



# Solidigm™ D7-P5520/D7-P5620 (Formerly Intel®)

Sightings Report

July 2024

Revision 011

Solidigm Confidential

# SOLIDIGM™

## Revision History

Revision	Description	Revision Date
001	<ul style="list-style-type: none"> <li>Initial Release</li> </ul>	December 2021
002	<ul style="list-style-type: none"> <li>Added new sightings: NSGSE-203817, NSGSE-205704, NSGSE-207000, NSGSE-208801, NSGSE-208892</li> <li>Moved NSGSE-195807 from Open - Intend to Fix, to Closed – Fixed with Firmware/Bootloader Change</li> <li>Moved NSGSE-195579 from Open – Under Investigation, to Closed – No Plan to Fix</li> </ul>	January 2022
003	<ul style="list-style-type: none"> <li>Moved NSGSE-200108, NSGSE-203817, NSGSE-207000, NSGSE-208801 and NSGSE-208892 from Open – Under Investigation to Closed - Fixed with Firmware/Bootloader Change</li> <li>Moved NSGSE-192480 and NSGSE-202068 from Open – Under Investigation to Closed – No Plan to Fix</li> <li>Added NSGSE-198625 and NSGSE-201063 under Open – Under Investigation</li> </ul>	March 2022
004	<ul style="list-style-type: none"> <li>Added NSGSE-212334 under Open – Under Investigation</li> <li>Added NSGSE-205946 under 3<sup>rd</sup> Party Sightings</li> <li>Updated Firmware Release Matrix</li> </ul>	March 2022
005	<ul style="list-style-type: none"> <li>Added NSGSE-211637 to Open – Under Investigation</li> <li>Moved NSGSE-205704 from Open – Under Investigation to Closed – No Plan to Fix</li> <li>Moved NSGSE-198265 from Open – Under Investigation to Open – Intend to Fix</li> <li>Moved NSGSE-211219 from Open – Under Investigation to Closed – Fixed with Firmware Change</li> <li>Moved NSGSE-212334 from Open - Under Investigation to Closed – Fixed with Hardware Change</li> </ul>	April 2022
006	<ul style="list-style-type: none"> <li>Moved NSGSE-212334 from Closed – fixed with Hardware change to Open – Intend to Fix</li> <li>Moved NSGSE-198265, NSGSE-212334 from Open – Intend to Fix to Closed – Fixed with Firmware Change</li> <li>Added NSGSE-212430 to Closed – No Plan to Fix</li> <li>Moved NSGSE-201063 from Open – Under Investigation to Closed – No Plan to Fix</li> <li>Added NSGSE-214347, NSGSE-217048 to Open – Under Investigation</li> <li>Moved NSGSE-211637 from Open – Under Investigation to Closed – Fixed with Firmware Change</li> <li>Added NSGSE-218406 to Closed – Fixed with Firmware Change</li> </ul>	July 2022
007	<ul style="list-style-type: none"> <li>Added NSGSE-220159 and NSGSE-221734 to Open – Intend to Fix</li> <li>Added NSGSE-223556 to Open – Under Investigation</li> <li>Moved NSGSE-217048 from Open – Under Investigation to Closed - Fixed with Firmware Change</li> </ul>	September 2022
008	<ul style="list-style-type: none"> <li>Added NSGSE-219822, NSGSE-220159 and NSGSE-221734 to Closed – Fixed with Firmware Change</li> </ul>	September 2022

Revision	Description	Revision Date
009	<ul style="list-style-type: none"> <li>MR1 Firmware Release</li> <li>Moved NSGSE-223556 from Open – Under Investigation to Closed - Fixed with Firmware Change</li> <li>Added NSGSE-224231, NSGSE-226270 and NSGSE-226849 to Closed - Fixed with Firmware Change</li> <li>Moved NSGSE-214347 from Open – Under Investigation to Closed – No Plan to Fix</li> <li>Added NSGSE-238367 to Open – Under Investigation</li> </ul>	February 2023
010	<ul style="list-style-type: none"> <li>MR2 Firmware Release</li> <li>Moved NSGSE-236980 and NSGSE-238637 from Open – Under Investigation to Closed – Fixed with Firmware Change</li> <li>Added NSGSE-245820, NSGSE-246182, NSGSE-246354, NSGSE-248503 and NSGSE-254953 to Closed – Fixed with Firmware Change</li> </ul>	October 2023
011	<ul style="list-style-type: none"> <li>Added new sighting category "Closed - Firmware Mitigation Implemented" to section 1.1.1</li> <li>Added NSGSE-267647, NSGSE-270971 and NSGSE-275552 to Closed – Fixed with Firmware Change</li> <li>Added NSGSE-271122 to Closed - Firmware Mitigation Implemented</li> <li>Added NSGSE-275020 to Closed - No Plan to Fix</li> </ul>	July 2024

All information provided here is subject to change without notice.

