

Improved Memory Reliability, Availability, and Serviceability (RAS) in PowerEdge YX5X servers

Abstract

This technical white paper describes enhancements made to the memory RAS features in PowerEdge YX5X servers that use Intel 3rd Generation ('Ice Lake') Xeon Scalable Processors. System configurations for achieving the maximum memory up time are called out.

September 2021

Revisions

Date	Author(s)	Description
Jun 2021	Jordan Chin	Initial release
Aug 2021	Jordan Chin	Revised language and wording, corrected typos
Sept 2021	Jordan Chin	Added content for new DIMM map-out feature

Acknowledgements

Author(s): Jordan Chin (Memory Systems Engineering)

Support: Mark Dykstra (Escalation Engineering), Kevin Cross (Memory Systems Engineering), Dave Chalfant (BIOS Engineering), Rene Franco (Memory Systems Engineering), and Trent Bates (Product Management)

Others: Sheshadri PR Rao (InfoDev)

The information in this publication is provided "as is." Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.




Copyright © 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [9/13/2021] [Technical White Paper] [Manager]

Contents

Revisions.....	2
Acknowledgements.....	2
Contents.....	3
Executive Summary	4
1 RAS improvements in the Chipset	5
1.1 Intel Partial Cache Line Sparing (PCLS).....	5
1.2 Improvements to Advanced ECC (or Single Device Data Correction).....	5
2 Smarter Self-healing.....	6
2.1 Memory Health Check	6
2.2 Self-healing during Full-Boot Memory tests	7
3 Other Memory Serviceability improvements	9
3.1 Redesigned Memory Event messaging.....	9
3.1.1 Standard Memory Event Message severity levels	9
3.1.2 Prescriptive Remediation.....	9
3.2 User-directed Memory Map-out.....	10
4 Legacy RAS features	12
5 Achieving Maximum Memory Uptime.....	13
6 Applicable PowerEdge YX5X servers	14

Executive Summary

In our *Memory Errors and Dell EMC PowerEdge YX4X Server Memory RAS Features* technical white paper, we provided a primer on memory errors and the Memory Reliability, Availability, and Serviceability (RAS) features on our PowerEdge YX4X products. Most of the principles and RAS feature information from that whitepaper continue to remain intact for our latest YX5X line of PowerEdge products. However, as with every generation of Dell EMC PowerEdge server product, we continue to innovate and improve. With this document, we would like to share with our customers the next generation of improvements to the memory RAS capabilities in our latest YX5X line of PowerEdge products. At a high-level, this includes the following new features or processes:

Reliability	Availability	Serviceability
 <i>Ability to Prevent Faults</i>	 <i>Ability to Recover from Faults</i>	 <i>Ability to Diagnose, Repair or Service Faults</i>
<ul style="list-style-type: none"> • Intel chipset feature: Improved Advanced ECC (SDDC) 	<ul style="list-style-type: none"> • Intel Partial Cache Line Sparing (PCLS) 	<ul style="list-style-type: none"> • Memory Health Check • Self-Healing During Full-Boot Memory Tests • Redesigned Memory Event Messaging • User-directed Memory Map out

The content covered in this technical white paper applies only to DRAM based memory modules and does not apply to Intel Optane Persistent Memory. Also, this content is specific to Dell EMC PowerEdge YX5X servers with 3rd generation Intel Xeon Scalable Processors.

1 RAS improvements in the Chipset

Intel introduced the following improvements in their third generation of Xeon Scalable Processors:

1.1 Intel Partial Cache Line Sparing (PCLS)

The third generation of Xeon Scalable Processors contains 16 segments of spare cache lines per DDR channel in their integrated Memory Controllers (iMC). When BIOS detects that DRAM on a DIMM is experiencing single-bit errors (SBEs) for a particular cache line, it may request the iMC to re-map the offending portion of the cache-line to use its internal spare cache instead. Single-bit errors are generally harmless, but if left unchecked, they could cause unwanted activation of preventative RAS features that can impact system performance.



This feature will prevent continued triggering of the Dell Predictive Failure Algorithm (PFA) that would otherwise activate other RAS features (such as ADDDC) or consume self-healing resources. This results in fewer performance impacts and corrective actions required by users. This feature is available on all third Gen Xeon Scalable Processors. There are no memory or system configuration requirements necessary to take advantage of PCLS.

