HPE Fast Fault Tolerance for Gen10, Gen10 Plus, Gen11 and Gen12 servers

Contents

Introduction

Server virtualization, cloud computing, high-performance computing, and other resource-intensive workloads demand a lot from server memory. And as workloads increase, so do the number of potential memory failures—failures that could result cause a system crash or the need to reboot.

Because workload demands affect a system's reliability, server memory is a critical component to server integrity. HPE ProLiant Gen10 servers feature HPE DDR4 SmartMemory, which supports faster memory speeds and offers larger capacities than previous generations. It also provides an increasingly comprehensive suite of memory error detection and correction technologies that are collectively called reliability, availability, and serviceability (RAS) features.

The newest HPE SmartMemory RAS technology—HPE Fast Fault Tolerance—enables a server to boot with full memory performance while monitoring for DRAM device failures. In the event of a memory failure, the memory subsystem automatically reorganizes the way data is stored in memory. This creates a protected region just large enough to correct the DRAM failure while allowing the remaining portions of memory to continue to run at peak performance.

Frequently asked questions

Q1. What are the benefits inherent in HPE SmartMemory featuring HPE Fast Fault Tolerance?
A1. If we had to answer that question with a single word, the key benefit of HPE SmartMemory featuring HPE Fast Fault Tolerance is uptime. HPE Gen10 servers configured with HPE SmartMemory offer an extra layer of protection against unplanned server downtime and server crashes. However, unlike the RAS feature called double-device data correction (DDDC), you no longer need to compromise on throughput.

Q2. How was HPE Fast Fault Tolerance developed?
A2. HPE Fast Fault Tolerance is a result of a joint collaboration between Intel® and Hewlett Packard Enterprise to improve the reliability and availability of all Intel® Xeon®-based servers. As a result of the collaboration and the significant role Hewlett Packard Enterprise had in the feature's development, we are able to offer an enhanced version of the feature. This results in significantly better server reliability than the rest of the industry can provide for a 2-year period.

Q3. What is ADDDC?
A3. Adaptive double-device data correction (ADDDC) is the next-generation version of DDDC. DDDC requires two DIMMs on two channels to be run in lockstep that reduces peak memory bandwidth to approximately 60% of single-device data correction (SDDC). ADDDC incorporates the performance benefits of SDDC with the availability of DDDC. In the event that a DRAM device fails, the memory subsystem adaptively reorganizes the way data is stored in memory to create a virtual lockstep region just large enough to correct the DRAM failure. Although the small virtual lockstep memory region will run with reduced memory bandwidth, the remaining portions of memory will continue to run at peak performance.

Q4. **Isn't HPE Fast Fault Tolerance just another name for ADDDC?**
A4. No. HPE Fast Fault Tolerance is an enhancement to ADDDC through the addition of firmware. There are more spare regions and more server memory options to map out bad memory sections, resulting in significantly better server reliability and availability than what the rest of the industry will get with ADDDC.

Q5. If HPE Fast Fault Tolerance is disabled, do you lose that data?
A5. Even when you don't enable HPE Fast Fault Tolerance feature, there is spare server memory integrated or an extra DRAM built into the solution. Therefore, two DRAMs have to fail simultaneously for the system to lose data, which is a highly unlikely event.

Q6. Are there any caveats/trade-offs when enabling HPE Fast Fault Tolerance?
A6. In order to enable the feature, the server has to run in closed-page mode. Some workloads (for example, streams) will see a reduction in throughput caused by this change. Other workloads will see very little impact. If a DRAM fails, there will be an additional throughput reduction in the affected memory region. The reduction in throughput is expected to be minimal for most customers but does depend on the application, the size of the affected region, and the memory configuration.

Q7. How will HPE Fast Fault Tolerance affect the performance of a server running SQL?
A7. Databases tend to use more of a random access pattern than a sequential one. Because of this, the impact would likely be minimal.

Q8. Is HPE Fast Fault Tolerance managed and reported by the HPE Integrated Lights Out 5 (HPE iLO 5) system?

A8. HPE Fast Fault Tolerance is managed by the ROM and reported through HPE iLO when serious errors occur. Hewlett Packard Enterprise does not report intermittent or other minor errors that are not a risk to customer uptime.

Q9. Can HPE Fast Fault Tolerance identify a single DIMM to replace, or can it only isolate to a DIMM pair?

A9. HPE Fast Fault Tolerance can identify a single DIMM to replace.

Q10. When Fast Fault Tolerance is invoked, system performance was restricted/throttled. Is this normal?

A10. Yes, when a Fast Fault Tolerance event occurs, not only will memory bandwidth on that channel be reduced, but a temporary reduction of processor performance will also occur, so the memory sparing event can complete successfully. This should cause minimal performance impact, but is necessary to ensure maximum system uptime.

Q11. I noticed a slight performance impact lasting one minute during a Fast Fault Tolerance event. Can I avoid that performance impact?

A11. While using the HPE Fast Fault Tolerant memory protection mode, this is normal behavior and cannot be changed.

Q12. If I cannot tolerate the performance impact incurred during a Fast Fault Tolerant event, what options do I have?

A12. While not recommended due to the potential for reduced system uptime, you can change the memory protection mode to Advanced ECC. This disables the Fast Fault Tolerant capability, so system downtime can occur during certain memory error events. Depending on your current configuration and requirements, you may be able to utilize other memory protection modes, such as memory mirroring. Please see the Server memory and persistent memory population rules for HPE Gen10 servers with Intel Xeon Scalable processors technical white paper (hpe.com/psnow/doc/a00017079enw) for more information.

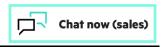### Q13. What is Minimum requirements?

A13. There must be a minimum of two ranks on each populated channel. Also, only HPE SmartMemory in x4 organization can be used. The x4 organization refers to the data width of the DRAM used to construct the DIMM.

Note
There are a number of applications that do not maximize use of the memory bus. For those applications, any impact to performance will be minimal when HPE Fast Fault Tolerance is enabled.

Learn more at
hpe.com/info/memory

Visit **HPE.com**

Chat now (sales)