The products described in this document may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Solidigm technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer.

Solidigm disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase.

Cost reduction scenarios described are intended as examples of how a given Solidigm-based product, in the specified circumstances and configurations, may affect future costs and provide cost -savings. Circumstances will vary. Solidigm does not guarantee any costs or cost reduction.

Solidigm does not control or audit the design or implementation of third-party benchmark data or Web sites referenced in this document. Solidigm encourages all of its customers to visit the referenced Web sites or others where similar performance benchmark data are reported and confirm whether the referenced benchmark data are accurate and reflect performance of systems available for purchase.

© Solidigm. Solidigm and the Solidigm logo are trademarks of Solidigm in the United States and other countries. Other names and brands may be claimed as the property of others.

# Contents

1.0	Overview .....	8
1.1	Nomenclature .....	8
1.1.1	Sighting Status Categories.....	9
2.0	Closed Sightings .....	10
2.1	Closed – Fixed with Hardware Changes.....	10
2.2	Closed – Fixed with Firmware Change.....	10
2.2.1	NSGSE-195807: NAND Byte Written (SMART) Increases after Drive Hot Swap .....	10
2.2.2	NSGSE-198265: Persistent Event Log (PEL) Logs May Be Lost if Accessed Out of Band.....	10
2.2.3	NSGSE-200108: Concurrent Sanitize Operation and Security Commands May Trigger Drive Hang [OPAL Enabled SKU Only].....	11
2.2.4	NSGSE-203817: Mismatch in Actual Measured Power compared to Identify Power State Descriptor Data Structure.....	11
2.2.5	NSGSE-207000: Device Self-Test Command May Fail if Issued too Quickly After a Sanitized Operation	12
2.2.6	NSGSE-208801: SMART Firmware Downgrade Count Still Increments When Firmware Downgrade Operation Fails .....	12
2.2.7	NSGSE-208892: Assert SR002 while Performing Firmware Update without Reset .....	13
2.2.8	NSGSE-211219: VMware Host: Drive not enumerating on PCI after power cycle .....	13
2.2.9	NSGSE-211637: OCP Smart Log Page LID=0xC0 not Working as Expected .....	14
2.2.10	NSGSE-212334: Unsafe Power Cycle Causes Disable Logical Mode BM_0226 On High-Capacity SKUs...	14
2.2.11	NSGSE-217048: Incorrect Value Set for LPA Bit 5 within Identify Controller .....	15
2.2.12	NSGSE-218406: Non-Graceful Shutdown Counter Not Incrementing During Unsafe Shutdowns within DEh Log Page .....	15
2.2.13	NSGSE-219822: UEFI Driver Health Status GetHealthStatus Returning Unsupported Status .....	16
2.2.14	NSGSE-220159: PCIe Unsupported Request Errors Observed by Host during Drive Power Cycle .....	16
2.2.15	NSGSE-221734: Drive De-Enumeration During PCIe Link Up Sequence .....	17
2.2.16	NSGSE-223556: Intermittent Failure During Drive Boot Sequence Causes No Drive Attached .....	17
2.2.17	NSGSE-224231: Use of 8K MPS Could Cause PRP Access Issues (if DAS is not Enabled) .....	18
2.2.18	NSGSE-226270: Incorrect NVMe MI Response Message Status During Drive's Time To Ready (TTR).....	18
2.2.19	NSGSE-226849: First Time LBA Format Change Failure Following a Workload .....	19
2.2.20	NSGSE-236980: NVMe-MI Basic Throttle State Value May Persist After Exiting Throttling State .....	19
2.2.21	NSGSE-238637: Solidigm PCLMT Eye Measurement Unexpected Lane Margin Behavior .....	20
2.2.22	NSGSE-245820: Environment that does not Implement Required Command Ordering Enforcement May Cause Drive to Return Last Data Written to NAND.....	20
2.2.23	NSGSE-246182: Power Management Set Feature Command sent Over NVMe MI Causes Drive to go into Disable Logical Mode (SR003) .....	21
2.2.24	NSGSE-246354: Vendor Unique Feature Identifier C6h Not Reflecting Power Mode Setting Changes Made with Feature Identifier 02h .....	21
2.2.25	NSGSE-248503: Drive May Fail to enumerate During Drive Boot Up Sequence .....	22
2.2.26	NSGSE-254953: PCIe link up sequence may fail during AC/DC cycle .....	22
2.2.27	NSGSE-267647: Admin Command Timeout During Read Intensive Workloads .....	23
2.2.28	NSGSE-270971: Drive May Fail to Enumerate Following Unplanned Power Cycle .....	23
2.2.29	NSGSE-275552: Previous Firmware Fix May Lead to Drive in Disable Logical Mode .....	24
2.3	Closed - Firmware Mitigation Implemented.....	24
2.3.1	NSGSE-273122: Failure of MOS-FET Component Leads to Higher Drive Failure Rate .....	24
2.4	Closed – Fixed with Driver Change .....	25
2.5	Closed – Documentation / Specification Change .....	25
2.6	Closed – No Plan to Fix.....	25

