

Cisco UCS M7 and M8 Compute Servers Memory Technical Overview – Memory RAS Features



Contents

Abstract	3
Overview of memory errors.....	4
Memory RAS features	6
Memory error handling through RAS features	7
Intel Platform-specific memory RAS features	8
Conclusions	10
References.....	10

Abstract

Modern servers, including Cisco UCS® M7 and M8 compute servers, provide increased memory capacities that run at higher bandwidths and lower voltages. These trends, along with higher application memory demands and advanced multicore processors, require additional memory RAS features to maintain system operation and performance.

Intel® 4th Gen Xeon® Scalable Processors (formerly “Sapphire Rapids Servers”) and 5th Gen Xeon Scalable Processors (formerly “Emerald Rapids Servers”) have implemented changes to incorporate DDR5-supported Reliability Availability and Serviceability (RAS) features. And these features are now available on all Cisco UCS M7 platforms.

Intel Xeon 6 Processors (formerly “Granite Rapids Servers”) have implemented changes further incorporating some additional DDR5-supported Reliability Availability and Serviceability (RAS) features. And these features are now available on all Cisco UCS M8 compute platforms.

AMD EPYC 4th and 5th generation processors support RAS features, although their RAS feature set differs slightly from that of Intel Xeon processors.

Cisco UCS platforms support both x4 and x8^{1, 2} Dual In-line Memory Modules (DIMMs). It is important to note that some of the RAS features listed below are supported only on x4 DIMMs and have been highlighted accordingly.

Refer to the appropriate Cisco UCS memory guides for details on supported memory modules on Cisco UCS platforms:

Cisco UCS M7 Memory Guide: [Cisco UCS/UCSX M7 Memory Guide](#)

Cisco UCS M8 Memory Guide (AMD): [Cisco UCS AMD M8 Memory Guide](#)

Cisco UCS M8 Memory Guide (Intel): [Cisco UCS Intel M8 Memory Guide](#)

This paper describes the classification and handling of memory errors on the following Cisco UCS servers:

- Cisco UCS M7 servers with Intel 4th and 5th Gen Xeon Scalable Processors
- Cisco UCS M8 servers with Intel Xeon 6 Processors
- Cisco UCS M8 servers with 4th and 5th Generation AMD EPYC Processors

Cisco UCS firmware supports advanced platform-specific features such as Adaptive Double Device Data Correction (ADDDC) sparing (on Intel platforms) and Post-Package Repair (PPR) on both AMD- and Intel-based platforms. These features provide an optimal balance of performance, memory capacity, and error resilience. Memory mirroring is recommended for mission-critical applications unable to tolerate memory failure.

¹ x8 16GB DIMMs are supported on Cisco UCS M7 and Cisco UCS M8 servers.

² x8 32GB DIMMs are supported only on Intel-based Cisco UCS M8 servers.

Software requirements and recommendations

Refer to Table 1, below, for the baseline firmware version supporting the RAS features that are described in this document. Before upgrading firmware, the memory RAS BIOS policy should be reviewed and set appropriately to avoid additional reboots after upgrading the firmware. See the “Memory error handling through RAS features” section below, for each of the supported platforms, for details and recommendations on selecting a policy setting.

Table 1. Baseline firmware

Software product	Minimum server firmware supporting ADDDC sparing and PPR	Recommended server firmware
Cisco UCS M7 blade servers and integrated Cisco UCS M7 rack servers	4.3(3a)	4.3(5e) or higher
Cisco UCS M8 blade servers and integrated Cisco UCS M8 rack servers (Intel platform)	4.3(5a)	4.3(5e) or higher
Cisco UCS M8 blade servers and integrated Cisco UCS M8 rack servers (AMD platform)	4.3(6a)	4.3(5e) or higher

Overview of memory errors

Memory errors are among the most common types of errors on modern servers. Errors are often discovered when an attempt is made to read a memory location and the value read does not match the value last written.

Memory errors can be soft or hard. Some errors are correctable, but multiple simultaneous soft or hard errors on a single memory access may be uncorrectable. Overall error rates will scale with total memory capacity in a system, both by individual DIMM capacity and total quantity of DIMMs. The impacts of soft and hard errors are mitigated through hardware and firmware features as described in this document. Soft and hard errors can be introduced post manufacturing by high-energy-particle strikes or normal device wear over time. DDR memory is increasingly made with smaller process nodes running at higher frequency with reduced voltages, which increases the potential for errors. Mitigating these errors is critical to the basic functionality of the system to provide availability and resiliency against system downtime.