1.2 Improvements to Advanced ECC (or Single Device Data Correction)

Intel has redesigned and optimized their Advanced Error Correcting Code in 3rd Gen Xeon Scalable Processors to handle the most common failure patterns known among the major DRAM suppliers. In doing this, many of the multi-bit error patterns that were uncorrectable by previous generations of Intel Xeon Scalable Processors are now correctable by 3rd Gen Xeon SPs. This uplift will result in a significant decrease in uncorrectable memory errors.



This enhancement is available on all 3rd Gen Xeon Scalable Processors. There are no memory or system configuration requirements necessary to take advantage of the improved Advanced ECC (or SDDC).

For more information about Advanced ECC, see the *Memory Errors and Dell EMC PowerEdge YX4X Server Memory RAS Features* technical white paper available on the Dell support site.

2 Smarter Self-healing

Dell Technologies has introduced Memory Health Check and self-healing on full-boots as improvements to memory self-healing in their PowerEdge YX5X servers. When used together, these features should provide a multiplicative effect on the testing and self-healing coverage for memory while minimally impacting the server boot time.

2.1 Memory Health Check

Previous Dell EMC PowerEdge implementations of memory self-healing was based on targeting repairs at a specific memory location where an uncorrectable error or multiple correctable errors were detected. For self-healing to take effect, the system must be rebooted—thereby, resulting in some downtime. Understanding that downtime can be impactful and difficult to arrange, it made sense to make use of any downtime because of self-healing and use it as an opportunity to find other questionable memory locations and repair them. This would prevent the need for future reboot operations and downtime to conduct further self-healing. Thus, Dell Technologies has developed Memory Health Check.



Dell Memory Health Check works by first allowing the original self-healing process to execute and complete as described in the original *Memory Errors and Dell EMC PowerEdge YX4X Server Memory RAS Features* technical white paper. Next, during the same boot cycle as the original self-heal, BIOS executes rigorous testing throughout the DRAM rank where memory errors have occurred. The rigorous testing used by Dell Memory Health Check is the same used by the Dell factory memory screening process. BIOS on the PowerEdge server can quickly and efficiently execute this type of rigorous testing (which can normally take hours) by intelligently testing the affected DRAM rank instead of the entire system memory range.

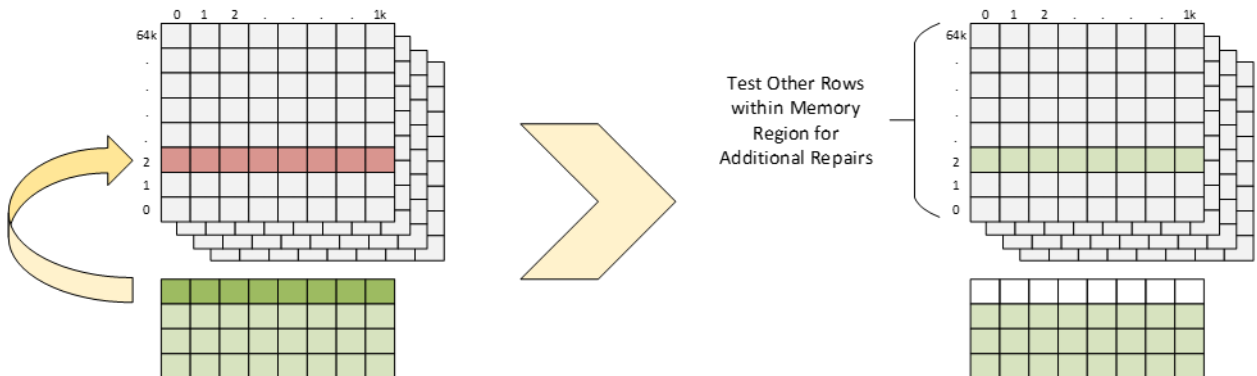


Figure 1 The Memory Health Check process in PowerEdge YX5X servers

Lastly, if any additional questionable locations are found by this enhanced screening, then those locations are automatically repaired within the same boot cycle, thereby incurring no additional reboot penalty.

If BIOS detects that there are more questionable locations than there are self-healing resources available to repair them, then BIOS will automatically map-out that memory region and message to the user that a DIMM replacement is necessary.

The execution of Dell Memory Health Check adds several seconds of boot time and is only executed as needed based on certain types of memory health events from the previous boot. Normal boots should not see an impact to completion time.

This feature is available on all 3rd Gen Xeon Scalable Processors. There are no memory or system configuration requirements necessary to take advantage of Dell Memory Health Check.