2.6.1	NSGSE-192480: MBR Set / Get Methods Cause Concurrent IO Failures .....	25
2.6.2	NSGSE-195579: Error during simultaneous Extended Device Self-Test and Firmware Upgrade Without Reset.....	25
2.6.3	NSGSE-201063: NVMe-MI Command Following a Firmware Update and NVMe Reset exceeds MCTP command timeout.....	26
2.6.4	NSGSE-202068: Concurrent IO interrupted by Set Method [OPAL Enabled SKU Only].....	26
2.6.5	NSGSE-205704: MNS Trim with VSS – NS MGMT Commands Running Long / Never Complete .....	27
2.6.6	NSGSE-212430: Level 0 Discovery Over NVMe-MI May Fail During Execution of Concurrent Commands [OPAL Enabled SKU Only].....	27
2.6.7	NSGSE-214347: Namespace Attach Operation May Fail When Issued Through NVMe-MI .....	28
2.6.8	NSGSE-275020: Drive May Fail to Be Detected Following a Power Loss During Firmware Upgrade.....	28
2.7	Closed – Non-Solidigm SSD Sighting .....	28
3.0	Open Sightings .....	29
3.1	Open – Intend to Fix.....	29
3.2	Open – Under Investigation .....	29
3.3	3rd Party Sightings .....	29
3.3.1	NSGSE-205946: FW Upgrade/Downgrade with IMAS or NVMe-cli will cause error message in dmesg...	29

## Firmware Release Matrix

Product Family	Capacity	Firmware Version	Milestone	Release Date
D7-P5520 – U.2 15mm	1.92, 3.84, 7.68 & 15.36 TB	9CV10015	PRQ1 ES2	September 2021
D7-P5520 – U.2 15mm	1.92 TB	9CV10100	PRQ1	January 2022
D7-P5520 – U.2 15mm D7-P5520 – E1.S 9.5mm D7-P5520 – E1.S 15mm D7-P5620 – U.2 15mm	1.92, 3.84, 7.68 & 15.36 TB 1.92, 3.84 & 7.68 TB 1.92, 3.84 & 7.68 TB 1.6, 3.2, 6.4 & 12.8 TB	9CV10200	PRQ2	April 2022
D7-P5520 – E1.L 9.5mm	15.36 TB	9CV10310	PRQ3	July 2022
D7-P5520 – U.2 15mm D7-P5520 – E1.S 9.5mm D7-P5520 – E1.S 15mm D7-P5620 – U.2 15mm D7-P5520 – E1.L 9.5mm	1.92, 3.84, 7.68 & 15.36 TB 1.92, 3.84 & 7.68 TB 1.92, 3.84 & 7.68 TB 1.6, 3.2, 6.4 & 12.8 TB 15.36TB	9CV10410	MR1	February 2023
D7-P5520 – U.2 15mm D7-P5520 – E1.S 9.5mm D7-P5520 – E1.S 15mm D7-P5620 – U.2 15mm D7-P5520 – E1.L 9.5mm	1.92, 3.84, 7.68 & 15.36 TB 1.92, 3.84 & 7.68 TB 1.92, 3.84 & 7.68 TB 1.6, 3.2, 6.4 & 12.8 TB 15.36TB	9CV10450	MR2	October 2023
D7-P5520 – U.2 15mm D7-P5520 – E1.S 9.5mm D7-P5520 – E1.S 15mm D7-P5620 – U.2 15mm D7-P5520 – E1.L 9.5mm	1.92, 3.84, 7.68 & 15.36 TB 1.92, 3.84 & 7.68 TB 1.92, 3.84 & 7.68 TB 1.6, 3.2, 6.4 & 12.8 TB 15.36TB	9CV10490	MR3	July 2024

