



Memory RAS technologies for HPE ProLiant/Synergy/Blade Gen10 servers with Intel Xeon Scalable processors

Contents

Introduction.....	2
Why memory RAS is needed.....	2
Memory RAS technologies in HPE ProLiant/Synergy/Blade servers.....	3
Fast Fault Tolerance.....	3
Advanced ECC support.....	4
Online spare with Advanced ECC support.....	4
Mirrored memory with advanced ECC support.....	5
Memory scrubbing (Patrol and Demand).....	7
Conclusion.....	8



View our **DL580 Gen10**
inventory [here](#)

Introduction

Memory device failures—if not corrected—can result in service events or even server crashes. As modern servers implement ever-larger memory arrays, the likelihood of a memory device failure increases with higher memory capacity. Since memory device failures are some of the most frequent types of failure for servers besides storage failures, HPE ProLiant Gen10 servers using Intel® Xeon® Scalable processors provide an increasingly comprehensive suite of memory RAS (reliability, availability, and serviceability) features split into the following categories:

- Error detection and correction
- Redundancy and resiliency
- Maintenance

This paper provides a quick overview of selected memory RAS technologies for HPE ProLiant Gen10 servers, their characteristics, minimum requirements, and how to enable them. The information presented will help you select the most appropriate memory RAS technologies to meet your demanding workload and data center service level requirements, especially for business-critical workloads.

Note

This paper focuses solely on server memory RAS features. It does not review the comprehensive suite of other RAS technologies found throughout the ProLiant/Synergy/Blade portfolio.

Why memory RAS is needed

Server uptime is still one of the most critical aspects of data center maintenance. Unfortunately, servers can run into trouble from time to time due to software issues, power outages, or memory errors. The three major categories of memory errors we track and manage include correctable errors, uncorrectable errors, and recoverable errors. The determination of which errors are correctable and uncorrectable is completely dependent on the capability of the memory controller.

Correctable errors are, by definition, errors that can be detected and corrected by the chipset. Correctable errors are generally single-bit errors. All HPE servers are capable of detecting and correcting single-bit errors and with advanced error-correcting code (ECC) support. On HPE systems, the user is warned about a DIMM exceeding the correctable error threshold (maximum amount of correctable errors tolerated in a certain time window) either through lights on the front panel or system board (if available), or the HPE Integrated Management Log (IML).

Uncorrectable errors are errors that can be detected but not corrected by the chipset. These are always multi-bit memory errors. The error will be logged in the IML. Uncorrectable errors can typically be isolated down to a single DIMM. Uncorrectable errors will usually immediately result in a system crash or shutdown. In some cases, with operating system (OS) support and advanced SKU processors (Intel Xeon Platinum and Gold processors), uncorrectable errors do not result in a system crash. We call these recoverable errors. For error recovery details, consult OS documentation.

DRAM errors generally come in two different types—hard errors and soft errors.

- **Hard errors** typically indicate a problem with the DIMM itself. Although hard correctable errors are corrected by the system and will not result in system downtime or data corruption, they still indicate a hardware problem. Hard errors will typically cause a DIMM to exceed HPE systems' correctable error threshold. The user is warned about those errors.
- **Soft errors** do not indicate any issues with the DIMM. They occur when the data and/or ECC bits on the DIMM are incorrect, but the error will not continue to occur once the data and/or ECC bits on the DIMM have been corrected. Soft errors will not typically cause a DIMM to exceed HPE systems' correctable error threshold and therefore, no indication of a hardware issue is shown.

Any kind of error, if not handled correctly, can eventually cause a system shutdown. In the early days of servers, basic ECC was sufficient to resolve most DRAM failures. However, today's servers present a completely different challenge, so additional RAS features are necessary to maintain expected server stability and uptime. It is important to note that by avoiding a critical failure, a system crash can be avoided. Failed memory devices are replaced as part of periodic service. Also, memory RAS technologies can detect a DRAM device on a DIMM that has had numerous soft errors, and recommend replacing it before it has a hard failure.