Soft errors

Errors caused by brief electrical disturbances within the DRAM or on an external memory interface are referred to as “soft” errors. Soft errors are often transient and do not always repeat. If the error was the result of a disturbance during the read operation, then retrying the read may yield correct data. If the error was caused by a disturbance that upset the contents of the memory, then rewriting the memory location will correct the error.

Soft error rates can be affected by temperature, altitude, and memory access patterns specific to particular workloads. Note that workloads do not correlate directly to applications. The same application that might trigger a correctable error on a Dual In-line Memory Module (DIMM) in one server may not generate an error with a different data set. Memory-test algorithms are tuned to represent worst-case workload behavior, but previously undetected errors may still occur during runtime. Cisco reviews and revises test algorithms to improve fault detection.

Hard errors

Errors caused by persistent physical defects are traditionally referred to as “hard” errors. Hard errors may be caused by an assembly defect such as a solder bridge or cracked solder joint or may be the result of a defect in the memory chip. Rewriting the affected memory contents and retrying read access will not eliminate a hard error. The error will persist.

Correctable errors

If errors are detected and corrected, they are considered correctable. This can be accomplished by retrying the read or by calculating the correct memory contents using ECC data and writing the correct data back into memory. After an error is detected and corrected, the Cisco® Integrated Management Controller (IMC) will log the event in the System Event Log.

Typically, correctable errors are the result of soft errors. If correctable errors persist within the same memory location over an extended period, it may indicate potential hard error.

Uncorrectable errors

An error is deemed uncorrectable when it exceeds the correction capability of the processor’s Error Correcting Code (ECC) engine.

An uncorrectable error experienced during runtime usually results in a catastrophic processor crash or hang, which will cause a server outage. This requires a reboot of the affected server and a replacement of the component that is at the root of the error. Usually this is the memory module, but the root cause could also be tied to a processor, processor socket, or DIMM socket.

After experiencing an uncorrectable error and rebooting the server, Cisco UCS will automatically map out (disable) the affected DIMM. This allows the server to return to service while preventing a second failure from the same module.

If any memory errors are detected during system power on testing, they are deemed uncorrectable, and the module will be mapped out. This is often an indication of a hard error, and the module should be replaced.

Minimizing early production runtime errors

Some errors are most likely to occur earlier than expected in the DIMM lifecycle. As noted, errors can be caused by electrical disturbances or physical defects. Some electrical disturbances can degrade a device and introduce a physical fault (hard errors). Prior to shipment, Cisco performs comprehensive system-level testing on all servers. Analysis of field errors has shown a correlation between physical shipping and new errors. Based on this analysis, Cisco recommends running an Enhanced Memory Test (EMT) in combination with memory diagnostics prior to placing servers into production. This minimizes the time interval and potential introduction of errors between the latest memory testing and execution of production workloads.

Likewise, after upgrading or swapping memory, Cisco recommends running the EMT tool in combination with memory diagnostics to help identify installation errors and minimize potential early failures. The most common errors observed during memory upgrades are attributed to DIMM-seating or installation issues.

It is recommended to enable PPR while running EMT. This gives the system the ability to immediately repair weak or faulty DIMM rows instead of experiencing failures during runtime.

Memory RAS features

The following section will discuss the available RAS features on Cisco UCS M7 and M8 compute servers. Cisco UCS M7 compute servers are platforms using Intel Xeon Processors (4th and 5th generation CPUs) while UCS M8 compute servers offers platforms using AMD CPUs and Intel CPUs. Certain memory RAS features are common to both these platforms. These are mentioned in the “Common memory RAS features” section in this document, whereas CPU-specific memory RAS features are discussed separately under Intel-based platforms.

Common Memory RAS features

The following memory RAS features are supported on both Cisco UCS M7 and M8 (AMD- and Intel-based) compute servers.

