



Lenovo AI-Based Memory Predictive Failure Analysis Planning / Implementation

While components like CPUs, storage, and networking are essential, server performance and reliability are heavily dependent on system memory. Memory errors—caused by electrical noise, cosmic rays, or hardware degradation—are categorized as either correctable or uncorrectable. While processors can seamlessly fix correctable errors, uncorrectable errors can lead to application crashes or complete system hangs, making robust memory a cornerstone of server stability.

Lenovo's ThinkSystem V3 generation servers, introduced in 2023, feature DDR5 memory. DDR5 includes advancements like on-die Error-Correcting Code (ECC) to proactively correct single-bit errors within the DRAM chip, and support ECC Error Check and Scrub (ECS) mode with an error counting scheme for transparency. However, its higher speeds and doubled bandwidth compared to DDR4 also introduce new signal integrity challenges, making sophisticated error handling more critical than ever.

Traditionally, memory errors are managed by a UEFI SMI handler, which is triggered when correctable errors reach a predefined threshold. The drawback is that these SMI interrupts can impact OS workload performance, preventing the threshold from being set low enough for early detection. A more effective solution is to use the server's Baseboard Management Controller (BMC) to gather detailed memory telemetry without impacting the OS, enabling more accurate error handling. Intel's [Memory Resilience Technology \(MRT\)](#) is one such example.

To provide this advanced capability, Lenovo has implemented *Memory Predictive Failure Analysis* (Memory PFA, or simply MPFA) across its ThinkSystem V3 and V4 servers for both Intel Xeon and AMD EPYC platforms. MPFA integrates UEFI and BMC firmware to collect comprehensive memory error data. This data is fed into a Lenovo AI-based engine that uses historical data to predict failures. Based on this predictive analysis, the UEFI and BMC work in concert to initiate proactive memory recovery actions, significantly enhancing server uptime and reliability.

The following table shows key differences between Lenovo MPFA and Intel MRT technology. Lenovo will continue enhance the MPFA design and AI based heuristic algorithm for this feature.

Table 1. Key differences between Lenovo MPFA and Intel MRT

Features	Lenovo Memory Predictive Failure Analysis (MPFA)	Intel Memory Resilience Technology (MRT)
Processor Support	Both Intel Xeon and AMD EPYC	Intel Xeon only
Data Collection	Correctable errors, Uncorrectable errors, DDR5 ECS errors	Limited
AI Prediction	Yes (heuristic), and continuous refinement with data mining from Lenovo ThinkSystem servers	Basic Thresholds
Recovery Methods	Page Retire, sPPR/hPPR, ADDDC for Intel	Limited

Lenovo MPFA Infrastructure

The Lenovo AI-powered Memory Predictive Failure Analysis (MPFA) system is an advanced predictive analytics engine designed to enhance server reliability by anticipating DIMM failures before they occur. The core of this system is the MPFA Analyzer, which utilizes a sophisticated heuristic algorithm.

This algorithm is the result of model distillation from a comprehensive, offline machine learning (ML) model. The ML model is continuously trained and refined using Lenovo's extensive service database, which contains a vast collection of memory error telemetries from Lenovo servers in the field. This database includes detailed parameters such as platform type, DIMM manufacturer and model, slot location, and specific error signatures. As the dataset expands, the ML model's accuracy improves, allowing for the generation of increasingly precise heuristic algorithms for deployment on production servers.

The Memory Predictive Failure Analysis (MPFA) feature, implemented on Lenovo ThinkSystem V3 and V4 DDR5-based servers, is a cooperative system that relies on two core components:

- **UEFI Firmware:**
Responsible for initial configuration, collecting memory errors via System Management Interrupts (SMI), and executing hardware-level recovery actions.
- **XClarity Controller (XCC):**
The Lenovo Baseboard Management Controller (BMC) that continuously polls for memory errors, stores all telemetry in an MPFA database, and hosts the AI-based analyzer to predict DIMM failures.