Memory RAS technologies in HPE ProLiant/Synergy/Blade servers

The following descriptions provide an overview of the functionality of selected memory RAS technologies.

Fast Fault Tolerance

Overview

Fast Fault Tolerance is a new HPE Memory RAS feature first introduced in [HPE Gen10 servers](#) with Intel Xeon Scalable processors. Those servers configured with HPE SmartMemory and Fast Fault Tolerance offer an extra layer of protection against planned server downtime and server crashes. Fast Fault Tolerance, an enhanced version of adaptive double device data correction (ADDDC), is a result of a joint Intel® and Hewlett Packard Enterprise Collaboration. Fast Fault Tolerance has more spare regions (part of memory allocated only for replacing bad memory areas) and more options to map out bad sections of memory. This results in significantly better memory reliability and availability than what the rest of the industry can provide using ADDDC only.

Characteristics

In the past server generations, the most advanced memory protection technology in ProLiant servers was double device data correction (DDDC). The biggest issue with this was that it had to be enabled at boot and it significantly reduced memory throughput when enabled. Customers had to choose between resiliency and performance. Fast Fault Tolerance provides significant improvement over DDDC because it incorporates the performance benefits of single device data correction (SDDC) with the availability of DDDC. Fast Fault Tolerance allows the system to boot with full-memory performance and only puts small sections (banks) of memory into lockstep when needed to correct failures resulting in a significantly better performance than DDDC. When the failing section is larger than a bank, a larger negative impact on performance may be observed.

Summary:

- Fast Fault Tolerance survives up to two DRAM failures (detect and correct).
- The RAS feature combines the resiliency of DDDC with the performance of SDDC.

Minimum requirements

There must be a minimum of two ranks on each populated channel. Furthermore, only HPE SmartMemory in a x4 organization can be used.

How to enable Fast Fault Tolerance

Fast Fault Tolerance is enabled by default when the “mission-critical” profile is selected in the ROM-based setup utility (RBSU). Fast Fault Tolerance can also be enabled on Gen10 servers at the time of purchase (refer to platform QuickSpecs), but otherwise is disabled by default in all other workload profiles. If a server has Fast Fault Tolerance disabled, then the user can configure the system for Fast Fault Tolerance mode through the RBSU. Fast Fault Tolerance configuration requirements vary for each server series, but it does not require OS support or special software beyond the basic input/output system (BIOS).




Figure 1. Fast Fault Tolerance is enabled in the RBSU

Technical details

Currently, Fast Fault Tolerance does require that the server runs in “closed-page” mode and some workloads will see a reduction in throughput. Closed-page mode is not expected to have a significant performance loss for random access memory patterns (e.g., SQL or other databases), but there will be a performance loss for sequential access memory patterns (e.g., data streams).



There will also be a minimal performance reduction in throughput if a DRAM fails but only in the typically very small region (most common size is a bank) of memory that is affected. No significant loss is expected for random-access memory patterns because the region in lockstep will be accessed infrequently. The loss can be significant if you have rank level virtual lockstep or if an application accesses the region frequently until the DIMM is replaced. The overall reduction in throughput from Fast Fault Tolerance is expected to be minimal for the vast majority of customers but does depend on the application, the size of the affected region, and the memory configuration.

Advanced ECC support

Overview

Advanced ECC memory is the default memory protection mode for HPE servers. Standard ECC can correct single-bit memory errors and detect multi-bit memory errors. When multi-bit errors are detected using standard ECC, the error is signaled to the server and causes the server to halt.

Advanced ECC has been the default error correction scheme in HPE servers for over two decades. It not only protects servers against single-bit errors, it also protects against some multi-bit memory errors—specifically those that occur within a single DRAM chip.

Advanced ECC can correct both single-bit memory errors and 4-bit memory errors if all failed bits are on the same DRAM device on the DIMM. Advanced ECC provides more protection than standard ECC because it is possible to correct certain memory errors that would otherwise be uncorrected and result in a server failure. Using HPE advanced memory error detection technology, the server provides notification when a DIMM is degrading and has a higher probability of an uncorrectable memory error.