System-level ECC

All Cisco UCS M7 and M8 compute servers use memory modules with ECC codes that can correct any error confined to a single x4 DRAM chip and detect any double-bit error in up to two devices. This is now referred to as system-level ECC, as in older generation servers.

On-die ECC

On-die ECC is a new feature in DDR5. This feature is enabled by default. All single-bit errors (hard and soft) are corrected by DRAM before data is transmitted to the host. However, this corrected data is not written back to DRAM. Error Check and Scrub (ECS) is the feature used to scrub and correct single-bit errors in memory.

Error Check and Scrub (ECS)

Error Check and Scrub (ECS) checks for errors in the background by scrubbing each DRAM die periodically (every 24 hours), correcting the errors by writing data back to the array and providing a count of the errors found during the scrub. This feature is enabled by default.

Boot-time Post Package Repair (PPR)

Post-Package Repair (PPR) is a feature where spare rows are used to replace a bad cell or row in a DRAM device. Cisco UCS M7 and M8 servers with Intel and AMD CPUs support “hard” PPR. This is a permanent repair and is carried out during reboot based on the error data collected during the previous runtime or if any row errors are encountered during EMT.

Note:

1. Boot-time PPR relies on data stored from the previous boot to repair the DIMMs, so rebooting the system once before doing a BIOS update is highly recommended so that data requesting PPR is not lost.
2. PPR status is logged to BMC.

Runtime Post-package-repair

Cisco UCS M8 servers support runtime PPR. This feature repairs the errors during system runtime and is considered a soft repair. The repair is valid only for the current boot cycle. The PPR option should be selected in the BIOS for runtime PPR feature to be enabled.

Command/Address Parity check and retry

Command/Address (CA) parity protects the command and address interface. If a CA parity error occurs on a write cycle, the DIMM will alert the host and will not write data to memory. The host will have an opportunity to retry the operation. If a CA parity error occurs on a read cycle, the host will drop the data and retry the read. This feature is enabled by default.

Memory error handling through RAS features

Cisco UCS M7 and M8 servers have additional features, described below, that are used to decide when a particular RAS feature is utilized to manage DIMM-related errors.

Enhanced Memory Test

Enhanced Memory Test (EMT) is a suite of test patterns provided by CPU vendors AMD and Intel, to stress the system memory at boot time.

If correctable or uncorrectable errors occur at runtime, by default the EMT, if set to ‘Auto’ in BIOS, is enabled to run on the next boot to test the entire memory. If PPR is enabled, EMT PPR can attempt to repair damaged regions on the DIMM immediately and return the DIMM to operation. Note that this operation will add a significant delay in the boot time. The repair during boot time is considered hard PPR. Note that, depending on the outcome, the DIMM may become disabled if the DIMM fails EMT and cannot be repaired through PPR.

DIMM blocklisting

Cisco Memory blocklisting, when enabled, will disable any DIMM that previously had an uncorrectable error until it is either removed from the system or PPR is successful, and memory errors are cleared from UCSM/IMM.

Intel Platform-specific memory RAS features

Cisco UCS M7 and M8 Intel-based servers have additional RAS features, which are discussed below.

Single Device Data Correction (SDDC)

Single Device Data Correction (SDDC) is a multibit correction code developed by Intel. Most of the failures in a single DRAM can be corrected using this code. This feature is supported on Cisco UCS M7 and M8 servers when using DDR5 x4 DIMMs.

Correction on x8 DIMMs

Intel CPUs support half-device correction when x8 DIMMs are used, such as 32GB 2R x8, on Cisco UCS M8 servers. SDDC code corrects errors that fall on either right x4 or left x4 of the DRAM.

Memory Mirroring

Memory mirroring is a final backup to handle memory failures that cannot be handled by the features described above. A duplicate copy of all data is maintained in case the original data is corrupted beyond repair.

Virtual Lockstep (VLS)/Adaptive Double Device Data Correction (ADDDC sparing)

Adaptive Double Device Data Correction (ADDDC) sparing can correct two successive DRAM failures if they reside in the same region. This feature tracks Correctable Errors (CE) and dynamically maps out failing bits by spare-copying contents into a “buddy” cache line. This mechanism can mitigate correctable errors that, if left untreated, could become uncorrectable. This feature uses Virtual Lockstep (VLS) to assign cache-line buddy pairs within the same memory channel at either DRAM bank-level using bank VLS or DRAM device-level using rank VLS.

