

RAS Features of the Lenovo ThinkSystem Intel Servers

Article

Applications such as database, enterprise resource planning (ERP), customer resource management (CRM), and business intelligence (BI) applications need to be available 24/7 on a wide area or global basis. In addition, the likelihood of such failures increases statistically with the size of the servers, data, and memory required for these deployments.

While clustering and virtualization can help meet availability requirements, they are not adequate solutions for very large databases, BI, and high-end transactional systems. A failure affecting a single core business application can easily cost hundreds of thousands or even millions of dollars per hour. All this leads to a need for scalable and highly resilient servers that are well suited for critical business applications and large-scale consolidation.

Always On

Time is money. Even a few minutes of downtime can result in significant costs and cause internal business operations to come to a standstill. Downtime can also adversely impact a company's relationship with its customers, business suppliers and partners. Reliability or lack thereof can potentially damage a company's reputation and result in lost business.

The growth of new applications has ratcheted database processing and business analytics to the top of the list for server workloads. These workloads demand continuous availability from the enterprise platforms on which they run.

"Always on" has become a global requirement and impacts many aspects of our lives:

- **Maximize productivity** - Manufacturers need to keep their production line up and running. System downtime should not interrupt it.
- **Control access** - Facility Security companies prevent external threats to organizations. Security application downtime shouldn't be an internal threat.
- **Protect profit** - Retailers have sales targets to meet day in, day out. Transaction system downtime shouldn't get in the way.
- **Protect lives** - First Responders take care of emergencies 24 x 7 x 365. Application downtime shouldn't be one of them.
- **Ensure quality care and privacy** - Healthcare Institutions need to access patient information and be HIPPA compliant all the time. System downtime shouldn't compromise either one.
- **Process transactions** - Financial Services organizations manage thousands of transactions a second. Processing system downtime simply can't happen

Server RAS Defined

RAS in relation to servers is defined as follows:

Reliability – Reducing the mean time between hardware failures and ensuring data integrity. Data integrity is protected through error detection and correction — or, if not correctable, error containment

- Error Detection and Self-Healing
- Minimizes outage opportunities
- Correct results continuously

Availability – Refers to uninterrupted system and application operation even in the presence of uncorrectable errors

- Reduce frequency and duration of outages
- Self-diagnosing: work around faulty components or “self-heal”
- Never stops or slows down

Serviceability – Means a system can be maintained without disrupting operation. This capability requires both thoughtful platform design and innovative systems management.

- Avoid repeat failures with accurate diagnostics
- Concurrent repair on higher failure rate items
- Easy to repair and upgrade

Industry Cost of Downtime

90% of Companies surveyed say that 1 hour of downtime costs exceed \$300K.