The MPFA process operates through the following coordinated workflow, as illustrated in the figure below:

1. **System Initialization:**
During boot, the UEFI provides the XCC with critical memory configuration details, such as mirroring and ECC modes. This information serves as a baseline for the AI-based MPFA Analyzer.
2. **Continuous Data Polling:**
The XCC periodically polls memory controller registers using an out-of-band hardware interface (e.g., Intel PECL or AMD APML). This allows it to gather data on correctable and uncorrectable errors without impacting the host system. All findings are stored in the MPFA database.
3. **Event-Driven Error Reporting:**
In addition to continuous polling, when a predefined memory error threshold is reached, the UEFI SMI handler is triggered. It collects detailed error information and reports it to the XCC, which adds it to the MPFA database.

4. Feeding Memory Errors into MPFA Analyzer:

All correctable and uncorrectable memory errors with detailed telemetries will be feed into MPFA analyzer, and these data will be used to monitor DIMM health status and predict faults in next step.

5. AI-Based Analysis and Prediction:

MPFA analyzer on the XCC processes the combined telemetry from both polling and SMI events. It uses a heuristic algorithm, derived from Lenovo's offline machine learning models, to monitor DIMM health and predict faults. If a recoverable fault is predicted, the XCC notifies the UEFI by triggering an SMI.

6. Proactive Recovery:

Upon receiving the SMI from the XCC, the UEFI handler identifies the specific memory fault address. It then initiates the appropriate recovery action, which may include requesting the OS to retire a memory page (via SCI interrupt) or executing silicon-level repairs like soft or hard Post-Package Repair (PPR) or Adaptive Double Device Data Correction (ADDDC) if applicable.




Figure 1. Lenovo AI based MPFA Infrastructure

Lenovo self-designed UEFI + BMC have the full solution of Lenovo AI based MPFA, and this MPFA solution has been implemented on all Lenovo DDR5 based ThinkSystem server products.

Key Features

The MPFA technology is implemented through a set of integrated features for data collection, analysis, and proactive fault mitigation.

Data Collection and Telemetry:

- Comprehensive Error Logging:

The system persistently collects and logs all correctable (CE) and uncorrectable (UE) memory errors directly on the Baseboard Management Controller (BMC). Each error entry is augmented with detailed system telemetry for contextual analysis.

- DIMM Self-Repair Monitoring:

Information regarding intrinsic DIMM recovery actions, such as Post Package Repair (PPR), is captured and correlated with error data to provide a holistic view of the DIMM's health status.

- Redfish API Support:

Provides standardized Redfish API endpoints for secure, remote access to all collected memory error records, enabling further offline analysis and integration with external data center management tools.

AI-Driven Predictive Analysis and Mitigation (advanced features):

- Heuristic Analysis Engine:

The onboard MPFA Analyzer ingests the collected DIMM error and telemetry data in real-time. It applies the AI-derived heuristic algorithm to predict imminent hardware failures, including specific DIMM cell and row faults.

- Proactive Fault Response:

Upon predicting a failure, the MPFA Analyzer alerts the Lenovo XClarity Controller (XCC), which then triggers a System Management Interrupt (SMI) to initiate an immediate, out-of-band response.

- UEFI-Based Remediation:

A dedicated UEFI SMI handler executes a suite of memory recovery procedures to mitigate the predicted fault and preserve system stability. These automated actions include:

- Requesting page retirement from the operating system to isolate the failing memory address.
- Applying a runtime soft PPR to failing row for immediate, non-disruptive remediation.
- Applying ADDDC to failing banks if applicable for immediate, non-disruptive remediation.
- Scheduling a hard PPR to be applied during the next system reboot for a permanent repair.

Redfish commands

Lenovo MPFA AI-Driven Predictive Analysis and Mitigation features are Advanced Features which are disabled by default, and these features are controlled by XCC Platinum license, so customer can only enable MPFA advanced features if that system has XCC Advanced or Enterprise license. XCC license status can be checked in “System Information and Settings” of Lenovo BMC web, as shown in the following figure.