2.2 Self-healing during Full-Boot Memory tests



On Dell EMC PowerEdge platforms, the default boot behavior is 'fast boot', where the DDR bus is trained only partially and memory testing phase is skipped. 'Full-boot' is executed after either a BIOS upgrade, memory configuration change, or after a memory health event (such as uncorrectable error or large number of correctable errors) or the system is configured to do so. During a full-boot cycle, BIOS performs memory retraining and a quick memory test among other activities.

On PowerEdge YX5X servers, Dell Technologies has added the ability to perform self-healing based on the results of this quick memory test which differs from the testing in Memory Health Check in the following ways:

	Memory Health Check	Full-Boot Memory Test
Testing Capability	Rigorous (Multiple Patterns)	Light (Single Pattern)
Testing Range	Targeted Single Rank	Entire System Memory Range
Execution	After a Memory Health Event	Full-Boots Only
Self-healing	Yes	Yes

Despite the differences in testing, the self-healing and map-out functions of the two features are identical. If any questionable locations are detected, then those locations are automatically repaired. If BIOS detects that there are more questionable locations than there are self-healing resources available to repair them, then BIOS will automatically map-out that memory region and indicates to the user (through a message) that a DIMM replacement is necessary.

As mentioned earlier, by default, full-boot is executed only in certain conditions. However, full-boot can be enabled to always execute in the BIOS Setup menu. Also, a user can also request for a one-time execution of full-boot. The boot options are accessed as 'Memory Training' settings in the **System BIOS Settings** → **Memory Settings**. Setting 'Memory Training' to 'Enable' will cause the system to always perform full-boot, while 'Retrain at Next Boot' will perform a one-time full-boot operation during the next reboot cycle. Access to these settings are also available in iDRAC9 by clicking **Configuration** → **BIOS Settings** → **Memory Settings**.

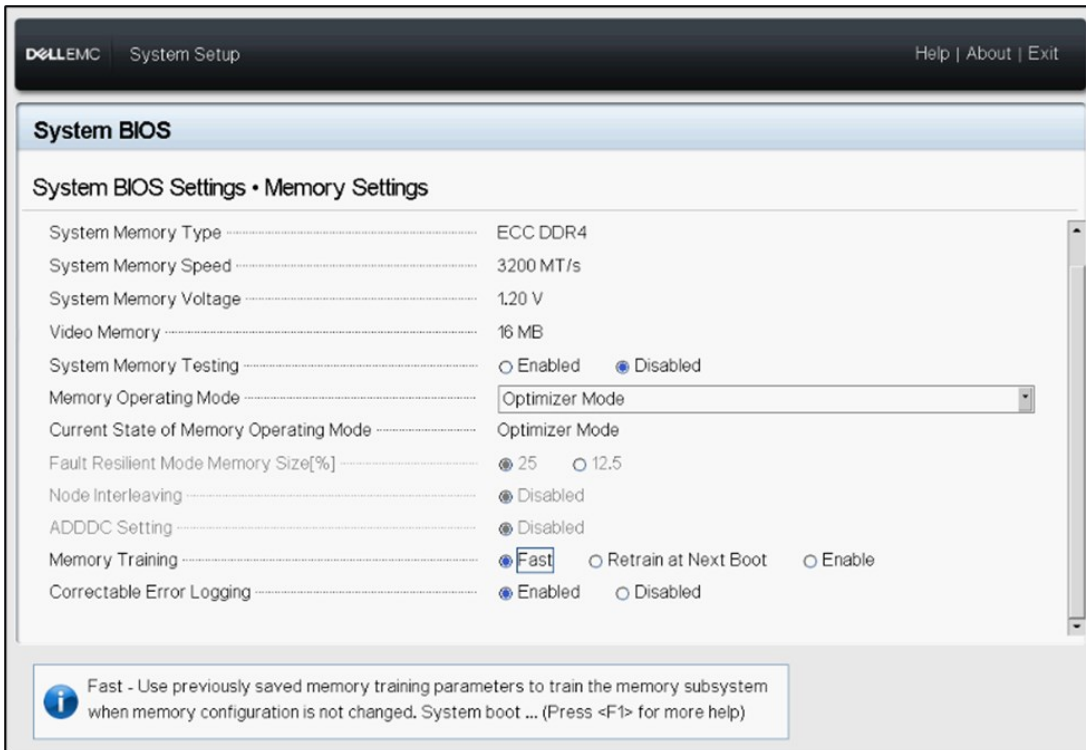


Figure 2 Access Boot options on the System Setup page

This feature is available on all 3rd Gen Xeon Scalable Processors. There are no memory requirements necessary to take advantage of self-healing during full-boot.