Minimum requirements

There are no specific memory population rules or RBSU settings required for advanced ECC support. It's enabled as the default on Purley platforms.

How to enable advanced ECC support

Advanced ECC support is the advanced memory protection mode default in the RBSU >> Memory Options.




Figure 2. Advanced ECC support is a RBSU default feature

Technical details

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 HPE ProLiant/Synergy/Blade servers using Intel Xeon Scalable processors offers three levels of advanced memory protection (including Fast Fault Tolerance) that provide increased fault tolerance for applications requiring higher levels of availability.

Online spare with Advanced ECC support

Overview

Online sparing provides protection against persistent DRAM failure. It tracks excessive numbers of correctable errors and copies the contents of an unhealthy rank to an available spare rank in advance of multi-bit or persistent single-bit failures that may result in future uncorrectable faults. It does not identify or disable individual failed DRAMs, but instead it disables the DIMM rank. Since a DIMM rank is needed to perform sparing, this technique reduces the total amount of available memory by the amount of memory used for sparing. Sparing can only handle one failure per channel. Ranks within a DIMM that are likely to receive a fatal/uncorrectable memory error are automatically removed from operation, resulting in less system downtime. DIMM sparing and rank sparing are not compatible with mirroring.



Characteristics

Technically, a whole DIMM is not required for sparing unless a single rank DIMM is used. Therefore, the cost of implementation for online spare over advanced ECC in additional hardware cost usually runs between 10% to 50%.

Minimum requirements

Online spare memory protection dedicates at least one rank of each memory channel for use as spare memory, so a minimum of two ranks per populated channel have to be present.

How to enable online spare

The user configures the system for online spare mode through the RBSU. Online spare configuration requirements vary for each server series, but it does not require OS support or special software beyond the BIOS.




Figure 3. The online spare feature is configured through the RBSU

Mirrored memory with advanced ECC support

Overview

Mirrored memory with advanced ECC support provides protection against uncorrectable errors that would otherwise result in system failure. There are two modes available—fully and partially mirrored memory support.

- **Full mirrored memory support** uses half of the system memory capacity to maintain one copy of all data.
- **Partial mirrored memory support** enables the user to assign a smaller amount of the system memory for mirroring. This feature is supported with advanced CPU SKUs for Intel Xeon Platinum and Gold processors.

If an uncorrectable error occurs in the mirrored memory protected region, the system automatically retrieves the good data from the redundant copy. The system continues to operate normally without any user intervention. By providing added redundancy in the memory subsystem, memory mirroring provides the greatest protection against memory failure not corrected by ECC, SDDC, DDDC, ADDDC, and online spare memory.

Characteristics

By enabling full mirrored memory, only half the populated memory is usable as system memory. Since full memory mirroring consumes 50% of the system memory capacity, it is targeted for server workloads that must receive the highest level of protection from memory device failures. You might want to consider memory mirroring for workloads that cannot have downtime and cannot risk waiting until scheduled downtime to replace degraded memory modules.

Partial memory mirroring can be configured by the customer and supports different modes:

- OS configured
- First 4 GB of server memory
- 10% or 20% of memory above 4 GB

For more information on partial memory mirroring support, please consult OS documentation.




The performance impact for implementing memory mirroring is generally very small. Because partial memory mirroring uses less memory, the cost of implementation can be significantly lower than full memory mirroring.

Minimum requirements


Gen10 server platforms support two different methods of full mirroring—two- and three-channel mirroring. Three-channel memory mirroring will mirror memory across all three memory channels on each side of the CPU. Two-channel mirroring will mirror between two of the three channels on each side of the CPU (2 and 3; 5 and 6). In both modes, all populated channels on each side of the processor must be populated identically (same number of DIMMs, same DIMM type, same capacity). If DIMMs are populated on both sides of the CPU, the DIMMs on one side do not have to match those on the other, but they must match with the other DIMMs on the same side of the CPU.