## Glossary

Term	Definition
Disable Logical mode	If the drive senses an internal failure condition (e.g., unrecoverable NAND errors, missing meta-data), it saves log data and resets. After reset, the drive will enter Disable Logical mode. A Disable Logical drive will power-up and enumerate in the host system, but all read and write operations to the NAND Flash memory will be aborted. A drive in Disable Logical mode is faulted and may not be user recoverable.
DIX	Data Integrity Extensions
Drive	Unless otherwise specified, “drive” refers generically to all product SKUs, including 2.5” and M.2
FW	Firmware
JEDEC	Joint Electron Device Engineering Council: <a href="http://www.jedec.org">www.jedec.org</a>
LBA	Logical Block Address
Lightswitch	Factory Configuration. Firmware drive configuration parameters set by the factory.
NDA	No Drive Attached. Drive is not enumerating or detected by the host.
NVMe	Non-Volatile Memory Express (NVM Express)
OS	Operating System
PBA	Pre-boot Authentication
PCIe	Peripheral Component Interconnect Express (PCI Express)
RHEL	Red Hat Enterprise Linux: <a href="http://www.redhat.com">www.redhat.com</a>
SSD	Solid State Drive
VPD	Vital Product Data

## 1.0 Overview

This document is a compilation of device and document sightings as well as specification clarifications and changes for the Solidigm™ D7-P5520/D7-P5620 (formerly Intel®).

This document is intended for hardware system manufacturers and software developers of applications, operating systems and tools. Throughout the development process, Solidigm periodically updates the sightings document as new information becomes available. These updates are also available from your local Solidigm representative.

## 1.1 Nomenclature

The sightings in this document are separated into sub-groups to identify the status of the sightings and what action, if any, needs to be taken to address the sightings. The names and definitions are detailed in the following table.

Sightings will not be classified as Closed until the fix is verified with the appropriate hardware revision, firmware/driver revision, or disposition stated the following table. Sightings will be classified as Open until the fix is verified with the appropriate silicon stepping or firmware release.



### 1.1.1 Sighting Status Categories

Category/ Sub classification	Description
Closed - Fixed with Hardware Change	All sightings detailed in this section have been resolved in the hardware revision identified in the details for each sighting.
Closed - Fixed with Firmware/Bootloader Change	All sightings detailed in this section have been resolved in the firmware revision identified in the details for each sighting.
Closed - Firmware Mitigation Implemented	All sightings detailed in this section had their severity reduced in the firmware revision identified in the details for each sighting.
Closed - Fixed with Driver Change	All sightings detailed in this section have been resolved in the driver revision identified in the details for each sighting.
Closed - Documentation / Specification Change	All sightings detailed in this section require a change to either the product manual or specification. The specific revision to the appropriate documentation will be identified for each sighting.
Closed - No Plan to Fix	All sightings detailed in this section are not planned to be fixed in any version of this drive.
Closed - Non-Solidigm SSD Sighting	All sightings detailed in this section are items that are not related to these drives. These sightings may be due to BIOS, driver, OS, system software, or other considerations.
Open - Intend to Fix	All sightings detailed in this section have been correlated and a potential fix has been identified. Once the potential fix has been validated on the appropriate firmware version, the fix will be implemented, and the sighting will be classified as "Closed."
Open - Under Investigation	All sightings detailed in this section are still under investigation. The root cause for each sighting may or may not have been identified.
3 <sup>rd</sup> Party Sightings	All sightings detailed in this section are related to 3 <sup>rd</sup> party tools that affect this drive.

## 2.0 Closed Sightings

### 2.1 Closed – Fixed with Hardware Changes

None

### 2.2 Closed – Fixed with Firmware Change

#### 2.2.1 NSGSE-195807: NAND Byte Written (SMART) Increases after Drive Hot Swap

	Description
Problem Statement	NAND Byte Written count increases after drive hot swap.
Details	Performing a drive hot swap will result in the SMART NAND Byte Written count of that drive to increase. NAND Byte Written count should be the same before and after hot swap.
Impact	No functional impact – will observe SMART count increase.
Workaround	Avoid drive hot swap operation.
Root Cause	Byte written count was not being reset correctly following a drive hot swap.
Firmware Version Found / Affected	9CV10015
Status	Closed – Fixed with Firmware 9CV10100

#### 2.2.2 NSGSE-198265: Persistent Event Log (PEL) Logs May Be Lost if Accessed Out of Band

	Description
Problem Statement	Persistent Event Logs (PEL) logs may be lost if accessed out-of-band
Details	In rare occurrences, out-of-band PEL log request following out-of-band issue of admin commands and an in-band NVMe Reset may cause PEL log corruption leading to loss of PEL logs.
Impact	In rare occurrence, PEL logs may be lost if accessed out of band. Once lost, those logs cannot be retrieved in band or out of band.
Workaround	Recommending to access PEL logs in band only.
Root Cause	Incorrect implementation of PEL blob save time.
Firmware Version Found / Affected	All firmware versions prior to 9CV10310
Status	Closed – Fixed with Firmware 9CV10310