3 Other Memory Serviceability improvements

3.1 Redesigned Memory Event messaging

Dell Technologies has approached its YX5X PowerEdge servers with the goal of updating memory event messaging to improve reliability, usability, and serviceability. Run-time memory event messages in the System Event Log and Lifecycle Logs have been redesigned to focus on standardization of memory event messaging and prescriptive remediation.

3.1.1 Standard Memory Event Message severity levels

All run-time memory events are now designated by message ID prefixes MEM51xx, MEM61xx, and MEM71xx for alignment according to event severity. Severity is determined based the long-term health outlook of the memory device and whether the event or device correctable by the user.



MEM51xx (Informational Severity) Events—Memory device is operational but may benefit from additional user action to improve performance or health. The specified user action in the event message is optional. These event messages will appear only in the Lifecycle Log.



MEM61xx (Warning Severity) Events—Memory device has encountered an event that can be corrected with additional user action. The specified action in the event message is required to be performed by the user to prevent potential system outage or resolve an existing outage. If the event persists after performing the specified user action, users should contact the technical support teams for further assistance.

These events messages will appear in both the Lifecycle Log and System Event Log.



MEM71xx (Critical Severity) Events—Memory device has detected an event that has compromised the reliability or usability of the device. Immediate replacement of the device is either required or highly recommended. The specified user action in the event message, if provided, is intended only for use as a short-term remedy.

These event messages will appear in both the Lifecycle Log and System Event Log.

3.1.2 Prescriptive Remediation

After broad review of memory event messaging throughout the server industry, we found it extremely common to recommend a general set of user actions to resolve all memory event types. These actions were typically to reseat the memory device and to update BIOS firmware. Both actions are good practices in maintaining a healthy memory system. Reseating the memory device can reduce any device-to-socket connectivity issues, and while updating firmware, ensures that the system is patched for any known issues and is running the latest memory-training and error-handling algorithms. However, almost all solutions require some system downtime and/or physical access to execute. Thus, spending time strictly to complete good maintenance practices is an inefficient use of downtime. As mentioned earlier, downtime can be impactful and difficult to arrange, and as such, every solution prescribed should be carefully considered for the highest likelihood of resolving the problem.

Therefore, in addition to severity, run-time memory event messages on YX5X PowerEdge servers are now also categorized based on the recommended resolution. These recommended user actions are recommended as the best engineering recommendation to address each underlying event. Recommended

user action aligns with the last two digits of the message ID for all MEM51xx, MEM61xx, and MEM71xx events.

MEM 51xx, 61xx, 71xx	Recommended Response Action (RRA) for Issue Resolution
00	No action is required; the memory sub-system is operating normally.
01	Check the device or system configuration.
02	Reseat the device in the same slot location.
03	Perform a warm reboot.
04	Perform a cold reboot or system DC power cycle.
05	Perform a system AC or input power cycle.
06	No action is required; contact support if event re-occurs.
07	Retry operation.
08	Sanitize the persistent memory device.
09	Re-enter the encryption passphrase to unlock persistent memory device or perform a secure erase.
10	Update firmware for the memory device.
14	Contact support to replace the device.
99	Update iDRAC firmware. Note: This indicates that iDRAC could not properly identify the event and is caused by a firmware mismatch with BIOS.

Important—The table above is a simplified version of the listed Recommended Response Actions for MEM 51xx, 61xx, and 71xx. Refer to the appropriate product documentation for the fully detailed Recommended Response Actions which will include next steps for resolution beyond those listed here.

3.2 User-directed Memory Map-out

PowerEdge YX5X servers configured with BIOS version 1.3.x or later will enable users through firmware to directly remove a DIMM from the system memory pool without physically handling the device. When a DIMM is no longer part of the system memory pool, its DRAM components will not be used for any read/write operations. Total system memory capacity is reduced to exclude the mapped out DIMM. Any DRAM faults on this DIMM will no longer produce correctable or uncorrectable errors during run-time operation.



Important—DIMMs that are mapped out from the system memory pool are still configured, trained, and tested during the boot operation. Therefore, enabling memory map-out cannot prevent many memory related issues that may arise during boot.