Intel CPUs support ADDDC in Single Region (SR) and in Multi Regions (MRs). ADDDC-MR is a feature where ADDDC can correct CEs in different banks or ranks on two DIMM regions in a memory channel.

This feature is supported only while using DDR5 x4 DIMMs.

All Cisco UCS M8 CPU families support Advanced RAS features.

Cisco UCS M7 CPU families:

Platinum and Gold SKUs support Advanced RAS features supporting both ADDDC-SR and ADDDC-MR while Silver and Bronze support Standard RAS features supporting ADDDC-SR.

Platinum and Gold CPUs support both bank and rank VLS. Silver and Bronze CPUs support only bank VLS.

If errors persist after a sparing event, the process repeats as needed until all of the spare bits are consumed. Spare bits are obtained from buddy cacheline pairs by reusing redundant ECC bits produced by the lockstep process. ADDDC sparing does not require allocation or usage of spare main memory regions and does not reduce the overall memory available to the operating system.

Initiation of VLS/ADDDC

VLS will be triggered depending on the number of errors detected in the same row. VLS will attempt to temporarily swap the bank or rank with a spare at runtime.

Supported memory RAS features

Table 2 lists the memory RAS features supported on Cisco UCS M7 and M8 platforms.

Table 2. Support Memory RAS features on UCS M7 and M8

Feature	Intel Xeon 4 th and 5 th (Cisco UCS M7)	Intel Xeon 6 (Cisco UCS M8)	AMD 4 th /5 th Gen (Cisco UCS M8)	Notes/restrictions
System-level ECC	✓	✓	✓	
On-die ECC (DDR5)	✓	✓	✓	DDR5 DIMM feature
Error check and scrub (ECS)	✓	✓	✓	
Boot-time post-package repair (PPR)	✓	✓	✓	“Hard” PPR, all platforms
Runtime post-package repair (PPR)		✓	✓	Cisco UCS M8 only, BIOS option
CA parity check and retry	✓	✓	✓	
Enhanced memory test (EMT)	✓	✓	✓	
DIMM blocklisting	✓	✓	✓	

Intel CPU-specific features supported on Cisco UCS M7 and M8 platforms are listed in Table 3.

Table 3. Intel specific Memory RAS features on UCS M7 and M8

Feature	Intel Xeon 4 th and 5 th (Cisco UCS M7)	Intel Xeon 6 (Cisco UCS M8)	Notes/restrictions
Single device data correction (SDDC) (SDDC)	✓ (x4 DIMMs)	✓ (x4/x8 DIMMs)	x4: full-device correction x8: half-device correction
Memory mirroring	✓	✓	
ADDDC sparing (VLS/ adaptive double device)	✓*	✓	* x4 DIMMs only; Cisco UCS M7 Platinum/Gold: SR/MR Cisco UCS M7 Silver/Bronze: SR only
Bank and Rank VLS	✓**	✓	**Cisco UCS M7 Platinum/Gold: bank/rank Cisco UCS M7 Silver/Bronze: bank only

Cisco UCS M8 servers powered by Intel Xeon 6 Processors support two variants of 32GB DDR5 6400MT/s modules: 1Rx4 and 2Rx8. The 32GB 1Rx4 module enables full support for critical RAS features such as SDDC and ADDDC, while the 32GB 2Rx8 module offers greater flexibility in memory population. This distinction provides technical advantages in memory configuration and resiliency for enterprise workloads.

Conclusions

The features listed above work together to help prevent correctable and uncorrectable errors from severely degrading system performance. Uncorrectable errors may, though they do not always, bring a system down. Single-bit correctable errors are expected to occur sporadically and fixed by ECC before any corruption occurs. Single-bit errors on their own are not a cause for concern and do not require user intervention. Cisco UCS M7 and M8 compute servers employing DDR5 DIMMs use multiple memory RAS features for improved system uptime and reduced service actions compared to previous generation servers.

References

- Intel document Doc. No.: 638563, Rev.: 1.75
- Intel document Doc. No.: 742874, Rev.: 1.5
- AMD document Doc. No.: 57298, Rev.: 0.95
- AMD document Doc. No.: 57065, Rev.: 1.05