### 2.2.3 NSGSE-200108: Concurrent Sanitize Operation and Security Commands May Trigger Drive Hang [OPAL Enabled SKU Only]

	Description
Problem Statement	Concurrent Sanitize operation and security commands may trigger a drive hang.
Details	When a Sanitize operation targets a namespace that is already locked due to a security command in progress, it can cause the drive to hang.
Impact	Drive hang, power cycling needed to recover. Drive recovers healthy.
Workaround	Wait for security command to complete before issuing Sanitize operation (Refer to product spec for sanitize completion times).
Root Cause	Sanitize command will be blocked if there is an open session detected by firmware.
Firmware Version Found / Affected	9CV10100
Status	Closed – Fixed with Firmware 9CV10200

### 2.2.4 NSGSE-203817: Mismatch in Actual Measured Power compared to Identify Power State Descriptor Data Structure

	Description
Problem Statement	Mismatch in actual measured power compared to Identify Power State descriptor data structure.
Details	Values report in in Power State Descriptor 1 and 2 do not match the actual measured values. The typical difference is 1-2W higher than actual measured values.
Impact	Reported Power will be higher than measured power. No functional impact. Drive is healthy.
Workaround	None.
Root Cause	Drive Power tokens need additional tuning.
Firmware Version Found / Affected	9CV10100
Status	Closed – Fixed with Firmware 9CV10200

### 2.2.5 NSGSE-207000: Device Self-Test Command May Fail if Issued too Quickly After a Sanitized Operation

	Description
Problem Statement	Device Self-Test (DST) command may fail with “Sanitize in Progress” after a Sanitize operation completed if issued too quickly.
Details	If a DST command is issued too quickly after a Sanitize operation completed the command may fail. Reissuing the DST command will solve the issue.
Impact	A DST command after a Sanitize operation completion may fail, drive is healthy.
Workaround	Add a 1 second delay between Sanitize operation completion and DST command.
Root Cause	There is a race condition between reporting a Sanitize operation completed and disabling the filter preventing commands to be executed concurrently.
Firmware Version Found / Affected	9CV10100
Status	Closed – Fixed with Firmware 9CV10200

### 2.2.6 NSGSE-208801: SMART Firmware Downgrade Count Still Increments When Firmware Downgrade Operation Fails

	Description
Problem Statement	SMART Firmware Downgrade Count still increments when a firmware downgrade operation fails.
Details	If a firmware downgrade operation fails due to security revision that would block this operation, the SMART Firmware Downgrade Counter will still increment in the CAh log page (attribute F9h).
Impact	Excessive failed firmware downgrade operations (usually performed during qualification) would cause the Firmware Downgrade Count to hit its maximum limit of 166. This would prevent any additional firmware downgrade operations.
Workaround	Do not perform excessive firmware downgrade operations from firmware 9CV10100 to 9CV10015. To test firmware upgrade / downgrade operations use firmware 9CV10100 and 9CV10101 (buddy build).
Root Cause	Firmware downgrade count is incremented before firmware validation check is completed.
Firmware Version Found / Affected	9CV10100
Status	Closed – Fixed with Firmware 9CV10200

### 2.2.7 NSGSE-208892: Assert SR002 while Performing Firmware Update without Reset

	Description
Problem Statement	Drive may hit assert SR002 while performing a firmware update without reset.
Details	If there are any outstanding Asynchronous Event Requests (AERs) while performing a firmware without reset with Commit Action = 3, the drive may hit assert SR002.
Impact	Drive will be in disable logical state. hit an assert, need to use Host Data Recovery (VU Opcode xC8) to recover the drive.
Workaround	Perform firmware update with reset. Or recover the drive from disable logical state by using Host Data Recovery (VU Opcode xC8) to recover the drive.
Root Cause	Due to writing the AER data structure to an invalid command buffer space.
Firmware Version Found / Affected	9CV10100
Status	Closed – Fixed with Firmware 9CV10200

### 2.2.8 NSGSE-211219: VMware Host: Drive not enumerating on PCI after power cycle

	Description
Problem Statement	VMware Host: Drive not enumerating on PCI after power cycle
Details	Under certain conditions, the drive may transmit MCTP over PCIe VDM packets prior to the host assigning a PCIe BDF (BUS : DEVICE : FUNCTION) to the drive during hot add of the SSD
Impact	The drive may get a broadcast BDF which 0:0:0 and will lead the PSOD of the host
Workaround	None
Root Cause	VDM was enabled before PCIe BDF assignment from the host.
Firmware Version Found / Affected	9CV10100
Status	Closed – Fixed with Firmware 9CV10200