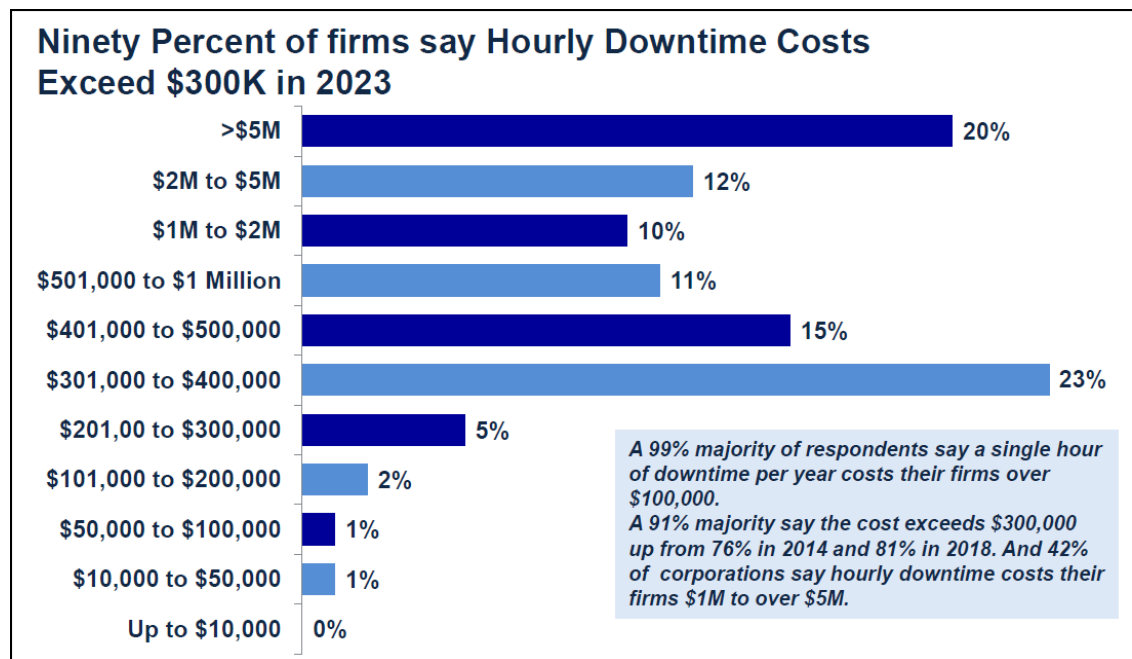


Figure 1. Hourly Cost of Downtime (from ITIC 2022-2023 Global Server Hardware reliability & Server Security Survey Results)

Unplanned Downtime by Hardware Platform

According to ITIC, Lenovo ThinkSystem and IBM lead the way for unplanned downtime

- Lenovo ThinkSystem: 1.1 minutes
- Cisco UCS: 2 minutes
- Dell PowerEdge: 26 minutes
- HPE ProLiant: 39 minutes




Figure 2. Unplanned Downtime by Server Hardware Platform in January 2023 (from ITIC 2022-2023 Global Server Hardware reliability & Server Security Survey Results)

RAS features with Intel Xeon Scalable processors

The following table is a list of Key RAS Features of the Intel Xeon Scalable processors on the Lenovo ThinkSystem servers.

Table 1. RAS features of Lenovo ThinkSystem servers with Intel Xeon Scalable processors

Feature	Category	2nd Gen Intel (CLX)	3rd Gen Intel 4S (CPX)	3rd Gen Intel 2S (ICX)	4th Gen Intel (SPR)	Intel Xeon Max Series (SPR HBM)	5th Gen Intel (EMR)	Benefit
Advanced RAS features								
Viral Mode of error containment	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	Enhanced error containment to improve data integrity, complimentary to corrupt data containment mode
Local Machine Check (LMCE) based Recovery	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	Enhances MCA recovery-Execution path event, and increases the possibility of recovery

Feature	Category	2nd Gen Intel (CLX)	3rd Gen Intel 4S (CPX)	3rd Gen Intel 2S (ICX)	4th Gen Intel (SPR)	Intel Xeon Max Series (SPR HBM)	5th Gen Intel (EMR)	Benefit
SDDC +1, Adaptive DDDC (ADDDC) (MR) +1*	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	Adaptive virtual lockstep delivers up to two DRAM Device sub-region(s) sparing/correction at bank and/or full rank granularity. Also supports Single DRAM correction, as well as single bit correction post final DRAM device map out.
PCI Express Live Error Recovery	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	PCI-e root port error containment, and the opportunity to dynamically recover from the error
Intel® UPI Dynamic Link width reduction	Availability	Yes	Yes	Yes	Yes	Yes	Yes	Enables interconnect to continue operation in presence of Interconnect link persistent failure
Address range/Partial Memory Mirroring	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	OS managed memory mirroring of selective ranges, increases data integrity at efficient cost
MCA 2.0 Recovery (as per eMCA gen2 architecture)	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	Firmware first model enables a reliable error sourcing capability with the ability to write to the MSR
Standard RAS features								
Advanced Error Detection and Correction (AEDC)	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	Enhanced fault coverage within processor cores, and attempt to recover via instruction retry
Error Detection and Correction	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	Extensive Error detection and correction capability across the silicon, and the interconnects.
Corrupt Data containment-Core	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	Uncorrectable data explicitly marked and delivered synchronously to the consuming core to assist error containment and increase system reliability
Corrupt Data containment-UnCore	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	Uncorrectable data explicitly marked and delivered synchronously to the requestor, to assist error containment and increase system reliability
SDDC, Adaptive Data Correction (ADC) (SR)*	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	Adaptive virtual lockstep delivers single DRAM Device sub-region sparing/correction at bank granularity. Also supports Single DRAM correction.
PCIe "Stop and Scream"	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	PCI-e root port corrupt data containment feature, increases data integrity
Memory Mirroring-Intra iMC	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	Increase data integrity by creating a redundant/mirrored copy of data in system DRAM
Rank/Multi Rank Sparing & DRAM sparing	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	Reserved/spare ranks are utilized to dynamically map out the failing rank into the spare ranks. Starting with 3rd Gen "Ice Lake" processors, this is implemented using ADDDC/ ADC-SR/ ADDDC-MR to provide DRAM-level sparing feature support.
Predictive Failure Analysis	Serviceability	Yes	Yes	Yes	Yes	Yes	Yes	Extensive error logs to assist software in predicting failures