Mapping out one or more DIMMs can produce an imbalanced memory configuration and significantly impact system performance. To retain a supported memory configuration, Dell EMC BIOS may automatically map out other DIMMs that are not targeted by you. This is necessary to maintain an optimized memory

configuration. Therefore, total system memory capacity may be reduced beyond the expected amount. This type of automatic map-out behavior also occurs if the DIMM was physically removed and the system is in an unsupported memory configuration.

DIMMs that have been manually mapped out will display the MEM6101 message in system logs to indicate successful map-out. Meanwhile, DIMMs that have been automatically mapped out to maintain a supported memory configuration will display the MEM6101 message and UEFI0339 message in Lifecycle Logs.

User directed memory map-out can be enabled and configured in the BIOS Setup menu under System BIOS Settings → Memory Settings → Memory Map Out. Access to these settings are also available in iDRAC9 via Configuration → BIOS Settings → Memory Settings → Memory Map Out.

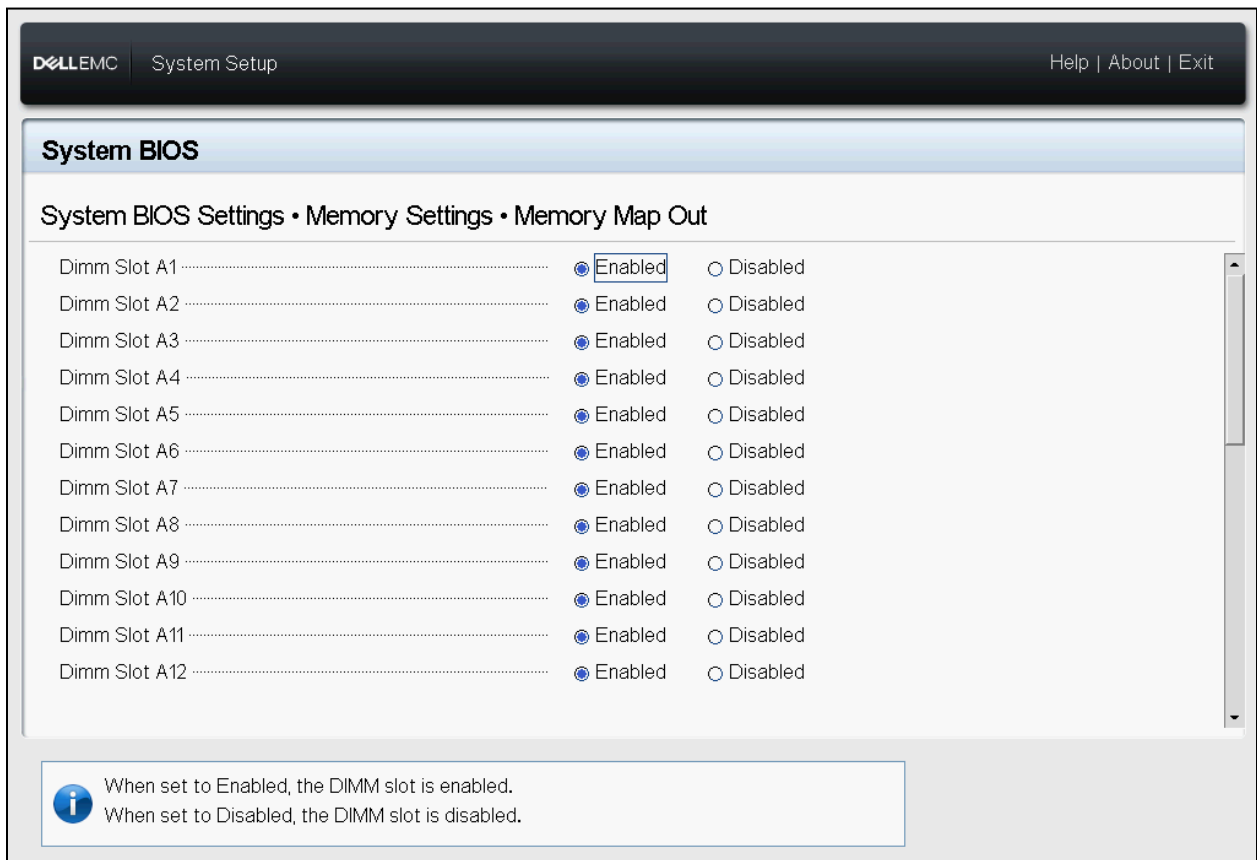


Figure 3 Enable the Memory Map Out feature on the System Setup page

Map-out takes effect during the next system reboot operation and will remain in effect until you manually reenables the DIMM. Map-out settings will persist through configuration changes and BIOS firmware updates. There are no CPU or memory configuration requirements necessary to take advantage of user directed memory map-out.

