



NVIDIA GPU Memory Error Management

Release 615

NVIDIA Corporation

Sep 09, 2026

Contents

1 Overview	3
2 Supported GPUs	5
3 Error Containment	7
4 Dynamic Page Offlining	9
5 Row Remapping	11
6 RAS Repair	13
6.1 GPU Memory Repair	13
7 Response to Uncorrectable Contained ECC Errors	15
8 Error Recovery and Response Flags	17
9 User Visible Statistics	19
10 RMA Policy	23
10.1 GPU DRAM Memory RMA Policy	23
10.2 GPU L2 Memory RMA Policy	23
11 SRAM Uncorrectable Errors	25
12 Notices	27
12.1 Trademarks	28

This document describes the new memory error recovery features introduced in the NVIDIA® 100 GPU and NVIDIA 800 GPU.

Chapter 1. Overview

NVIDIA® Ampere, Hopper, Blackwell architecture introduce new memory error recovery features that improve resilience and avoid impacting unaffected applications. The new features improve various aspects of the graphics processing units' (GPU) response to memory errors, which improves the overall robustness of the error handling and recovery process.

The error handling and response features are:

- ▶ Error-containment
- ▶ Dynamic page offlining
- ▶ Row-remapping
- ▶ RAS Repair for GPU Memory
- ▶ Uncorrectable error to correctable error coverage improved by 10%
- ▶ ECC correction of single bit errors (SBEs)

When referring to ECC errors in this application note, we focus on uncorrectable high bandwidth memory (HBM) memory errors. SRAM failure modes are discussed in RMA Policy Thresholds for SRAM Failure Modes. Correctable HBM errors are outside the scope of this application note.

Chapter 2. Supported GPUs

Error-containment and row-remapping are two separate features in the GPU architecture. The following table shows the GPUs that support these features.

Table 1: Supported GPUs

	GA100	GA10x	Ada AD10x	GH100	GB10x
Error Containment	X			X	X
Row remapping	X	X	X	X	X
Dynamic Page Offlining	X			X	X
RAS Repair – GPU Memory					X*

* Support is not guaranteed for all Blackwell GPUs.

Chapter 3. Error Containment

The NVIDIA Ampere architecture introduces the concept of error containment to NVIDIA GPUs. The benefit of error containment is being able to limit the impact of uncorrectable ECC errors on GPU applications. Uncorrectable ECC errors on prior architectures such as NVIDIA Volta™ impacted all of the currently executing GPU workloads. On NVIDIA data center-class GPUs starting with Ampere architecture the impact will be limited to the applications that encounter the error. All other workloads will continue running unaffected both in terms of accuracy and performance, and new workloads can be launched. Unlike earlier GPU architectures, NVIDIA Ampere and later generation class GPUs do not require a GPU reset when memory errors occur.

Note

While most frequently occurring classes of uncorrectable errors are contained, there can be rare cases where uncorrectable errors are still uncontained and might impact all the workloads being processed in the GPU.

Chapter 4. Dynamic Page Offlining

Dynamic Page Offlining improves resiliency and availability of NVIDIA GPUs to uncorrectable ECC errors. Once the NVIDIA driver identifies the location of an uncorrectable error in the frame buffer memory, it marks the page containing the error as unusable. Once the page is marked unusable, any of the currently executing or newly launched workloads will not be allocating this page in question.

Dynamic Page Offlining exists on NVIDIA GPUs starting with the NVIDIA Ampere architecture. It is not available on previous generations of NVIDIA GPUs that do not support error containment.

GPUs that support dynamic page offlining do not require a GPU reset to recover from most uncorrectable ECC errors.

After the page is marked as unusable, it will not be mapped to the address space of any currently running or newly launched CUDA kernels.