Feature	Category	2nd Gen Intel (CLX)	3rd Gen Intel 4S (CPX)	3rd Gen Intel 2S (ICX)	4th Gen Intel (SPR)	Intel Xeon Max Series (SPR HBM)	5th Gen Intel (EMR)	Benefit
Failed DIMM Isolation	Serviceability	Yes	Yes	Yes	Yes	Yes	Yes	Extensive error logs to help software identify the failing DIMM
Virtual (soft) Partitioning	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	Virtual Machine Monitor ability to make use of hardware recovery , signaling and error logs
Error reporting via IOMCA	Serviceability	Yes	Yes	Yes	Yes	Yes	Yes	Unified error reporting of the IIO logic to the OS
Error reporting through MCA 2.0 (eMCA gen2)	Serviceability	Yes	Yes	Yes	Yes	Yes	Yes	Firmware first model enables a reliable error sourcing capability
Error reporting through eMCA gen1	Serviceability	Yes	Yes	Yes	Yes	Yes	Yes	Firmware first model enables reliable error sourcing capability
PCIe Card Hot Plug NVMe (Add, Remove, Swap)	Serviceability	Yes	Yes	Yes	Yes	Yes	Yes	Hot add and repalce of NVMe drives
PCI Express ECRC	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	PCI Express End to end CRC checking, increasing system reliability
PCIe Corrupt Data Containment (Data Poisoning)	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	PCIe corrupt data mode of operation, synchronous signaling of the corrupted data along with data, increases system reliability
PCIe Link CRC Error Check and Retry	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	PCIe link CRC error check and retry, system reliability and recovery from transient errors
PCIe Link Retraining and Recovery	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	PCIe link retraining and attempted recovery from persistent link transient errors
Mem SMBus hang recovery	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	Software ability to reset memory SMBus interface to recover from hang condition
DDR4 Command/ Address Parity Check and Retry	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	DDR4 Address and command parity check and retry in the event of errors
Time-out timer Schemes	Serviceability	Yes	Yes	Yes	Yes	Yes	Yes	Hierarchy of transaction time outs to assist system debug and reliable error sourcing.
Intel® UPI Link Level Retry	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	Intel UPI link's ability to perform CRC check and retry on errors for higher degree of system reliability
Intel® UPI Protocol Protection via 16 bit Rolling CRC	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	Detection of transient data errors over Intel UPI interconnects, via 16bit CRC error checking
Processor BIST	Serviceability	Yes	Yes	Yes	Yes	Yes	Yes	At power up, Processor's built in self test engine performs test on the internal cache structure for and provides the results to the system BIOS
Socket disable for FRB	Availability	Yes	Yes	Yes	Yes	Yes	Yes	The capability to selectively disable socket at the boot time, and therefore allowing system to power-on in a failover configuration
Core disable for FRB	Availability	Yes	Yes	Yes	Yes	Yes	Yes	The capability to disable failing cores at boot time, map out the failing core