HPE Gen10 ProLiant DL360/DL380/DL560/DL580



Three-channel mirroring

HPE Gen10 ProLiant BL460c/XL170r/XL190r/ XL230k/XL450



Two-channel mirroring

Figure 4. Memory mirroring diagrams for HPE ProLiant Gen10 servers

Full memory mirroring is done on one integrated memory controller (IMC). It is not done across both IMC, so the IMC run mirroring independently from each other. Partial memory mirroring follows the same loading rules as full mirroring supported on the platform.

How to enable memory mirroring

Mirrored memory support can be enabled in the RBSU by configuring the advanced memory protection option to “mirrored memory with advanced ECC.” For full mirrored memory, the customer designates half of the memory banks as system memory and the remaining banks as mirrored memory. All banks must be configured identically.




Figure 5. The mirrored memory feature can be enabled in the RBSU



To configure partial memory mirroring, the advanced memory protection option has to be set to “mirrored memory with advanced ECC” and the memory mirroring mode to the appropriate setting as shown below.




Figure 6. Partial memory mirroring is an advanced feature enabled on RBSU

Please note that “partial mirroring—OS configured” is only supported by some OS. Please check OS documentation for details.

Memory scrubbing (Patrol and Demand)

Overview

Memory scrubbing is a standard RAS memory feature designed to prevent soft errors from accumulating in memory and eventually become an uncorrected error. It does this by proactively writing correct data back to memory every time an error is detected. There are two types of scrubbing in today’s systems—patrol scrubbing and demand scrubbing. Both of them do the same thing; once an error is found, they correct it in memory. The big difference is how the error is found. Patrol scrubbing is more of a proactive search for errors occurring continuously in the background, while demand scrubbing occurs only when memory is actually read by the OS or the application.

Characteristics

When patrol scrubbing is enabled, it proactively searches the system memory for correctable errors and repairs them. This prevents the accumulation of single-bit errors that become uncorrectable when the correctable threshold error count is exceeded or degrading into multi-bit errors. There is one patrol scrubber per IMC.

Minimum requirements

There are no specific memory population rules or RBSU settings required for patrol scrubbing. It’s enabled as default on Purley platforms and can be turned off by the user. Demand scrubbing is always enabled by default and cannot be turned off.

How to enable patrol scrubbing

The default is enabled for any advanced memory protection mode selected in RBSU >> Memory Options.






Figure 7. Patrol scrubbing is enabled in RBSU

Technical details

The BIOS enables the patrol scrubbing engine during boot and sets up the scrubbing interval. The scrubbing action involves:

- Reading every cache line once a day to check for errors.
- If errors are found, the correct data is written back to memory.

Patrol scrubs are intended to ensure that correctable errors do not remain in DRAM long enough to have a significant chance of combining with a transient error to cause an uncorrectable error. Patrol scrubbing works in all memory RAS modes, such as advanced ECC, mirroring or rank sparing and helps reduce uncorrectable events.

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.

Hewlett Packard Enterprise is meeting the challenge with fault-tolerant memory protection technologies such as online spare memory, mirrored memory, and Fast Fault Tolerance. 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. Mirrored memory provides a higher level of availability, with a more fault-tolerant option providing full protection against single-bit and multi-bit errors. Fast Fault Tolerance, the latest technology introduced in HPE Gen10 [ProLiant/Synergy/Blade](#) servers using Intel Xeon Scalable processors, combines significantly better memory reliability and availability to the customer.

These HPE 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.



Learn more at
hpe.com/info/memory



Sign up for updates

© Copyright 2013, 2016–2018 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise 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. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

HPE Fast Fault Tolerance is the result of a joint collaboration between Intel and Hewlett Packard Enterprise. Intel and Intel Xeon are trademarks of Intel Corporation in the U.S. and other countries. All other third-party marks are property of their respective owners.

4AA4-3490ENW, September 2018, Rev. 3