Figure 2. Lenovo BMC License Info in XClarity Controller

There are several Redfish commands to support MPFA related features, as described in the following topics:

Basic logging and Advanced Features

The Redfish commands are as follows:

- **BasicLoggingEnabled** – MPFA basic logging feature is enabled by default, and customer can disable it if customer does not want to collect memory related error information. But once it is disabled, it will impact Lenovo service support for DIMM related events.
- **AdvancedFeaturesEnabled** – MPFA advanced features can be enabled/disabled by this setting. Enabling MPFA advanced features can do better memory recovery and enhance system reliability. This command can be enabled successfully only when XCC has advanced or enterprise license.

Use the following URI and GET method to get the default settings:

- Redfish URI: /redfish/v1/Managers/1
- HTTP Method: GET

An example of this method is shown in the figure below.

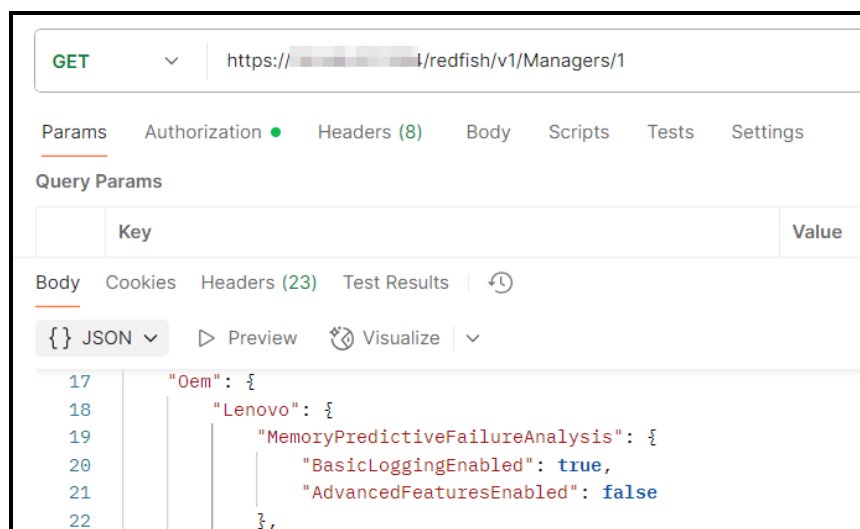


Figure 3. Using Redfish to GET MPFA Default Configuration

You can also use the URI and PATCH method to enable both basic logging feature and advanced features.

- Redfish URI: /redfish/v1/Managers/1
- HTTP Method: PATCH

An example of this method is shown in the figure below.




Figure 4. Using Redfish to PATCH MPFA Configuration

DIMM health status and Top faults information

The following commands can be used to check each DIMM's health status and top faults of that DIMM.

- **DimmHealthStatus** – it contains Major status and Minor status, and you can contact Lenovo support team for more details of status definition for ThinkSystem products if needed.
- **DimmTopFaults** – if the counter of DIMM fault is less than 5, then all faults will be reported to DimmTopFaults list; otherwise, only top 5 severe faults of each DIMM memory will be reported into DimmTopFaults list. And you can contact Lenovo support team for more details of fault type definition if needed.

The URI and Method are as follows:

- Redfish URI: /redfish/v1/Systems/1/Memory/{N} (where N is the DIMM number)
- HTTP Method: GET

The figure below is one example of DIMM 23 which contains one record of DimmTopFaults and DimmHealthStatus.




Figure 5. Using Redfish to GET DIMM Health Status and Top Faults

DIMM recovery information

The Redfish command for DIMM recovery information of the system can be used to fetch all memory recovery entries (including row recovery with runtime soft PPR recovery, row recovery with POST hard PPR recovery) which have been applied to DIMMs of the system.

The URI and Method are as follows:

- Redfish URI: `/redfish/v1/Systems/1/Memory/Oem/Lenovo/DimmRecoveryReport`
- HTTP Method: GET

The following figure shows one example of DIMM 23 which contains one hard PPR recovery entry.




Figure 6. Using Redfish to GET System DIMM Recovery Report

Summary

Lenovo MPFA feature is a crucial component in Lenovo's commitment to providing advanced, reliable, and high-performance server solutions for Lenovo ThinkSystem V3 and V4 servers. By anticipating and mitigating memory failures, MPFA helps enhance server reliability and performance by proactively addressing memory failures, ensuring optimal operation and reducing downtime. And this feature is planned to be supported in future generations of Lenovo ThinkSystem servers.