Feature	Category	2nd Gen Intel (CLX)	3rd Gen Intel 4S (CPX)	3rd Gen Intel 2S (ICX)	4th Gen Intel (SPR)	Intel Xeon Max Series (SPR HBM)	5th Gen Intel (EMR)	Benefit
Memory disable for FRB	Availability	Yes	Yes	Yes	Yes	Yes	Yes	The capability to disable failing DIMMs at boot time, map out the failing DIMMs
Memory demand and patrol Scrubbing	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	DRAM content is scrubbed with the corrected data due to Demand or patrol scrub operation. Scrubbing DRAM location can prevent accumulation of single-bit errors turning it into uncorrected error.
DDR Power Up Post Package Repair (PPR)	Availability	Yes	Yes	Yes	Yes	Yes	Yes	The capability will map out what platform identifies as bad rows with spare rows in DDR DIMMs, and the repair action will be executed during memory training phase at system power up.
PIROM for System Information Storage	Serviceability	Yes	Yes	Yes	Yes	Yes	Yes	On package Processor Information ROM
MCA Recovery-Execution path	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	OS layer assisted recovery from uncorrectable data errors to prevent system reset
MCA Recovery-Non execution path	Reliability	Yes	Yes	Yes	Yes	Yes	Yes	OS layer assisted recovery from uncorrectable data errors detected by Patrol scrubber or LLC Explicit Write Back
DCU Scrubbing	Reliability	No	No	Yes	Yes	Yes	Yes	Improves system uptime by minimizing the impact due to high energy particle strike (Soft Errors) within core DCU (aka L1D cache)
Partial Cache Line Sparing (PCLS)	Reliability	No	No	Yes	No	Yes	No	Extend system uptime in case a single bit memory persistent error is detected. It allows mapping out the cache line with failed single bit by using spare capacity available within the IMC
PCIe Enhanced Downstream Port Containment (EDPC)	Reliability	No	No	Yes	Yes	Yes	Yes	EDPC is an enhancement to the Downstream Port Containment (DPC) thereby adding Root Port Programmable IO (RPPIO) errors.
DCU/IFU Error Handling Enhancement	Reliability	No	No	Yes	Yes	Yes	Yes	Improving MCA Recovery (Execution path) coverage in case CPU core (DCU/IFU) receives multiple back-to-back read data with 'Poison Error' (aka Poison Storm).
Memory Permeant Fault Detection (PFD)	Reliability	No	No	No	Yes	Yes	Yes	Memory controller will use algorithm to detect and report permanent memory fault
DDR5 On-die Error Check and Scrub	Reliability	No	No	No	Yes	Yes	Yes	DDR5 has on-die ECC Error Check and Scrub (ECS) mode with an error counting scheme for transparency. The ECS mode allows the DRAM to internally read, correct single bit errors, and write back corrected data bits to the array (scrub errors) while providing transparency to error counts.

Feature	Category	2nd Gen Intel (CLX)	3rd Gen Intel 4S (CPX)	3rd Gen Intel 2S (ICX)	4th Gen Intel (SPR)	Intel Xeon Max Series (SPR HBM)	5th Gen Intel (EMR)	Benefit
DDR5 PMIC Error Handling	Reliability	No	No	No	Yes	Yes	Yes	DDR5 has Power Management IC (PMIC) on DIMM, this capability will isolate the DIMM with PMIC error.
RAS features for processors with High Bandwidth Memory (HBM)								
HBM Only mode, 1LM, 2LM	Availability	N/A	N/A	N/A	N/A	Yes	N/A	High Bandwidth Memory (HBM), it supports HBM Only mode (no DDR DIMMs in the system), 1LM mode (1 level memory: HBM acts same as DDR), and 2LM (2 levels memory: HBM acts as cache of DDR).
HBM Permeant Fault Detection (PFD)	Reliability	N/A	N/A	N/A	N/A	Yes	N/A	Memory controller will use algorithm to detect and report HBM permanent fault.
HBM Command/Address parity check and retry	Reliability	N/A	N/A	N/A	N/A	Yes	N/A	Memory controller will retry the transaction and try to recovery at CMD/ADDR parity errors on the CMD/ADDR bus.
HBM Power Up PPR	Availability	N/A	N/A	N/A	N/A	Yes	N/A	The capability will maps out what platform identifies as bad rows with spare rows in HBM memory, and the repair action will be executed during HBM training phase at system power up.
HBM Demand and Patrol Scrub	Reliability	N/A	N/A	N/A	N/A	Yes	N/A	Similar to DDR, HBM DRAM content is scrubbed with the corrected data due to Demand or patrol scrub operation. Scrubbing DRAM location can prevent accumulation of single-bit errors turning it into uncorrected error
HBM Memory bank sparing	Reliability	N/A	N/A	N/A	N/A	Yes	N/A	Legacy Bank sparing operation, each pseudo channel of HBM supports spare bank for UEFI to make use of after it detects bank failing on the pseudo channel.
HBM Corrected Error Reporting	Reliability	N/A	N/A	N/A	N/A	Yes	N/A	Memory Corrected Error reporting supports per rank corrected error counters with leaky bucket algorithm.
HBM Disable/Map out for FRB	Availability	N/A	N/A	N/A	N/A	Yes	N/A	The capability to disable failing HBM at boot time, map out the failing HBM.