Chapter 5. Row Remapping

Row-remapping is a hardware mechanism to improve the reliability of frame buffer memory on GPUs starting with the NVIDIA Ampere architecture. This feature is used to prevent known degraded memory locations from being used. The row-remapping feature is a replacement for the page retirement scheme used in prior generation GPUs. Every bank in DRAM is equipped with spare rows in hardware. As opposed to traditional page retirement, the row-remapper replaces degrading memory cells with spare ones to avoid offlining regions of memory in software. This differs from dynamic page offlining in that the memory is fixed at a hardware level and does not leave software visible holes in the address space. The process of row-remapping requires a GPU reset to take effect and will remain persistent throughout the life of the GPU.

The following table describes the differences between page retirement and row-remapping.

Table 2: Page Retirement vs. Row-Remapping

Feature	Page Retirement for Legacy GPUs	Row-Remapping for NVIDIA Ampere & Later Generation
Available remappings/retirements	Supported a maximum of 64 retirements for the frame buffer	Supports up to 512 remapping for the frame buffer.
Policy changes	Once a retirement takes effect, the page can never be unretired, regardless of correctable or uncorrectable errors	Remapping due to correctable errors can be replaced by uncorrectable error remapping when the memory bank's reserved rows are exhausted.
RMA criteria	A threshold of page retirements on a GPU usually resulted in investigation of whether the GPU was worthy of an RMA	Refer to RMA Policy .
Application of pending changes	Needed a kernel module reload or driver re-initialization or GPU reset	GPU reset is required.

Chapter 6. RAS Repair

6.1. GPU Memory Repair

GPU memory repair is another resiliency feature available on select Blackwell products. GPU memory repair consists of swapping spare DRAM channels or L2 cache slices.

There are multiple L2 slices that make up L2 cache in the NVIDIA GPU.

For example, if there are spare channels available and the DRAM has a bank that is trending towards failure, the channel in which the bank resides can be swapped out for a spare. After two row remappings in the same bank, the next occurrence of an uncorrectable ECC error in that bank will attempt to trigger repair if a spare is available. The same concept applies to failing L2 slice as well.

XID 160 will be printed if a DRAM channel or L2 slice is successfully marked for repair. The process of channel repair requires a reboot to take effect.

Chapter 7. Response to Uncorrectable Contained ECC Errors

Like previous GPU architectures, when an uncorrectable ECC error is detected, the NVIDIA driver software will perform error recovery. Error containment ensures that erroneous data does not continue to propagate, and the affected application is terminated.

- Uncorrectable contained ECC error are uncorrectable ECC errors where error containment process was successful.
- Uncorrectable uncontained ECC error are uncorrectable ECC errors where error containment process was not successful.

Dynamic page offlining marks the page containing the faulty memory as unusable. This ensures that new allocations do not land on the page that contains the faulty memory. Unaffected applications will continue to run and additional workloads can be launched on this GPU without requiring a GPU reset.

When GPU reset occurs as a part of the regular GPU/VM service window, row remapping fixes the memory in hardware without creating any holes in the address space and the offlined page is reclaimed.




Figure 1: NVIDIA GPU Response to Uncorrectable Contained ECC Error

Chapter 8. Error Recovery and Response Flags

Here is a list of flags that are required for client error recovery and response:

- ▶ Row-remapping pending flag
 - ▶ This flag indicates that row-remapping will happen at the next GPU reset.
 - ▶ Even with the flag set, unaffected applications can continue running without affecting accuracy and performance, and new workloads can be launched.
 - ▶ This is useful to identify readiness for live virtual machine (VM) migrations into this GPU, and this GPU should be reset if a live VM will be migrated to it.
- ▶ Row-remapping failure flag
 - ▶ For definition of row-remapping failure, see RMA Policy Thresholds for Row-Remapping.
- ▶ Recovery action flag
 - ▶ Recovery action flag will be triggered for uncorrectable errors. When set, this flag indicates the following user actions:
 - ▶ None (NVML_GPU_RECOVERY_ACTION_NONE)
 - ▶ No recovery action is needed. The GPU can be immediately used again by starting a new GPU process.
 - ▶ Reset (NVML_GPU_RECOVERY_ACTION_GPU_RESET)
 - ▶ The GPU has encountered a fault that requires a reset to recover. Terminate all GPU processes, reset the GPU using “nvidia-smi -r”, and the GPU can be used again by starting new GPU processes.
 - ▶ Reboot (NVML_GPU_RECOVERY_ACTION_NODE_REBOOT)
 - ▶ The GPU has encountered a fault may have left the OS in an inconsistent state. Reboot the operating system to restore the OS back to a consistent state.
 - ▶ Drain P2P (NVML_GPU_RECOVERY_ACTION_DRAIN_P2P)
 - ▶ The GPU has encountered a fault that requires all peer-to-peer traffic to be quiesced. Terminate all GPU processes that conduct peer-to-peer traffic and disable UVM persistence mode. Once all peer-to-peer traffic are drained, query NVML_FI_DEV_GET_GPU_RECOVERY_ACTION again, which will return one of the other actions.
 - ▶ Drain and Reset Flag (NVML_GPU_RECOVERY_ACTION_DRAIN_AND_RESET)

- ▶ The GPU has encountered a fault that results the GPU to temporarily operate at a reduced capacity, such as part of its frame buffer memory being offlined, or some of its MIG partitions down. No new work should be scheduled on the GPU, but existing work that didn't get affected are safe to continue until they finish or reach a good checkpoint. After all existing work have drained, reset the GPU to regain its full capacity.

Note

These client flags are currently exposed through SMBPBI.

Chapter 9. User Visible Statistics

Previously, end-users or sysadmins could use the page retirement count to monitor the health of the GPU (and possible RMA conditions) and determine whether to reset the GPU or reload the module for page offlining to take effect. For the same purpose, row-remapping statistics will be exposed to users to give an indication of the health of the GPU memory. This section describes the row-remapping statistics that are available via in-band and out-of-band reporting mechanisms.

- ▶ In-band reporting
 - ▶ XID error log (see Table 2 for a list of XID log examples for uncorrectable ECC errors)
 - ▶ XID 48: This XID indicates an uncorrectable ECC error has occurred. XID 171 indicates if it is DRAM and XID 172 indicates if it is SRAM
 - ▶ XID 94: This XID indicates a contained ECC error has occurred
 - ▶ XID 95: This XID indicates an uncontained ECC error has occurred
 - ▶ XID 63: This XID indicates successful recording of a row-remapping entry to the InfoROM
 - ▶ XID 64: This XID indicates a failure in recording a row-remapping entry to the InfoROM
 - ▶ XID 160: This XID indicates that GPU memory is successfully marked for repair.
 - ▶ NVML/nvidia-smi
 - ▶ Number of remapped rows (correctable and uncorrectable). This is the number of entries recorded in the InfoROM, not the ones remapped in hardware.
 - ▶ Row remapping pending Boolean
 - ▶ Row remapping failure Boolean
 - ▶ Bucketized counts
 - ▶ Refer to <https://docs.nvidia.com/deploy/nvml-api/index.html> for more information about NVML.
 - ▶ Refer to <https://developer.nvidia.com/nvidia-system-management-interface> for more information about nvidia-smi.
- ▶ Out-of-band reporting (SMBPBI)
 - ▶ Number of remapped rows (correctable and uncorrectable), and this is the number of entries recorded in the InfoROM, not the ones remapped in hardware.
 - ▶ Row remapping pending Boolean
 - ▶ Row remapping failure Boolean
 - ▶ Bucketized counts

- Table 3 lists the new SMBPBI APIs for error reporting. For additional details refer to the NVIDIA SMBus Post-Box Interface (SMBPBI) Software Design Guide (DG-06034-002).