### 2.2.9 NSGSE-211637: OCP Smart Log Page LID=0xC0 not Working as Expected

	Description
Problem Statement	NUSE field within OCP Smart Log Page LID=0xC0 does not return namespace size.
Details	OCP spec for C0 log page states that “NUSE” field for a multiple namespace device shall reflect the total utilization based on all the namespaces. Currently NUSE reports “0”.
Impact	Wrong value reported. Drive is healthy.
Workaround	No workaround.
Root Cause	Wrong implementation.
Firmware Version Found / Affected	All firmware versions prior to 9CV10310
Status	Closed- Fixed with Firmware 9CV10310

### 2.2.10 NSGSE-212334: Unsafe Power Cycle Causes Disable Logical Mode BM\_0226 On High-Capacity SKUs

	Description
Problem Statement	Unsafe Power Cycle causes disable logical state (BM0226) on 12.8TB & 15.36TB SKUs.
Details	In rare occurrences unsafe power cycles may cause drive to enter disable logical state BM0226 (Page Out of Range). SKUs affected are 12.8TB & 15.36TB.
Impact	Drive in disable logical state.
Workaround	None.
Root Cause	Context save requires higher budget on 12.8TB and 15.36TB.
Firmware Version Found / Affected	All firmware versions prior to 9CV10310
Status	Closed- Fixed with Firmware 9CV10310

### 2.2.11 NSGSE-217048: Incorrect Value Set for LPA Bit 5 within Identify Controller

	Description
Problem Statement	Incorrect Value Set for Log Page Attribute (LPA) Bit 5 within Identify Controller.
Details	LPA field within Identify Controller indicates optional attributes for log pages that are accessed via the Get Log Page Command. Log Page Identifier Bit 5 should be set to '0' instead of '1' since drives do not support returning the scope of each command in the Commands Supported and Effects log page.
Impact	Fails to meet NVMe Spec 1.4.
Workaround	None. Log pages log identifier 0h, 5h, 12h and 13h remain accessible
Root Cause	Incorrect value set for Bit 5 in firmware.
Firmware Version Found / Affected	All firmware versions prior to 9CV10310
Status	Closed – Fixed with Firmware 9CV10310

### 2.2.12 NSGSE-218406: Non-Graceful Shutdown Counter Not Incrementing During Unsafe Shutdowns within DEh Log Page

	Description
Problem Statement	Non-graceful shutdown counter not incrementing during unsafe shutdowns within DEh log page.
Details	Unsafe shutdowns such as hot-plug or removal of AC power do not contribute to non-graceful shutdown counter to be incremented within DEh log page.
Impact	Drive is healthy. Non-graceful shutdown counter not incrementing.
Workaround	Log page 02h can be used instead to track non-graceful shutdown count.
Root Cause	Incorrect firmware implementation of the SMART Attribute.
Firmware Version Found / Affected	All firmware versions prior to 9CV10310
Status	Closed – Fixed with Firmware 9CV10310

### 2.2.13 NSGSE-219822: UEFI Driver Health Status GetHealthStatus Returning Unsupported Status

	Description
Problem Statement	UEFI driver health status GetHealthStatus is returning the unsupported status.
Details	With more than one drive installed, the GetHealthStatus may incorrectly return the unsupported status.
Impact	Incorrect GetHealthStatus. Drive is healthy.
Workaround	None
Root Cause	UEFI driver bug, updating the UEFI driver to a newer version fixed the issue.
Firmware Version Found / Affected	All firmware versions prior to 9CV10410
Status	Closed- Fixed with Firmware 9CV10410

### 2.2.14 NSGSE-220159: PCIe Unsupported Request Errors Observed by Host during Drive Power Cycle

	Description
Problem Statement	PCIe Unsupported request errors observed by host during drive power cycle.
Details	PCIe Unsupported Request (UR) errors will be observed by the host through AER during power cycles as the PCIe link is under negotiation because Unsupported Request Reporting Enable (URRE) bit is set to 1.
Impact	No functional impact. Drive is healthy.
Workaround	None
Root Cause	URRE bit in the PCIe register incorrectly set to 1
Firmware Version Found / Affected	All firmware versions prior to 9CV10410
Status	Closed- Fixed with Firmware 9CV10410



