. While meeting the demand for more system memory, the challenge for server manufacturers is to maintain the reliability of the memory system, even though there is a higher probability of memory errors as memory densities and capacities climb.

HP is meeting the challenge with fault-tolerant memory protection technologies: Online Spare Memory, Hot Plug Mirrored Memory, and Hot Plug RAID Memory. Online Spare Memory is beneficial to customers with sites that cannot afford downtime from memory errors, yet can wait until a scheduled downtime to replace failed memory modules. Hot Plug Mirrored Memory provides a more fault-tolerant option for sites that cannot afford downtime from memory errors and do not want to wait until scheduled downtime to replace failed memory modules. It allows memory modules to be hot-replaced without shutting down the server.

HP Hot Plug RAID Memory provides the highest level of availability for customers who deploy industry-standard servers with large memory systems to run 24x7 applications. It enables the memory subsystem to operate continuously, even in the event of a complete memory device failure, by allowing DIMMs to be hot-replaced, hot-added, and hot-upgraded.

These HP Advanced Memory Protection technologies enable customers to choose a system with the level of memory availability they prefer to enhance the robustness of their final solution.

# For more information

For additional information, refer to the resources detailed below.

| Resource description | Web address |
|---|---|
| "Memory technology evolution: an overview of system memory technologies" technology brief | http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00256987/c00256987.pdf |
| "Technologies for the ProLiant ML570 G4 and ProLiant DL580 G4 servers" technology brief | http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00795607/c00795607.pdf |
| Industry Standard Server Technology Papers | www.hp.com/servers/technology/ |
| HP ProLiant DL380 G5 QuickSpecs (North America) | http://h18004.www1.hp.com/products/quickspecs/12477_na/12477_na.html |
| HP ProLiant DL360 G5 QuickSpecs (North America) | http://h18004.www1.hp.com/products/quickspecs/12476_na/12476_na.html |
| HP ProLiant ML370 G5 QuickSpecs (North America) | http://h18004.www1.hp.com/products/quickspecs/12478_na/12478_na.html |

# Call to action

Send questions and comments about this paper to: TechCom@HP.com.



TC070803TB, 8/2007