Table 3: Uncorrectable ECC Errors XID Log Examples

Error Type	XID Log
Contained error with MIG enabled	NVRM: Xid (PCI:0000:01:00 GPU-I:05): 94, pid=7194, Contained: CE User Channel (0x9). RST: No, D-RST: No
Contained error with MIG disabled	NVRM: Xid (PCI:0000:01:00): 94, pid=7062, Contained: CE User Channel (0x9). RST: No, D-RST: No
Uncontained error	NVRM: Xid (PCI:0000:01:00): 95, pid=7062, Uncontained: LTC TAG (0x2,0x0). RST: Yes, D-RST: No

Table 4: SMBPBI APIs for NVIDIA A100 Memory Error Reporting

Opcode	Description
0x1E	Request ECC statistics (format V6)
0x20	Request row-remapping related statistics

Table 5 shows an example of bucketized count where all remapping resources are available.

Table 5: Bucketized Count Example #1

Bucket Name	Row Remap Availability (shown as reference only)	Number of Banks per Bucket
Max	8	640
High	7	0
Partial	2 to 6	0
Low	1	0
None	0	0

Note

The Bank Remap Availability column is shown only for reference purposes. The information is not available through APIs or nvidia-smi.

Table 6 shows an example of bucketized count where 635 banks have 8, 3 banks have 7, and 2 banks have one remaining remapping resource available.

Table 6: Bucketized Count Example #2

Bucket Name	Row Remap Availability (shown as reference only)	Number of Banks per Bucket
Max	8	635
High	7	3
Partial	2 to 6	0
Low	1	2
None	0	0

Note

The Bank Remap Availability column is shown only for reference purposes. The information is not available through APIs or `nvidia-smi`.

Chapter 10. RMA Policy

10.1. GPU DRAM Memory RMA Policy

The NVIDIA Field Diagnostic tool determines whether a GPU qualifies for RMA. Regarding row-remapping failures, the RMA criteria is met when the row-remapping failure flag is set and validated by the field diagnostic.

For Blackwell GPUs, a third remapping attempt for an uncorrectable memory error on a bank invokes HBM channel repair if there is a spare channel available. If the channel repair is successful, the hardware will be recovered and ready to use. Otherwise, the GPU will continue remapping until Row-Remap Failure flag is set.

Any of the following events will trigger a row-remapping failure flag:

- ▶ A remapping attempt for an uncorrectable memory error on a bank that already has eight uncorrectable error rows remapped.
- ▶ A remapping attempt for an uncorrectable memory error on a row that was already remapped and can occur with less than eight total remaps to the same bank.
- ▶ After 512 total remappings for an uncorrectable memory error have occurred.

The row-remapping failure flag is available through in-band (NVML/nvidia-smi) and out-of-band (SMBPBI) tools.

10.2. GPU L2 Memory RMA Policy

The NVIDIA Field Diagnostic tool will determine whether a GPU qualifies for RMA. Regarding SRAM uncorrectable errors, the RMA criteria is met for events outlined below.

Any of the following events will trigger the SRAM Threshold Exceeded flag:

- ▶ More than 4 UCE Unique Count events within an address bank for parity protected SRAMs.
- ▶ More than 2 UCE Unique Count events within an address bank for SECDED ECC protected SRAMs.

Chapter 11. SRAM Uncorrectable Errors

The NVIDIA Field Diagnostic tool will determine whether a GPU qualifies for RMA. Regarding SRAM uncorrectable errors, the RMA criteria is met for events outlined below. Any of the following events will trigger the SRAM Threshold Exceeded flag:

- ▶ More than 4 UCE Unique Count events within an address bank for parity protected SRAMs.
- ▶ More than 2 UCE Unique Count events within an address bank for SECDED ECC protected SRAMs.

Chapter 12. Notices

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation (“NVIDIA”) makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer (“Terms of Sale”). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer’s own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer’s product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

12.1. Trademarks

NVIDIA and the NVIDIA logo are trademarks or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

©2021-2026, NVIDIA Corporation & affiliates. All rights reserved