* See additional information about ADDDC below

All ThinkSystem servers with Intel Xeon Scalable processors support the ADC (SR or Single Region) and ADDDC (MR or Multiple Region) feature which can support DRAM sub-region(s) sparing at Bank and/or full Rank granularity, compared to the memory rank sparing feature.

Details of Adaptive Double DRAM Device Correction (ADDDC) from the Intel article "New Reliability, Availability, and Serviceability (RAS) Features in the Intel® Xeon® Processor Family"

<https://www.intel.com/content/www/us/en/developer/articles/technical/new-reliability-availability-and-serviceability-ras-features-in-the-intel-xeon-processor.html>

Intel Xeon introduces an innovative approach in managing errors that the DDR4 DRAM DIMM may induce through the life of the product. ADDDC is deployed at runtime to dynamically map out the failing DRAM device and continue to provide SDDC ECC coverage on the DIMM, translating to longer DIMM longevity. The operation occurs at the fine granularity of DRAM Bank and/or Rank to have minimal impact on the overall system performance.

With the advent of ADDDC, the memory subsystem is always configured to operate in performance mode. When the number of corrections on a DRAM device reaches the targeted threshold value, with help from the UEFI runtime code, the identified failing DRAM region is adaptively placed in lockstep mode where the identified failing region of the DRAM device is mapped out of ECC. Once in ADDDC, cache line ECC continues to cover single DRAM (x4) error detection and apply a correction algorithm to the nibble.

Dependent on the processor SKU, each DDR4 channel supports one to two regions that can manage one or two faulty DRAMs, at Bank and/or full Rank granularity. The dynamic nature of the operation makes the performance implications of the lockstep operation on the system to be material only after the DRAM device is detected to be failing. The overall lockstep impact on system performance is now a function of the number of bad DRAM devices on the channel, with the worst-case scenario of two bad Ranks on every DDR4 channel.

The Silver/Bronze SKUs offer Adaptive Data Correction (ADC [SR]), at Bank granularity, and the Platinum/Gold SKUs offer Adaptive Double DRAM Device Correction (ADDDC [MR]), at Bank and Rank granularity, with additional hardware facilities for device map-out.

Further reading

For further analysis of the reliability of Lenovo servers, see the latest ITIC report, available from:

ITIC 2022-2023 Global Server Hardware Reliability & Server Security Survey Results:

<https://lenovopress.lenovo.com/lp1117-itic-reliability-study>

This article is one in a series on the ThinkSystem V3 servers:

- [Five Highlights of the Lenovo ThinkSystem SR630 V3 Server](#)
- [Five Highlights of the Lenovo ThinkSystem SR650 V3 Server](#)
- [Five Highlights of the Lenovo ThinkSystem SR850 V3 Server](#)
- [Five Highlights of the Lenovo ThinkSystem SR860 V3 Server](#)
- [RAS Features of the Lenovo ThinkSystem Intel Servers](#)

About the author

Randall Lundin is a Senior Product Manager in the Lenovo Infrastructure Solution Group. He is responsible for planning and managing ThinkSystem servers. Randall has also authored and contributed to numerous Lenovo Press publications on ThinkSystem products.

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 2024. All rights reserved.

This document, LP1711, was created or updated on April 22, 2024.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP1711>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP1711>.

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®

The following terms are trademarks of other companies:

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.