Warning—This feature should be used with care to ensure that at least one physically installed DIMMs is always enabled. Accidentally disabling all physically installed DIMMs will prevent the system from booting. The system BIOS NVRAM must be cleared to recover from this state. For information about clearing NVRAM, see the respective Installation and Service Manual available on the support site.

4 Legacy RAS features

The RAS features that were supported in the YX4X servers are still supported in YX4X. See the *Memory Errors and Dell EMC PowerEdge YX4X Server Memory RAS Features* technical white paper available on the support site. However, some updates to those features are listed here as sub-bullets:

- Single Error Correction—Double Error Detection (SEC-DED) ECC
- Advanced ECC
 - [Improvements to Advanced ECC \(or Single Device Data Correction\)](#)
- Adaptive Double Device Data Correction (ADDDC)
- Memory Patrol Scrub
- Operating System Memory Page Retirement / Off-lining
- Fault Resilient Memory (FRM) / Address Range Mirroring mode
 - A Non-Uniform Memory Access (NUMA) aware FRM option has been added. When NUMA aware FRM has been enabled, BIOS enforces creation of fault-resilient memory regions which span NUMA nodes, and if SNC is enabled, then also spans across sub-NUMA clusters. Earlier, YX4X servers created a fault-resilient region of memory starting from the lowest system memory address without awareness of NUMA. Address range mirroring requested and enabled by the operating system is and remains NUMA-aware.
- Memory Self-Healing
 - [Smarter Self-Healing](#)
- Machine Check Architecture Recovery

The following legacy RAS features have been deprecated in PowerEdge YX5X servers because of insufficient user enablement:

- Memory Rank Sparing
- Full Memory Mirroring

5 Achieving Maximum Memory Uptime

The following summarizes actions that users can take to achieve maximum memory uptime on their PowerEdge YX5X servers:

- Configure server using genuine DIMMs from Dell Technologies
 - Benefit—Memory modules are fully validated and assured by Dell Technologies. Additional self-healing (PPR) resources above and beyond industry standards. Rigorously screened using proprietary Dell Technologies test patterns and DRAM vendor unique test patterns.
- Configure server with x4 DRAM based DIMMs
 - Benefit—Single DRAM Device Correction and ADDDC.
- Configure server to operate in Fault Resilient Memory mode
 - Benefit—Significantly reduced probability of UCEs in critical portions of memory used by operating systems. Low memory capacity reduction overhead (depending on the system settings).
- Configure server to run memory patrol scrub in 'Extended Mode'
 - Benefit—Patrol-scrub will run after every four hours (instead of 24). Increased scrubbing frequency will reduce the accumulation of errors in areas of memory with low utilization that are not being corrected by demand scrub.
- Configure server to run 'Full-Boot' instead of 'Fast Boot'
 - Benefit—Server performs full DDR bus training and testing after every boot operation. Re-training ensuring that the bus is optimally tuned given the current environmental and operating conditions. The quick memory testing will self-heal questionable memory cells that may have developed over time.
- Regularly check for MEM events in Lifecycle Logs and perform the actions as specified—Even for MEM51xx events displayed for information purposes only.
 - Benefit—Actions specified by MEM61xx and 71xx events are required while those specified by MEM51xx are optional. Even though the MEM51xx events are optional, PowerEdge BIOS has made these recommendations based on monitoring memory health and performing the actions may improve memory performance or long-term reliability.
- Regularly update system firmware
 - Benefit—System firmware is patched for any known issues and is running the latest memory training and error handling algorithms.

6 Applicable PowerEdge YX5X servers

The following platforms are considered PowerEdge YX5X servers equipped with 3rd Gen Xeon Scalable Processors and are therefore covered by this technical white paper:

Important—Subsequent to the publication of this document, Dell may continue to add products to its YX5X server lineup. If a product is not listed below, please consult with a Dell sales or support representative to confirm the server generation.

PowerEdge leveraged products such as some Precision workstations may also be covered by this document. Please consult with a Dell sales or support representative to confirm.

- PowerEdge R750
- PowerEdge R750xa
- PowerEdge R750xs
- PowerEdge R650
- PowerEdge R650xs
- PowerEdge R550
- PowerEdge C6520
- PowerEdge MX750c
- PowerEdge R450
- PowerEdge T550
- PowerEdge XR11
- PowerEdge XR12