Resources

For more information, see the following resources:

- Intel Memory Resilience Technology (also known as Intel Memory Failure Prediction)
<https://www.intel.com/content/www/us/en/software/intel-memory-resilience-technology.html>
- AMD Advanced Platform Management Link (APML) Library
<https://www.amd.com/en/developer/e-smss/apml-library.html>
- Lenovo ThinkSystem Memory RAS Introduction
https://download.lenovo.com/servers_pdf/thinksystem_memory_ras_intro.pdf
- Redfish spec in DMTF
<https://www.dmtf.org/standards/redfish>

Abbreviations and terms

The following table lists relevant terms used in this document.

Table 2. Abbreviations and terms

Term	Meaning
ADDDC	Adaptive Double Device Data Correction – Intel-defined memory RAS feature which is implemented by adaptive virtual lockstep concept to map out one faulty region
APML	Advanced Platform Management Link – out of band interface for AMD EPYC telemetry.
ECS	Error Check and Scrub – DDR5 ECC monitoring that provides Error Counter (EC) and Error per Row Counter (EpRC).
FFDC	First Failure Data Capture – comprehensive BMC log used for post mortem analysis.
PECI	Platform Environment Control Interface – Intel defined out of band bus for processor BMC communication.
PPR	Post Package Repair – self healing memory technique (soft sPPR, hard hPPR).
RAS	Reliability, Availability, Serviceability – holistic metric of system robustness.
Redfish	DMTF standardized, RESTful management interface for remote monitoring and configuration.
SCI	System Control Interrupt – OS level interrupt used by UEFI to request page retirement.
SMI	System Management Interrupt – high priority, non maskable interrupt that drives firmware level error handling.

Authors

Jason (Zhijun) Liu is a Principal Engineer and Senior UEFI Architect at Lenovo Infrastructure Solutions Group. Jason provides high-level infrastructure design support for Lenovo ThinkSystem UEFI firmware and leads the enabling, customization and innovation of new technologies into UEFI firmware. Jason also leads Reliability, Availability and Serviceability (RAS) architecture design and Secure feature design for ThinkSystem firmware.

Hao Chen is a Senior Researcher leading the data processing and analytics at Lenovo Research, specializing in AI-driven predictive maintenance (PFA) and AI innovation for enterprise system management platforms such as XClarity Products. He leads the development of applying machine learning and deep learning-based predictive fault analytics across memory, drive, GPU and storage subsystems to enable early risk detection and proactive service automation at scale.

Da Li is a Senior Engineer and Senior XCC Technical Lead at Lenovo Infrastructure Solutions Group. Da leads XCC common function design and support For Lenovo ThinkSystem servers. Da also leads power management, RAS design in XCC side for for ThinkSystem firmware.

Joseph (Fuzhou) Liu is a Senior UEFI Development Engineer at Lenovo Infrastructure Solutions Group. He focuses on Reliability, Availability, and Serviceability (RAS) feature enabling, enhancement and customer support in UEFI firmware for both Intel and AMD based ThinkSystem servers.

Thanks to the following people for their contributions to this project:

- Scott Harsany
- Scott Faasse
- Chuang Zhang
- Jerry (Xinjie) Nie
- Gavin (Gaofeng) Zhang
- Minchao Zhao
- Rui Ma
- Shuangqing Zhang
- Bill Zevin
- Sumanta Bahali

- Benjamin Ming Lei
- Rob Tahamtan
- Paul Klustaitis

Related product families

Product families related to this document are the following:

- [Memory](#)

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2025. All rights reserved.

This document, LP2326, was created or updated on October 31, 2025.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP2326>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP2326>.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®

ThinkSystem®

XClarity®

The following terms are trademarks of other companies:

AMD and AMD EPYC™ are trademarks of Advanced Micro Devices, Inc.

Intel® and Xeon® are trademarks of Intel Corporation or its subsidiaries.

Other company, product, or service names may be trademarks or service marks of others.