### 2.2.15 NSGSE-221734: Drive De-Enumeration During PCIe Link Up Sequence

	Description
Problem Statement	Drive de-enumeration during PCIe link up sequence
Details	In rare occurrences, drive PCIe link up sequence can fail resulting in a no drive attached failure.
Impact	No drive attached. Drive comes back healthy after power cycle.
Workaround	None
Root Cause	During PCIe link up, unexpected timing correlation occurred where background smart logging feature and link up interrupt handling were accessing the same register. This simultaneous register access overwrote and cleared a valid link up enable bit operation resulting in no drive attached during link up sequence.
Firmware Version Found / Affected	All firmware versions prior to 9CV10410
Status	Closed- Fixed with Firmware 9CV10410

### 2.2.16 NSGSE-223556: Intermittent Failure During Drive Boot Sequence Causes No Drive Attached

	Description
Problem Statement	Intermittent failure of the drive boot-up sequence causes the drive to not enumerate.
Details	Failure during the drive boot-up sequence causes the drive to no enumerate. The failure happens during power up / power cycle only, no runtime failure or runtime impact.
Impact	Drive does not enumerate at power up or during a power cycle by the system.
Workaround	Firmware mitigation identified. Optimized ASIC ROM loading and access process to allow DRAM enumeration sufficient time to complete initialization. Power cycling in case of NDA for drives running FW prior to 9CV10410.
Root Cause	A call to DRAM during the drive initial boot up sequence prior to DRAM enumeration completion causes drive hang.
Firmware Version Found / Affected	All firmware versions prior to 9CV10410
Status	Closed- Fixed with Firmware 9CV10410

### 2.2.17 NSGSE-224231: Use of 8K MPS Could Cause PRP Access Issues (if DAS is not Enabled)

	Description
Problem Statement	Use of 8K Memory Page Size (MPS) could cause Physical Region Page (PRP) access issues (if DAS is not enabled)
Details	I/O tools could report DMA errors when Driver Assisted Striping (DAS) is not enabled due to incorrect implementation of 8K split.
Impact	I/O tools report DMA errors. Drive is healthy.
Workaround	Enabling Driver Assisted Striping (DAS) resolves the issue.
Root Cause	Internal split of 8K PRP was incorrectly programmed within firmware.
Firmware Version Found / Affected	All firmware versions prior to 9CV10410
Status	Closed- Fixed with Firmware 9CV10410

### 2.2.18 NSGSE-226270: Incorrect NVMe MI Response Message Status During Drive's Time To Ready (TTR)

	Description
Problem Statement	An incorrect NVMe MI response message status is returned during the drive's TTR.
Details	While the drive is coming out of reset, receiving a NVMe MI admin command from the host causes to return a success response message status (00h) with empty or erroneous value instead of an access denied status (07h).
Impact	NVMe MI commands issues during drive TTR may return empty or erroneous value with a success response message status value.
Workaround	Issue NVMe MI commands once the TTR has passed.
Root Cause	NVMe-MI Specification interpretation. Fixed by returning Access Denied Status (07h).
Firmware Version Found / Affected	All firmware versions prior to 9CV10410
Status	Closed- Fixed with Firmware 9CV10410

### 2.2.19 NSGSE-226849: First Time LBA Format Change Failure Following a Workload

	Description
Problem Statement	Issuing a LBA format change after running a workload on the drive may return an error.
Details	Issuing a LBA format change (nvme format -l xx) after running a workload on the drive will return an interrupt system call error and the LBA format change will not complete. Found by customer in Sep'22, failure rate is about 50%.
Impact	LBA format not changed after issuing the command.
Workaround	Issuing the LBA format change command again.
Root Cause	Conflict with background defrag operation.
Firmware Version Found / Affected	All firmware versions prior to 9CV10410
Status	Closed –Fixed with Firmware 9CV10410

### 2.2.20 NSGSE-236980: NVMe-MI Basic Throttle State Value May Persist After Exiting Throttling State

	Description
Problem Statement	In the event of a power cycle or NVMe reset while the drive is in throttling state, the NVMe-MI Basic Throttle State value will not refresh when the drive re-enumerates.
Details	When throttling the NVMe-MI Basic Throttle State value will be <100%. If a power cycle or NVMe reset is issued while the drive is throttling, the Throttle State value will still be <100% when the drive re-enumerates. This remains true even if the drive has exited throttling state.
Impact	The NVMe-MI Basic Throttle State will have a value of <100% even if the drive is not throttling anymore.
Workaround	Bring the drive to throttle state again and make it exit throttle state gracefully with no power cycle or NVMe reset.
Root Cause	The Power Loss Imminent (PLI) following a power cycle or NVMe reset restores the last known NVMe-MI Basic Throttle State value.
Firmware Version Found / Affected	9CV10410
Status	Closed –Fixed with Firmware 9CV10450

### 2.2.21 NSGSE-238637: Solidigm PCLMT Eye Measurement Unexpected Lane Margin Behavior

	Description
Problem Statement	Margins reported by the Solidigm Platform Connected Lane Margin Tool (PCLMT) tool are not accurate.
Details	During eye measurement using PCLMT, the drive exhibits perfect eye margins which is unexpected behavior.
Impact	PCLMT will report perfect eye margins, this result is not accurate.
Workaround	Use previously released firmware 9CV10200 or 9CV10320 to get accurate eye margins data.
Root Cause	Under investigation
Firmware Version Found / Affected	9CV10410
Status	Closed –Fixed with Firmware 9CV10450

### 2.2.22 NSGSE-245820: Environment that does not Implement Required Command Ordering Enforcement May Cause Drive to Return Last Data Written to NAND

	Description
Problem Statement	If the host software or associated application does not enforce any required command ordering and is running a workload that targets a single location with multiple Write commands followed immediately by a Read command to the same location, the drive may return the last data written to NAND.
Details	<p>Per NVMe Specification 1.3/1.4 Section 6.3: "...the controller is not responsible for checking the LBA of a Read or Write command to ensure any type of ordering between commands. If there are any ordering requirements between ... commands, host software or associated application is required to enforce that ordering above the level of the controller." (<a href="https://nvmexpress.org/wp-content/uploads/NVM-Express-1.4c-2021.06.28-Ratified.pdf">https://nvmexpress.org/wp-content/uploads/NVM-Express-1.4c-2021.06.28-Ratified.pdf</a>)</p> <p>If the host software or associated application does not enforce any required command ordering and is running a workload that targets a single location with multiple Write commands followed immediately by a Read command to the same location, the drive may return the last data written to NAND.</p>
Impact	When the host software or associated application does not enforce any required command ordering and is running a workload as described above, the drive may return the last data written to NAND.
Workaround	Following NVMe Specification 1.3/1.4 Section 6.3 recommendation to use a host software or associated application that enforces required command ordering will mitigate the issue. If it is not possible, the issue can be mitigated by ensuring Completion Queue entry is received for in-flight Write command before executing a Read command to the same location.
Root Cause	Due to a firmware logic error, when servicing a Read command to a target location with Write commands in-flight to the same location, the drive may return the last data written to NAND for the target location if the host software or associated application does not enforce required command ordering.
Firmware Version Found / Affected	All firmware versions prior to 9CV10450
Status	Closed –Fixed with Firmware 9CV10450

### 2.2.23 NSGSE-246182: Power Management Set Feature Command sent Over NVMe MI Causes Drive to go into Disable Logical Mode (SR003)

	Description
Problem Statement	Power management set feature command sent over NVMe MI causes drive to go into Disable Logical mode (SR003)
Details	Power management set feature command sent over NVMe MI before the config ready bit is set causes drive to go into Disable Logical mode (SR003)
Impact	Drive enters Disable Logical mode
Workaround	Allowing BMC to wait for config ready bit to be set before sending NVMe MI commands to the drive.
Root Cause	In certain cases, the BMC may issue a power management set feature over NVMe MI before the config ready bit is set during the transition from boot firmware to drive firmware. In that case the drive is not ready to receive this command which is resulting in disable logical mode.
Firmware Version Found / Affected	9CV10410
Status	Closed –Fixed with Firmware 9CV10450

### 2.2.24 NSGSE-246354: Vendor Unique Feature Identifier C6h Not Reflecting Power Mode Setting Changes Made with Feature Identifier 02h

	Description
Problem Statement	Vendor Unique Power Governor Setting (Feature Identifier C6h) not reflecting power mode setting changes made with Power Management (Feature Identifier 02h).
Details	When changing the power mode setting using Power Management (Feature Identifier 02h), the new power mode setting will not be reflected in the Vendor Unique Power Governor Setting (Feature Identifier C6h).
Impact	Vendor Unique Power Governor Setting (Feature Identifier C6h) may not return the current power mode setting.
Workaround	Check power mode setting using Power Management (Feature Identifier 02h).
Root Cause	Vendor Unique Power Governor Setting (Feature Identifier C6h) behavior was not in line with Power Management (Feature Identifier 02h).
Firmware Version Found / Affected	All firmware versions prior to 9CV10450
Status	Closed- Fixed with Firmware 9CV10450

### 2.2.25 NSGSE-248503: Drive May Fail to enumerate During Drive Boot Up Sequence

	Description
Problem Statement	Drive may fail to enumerate after a power cycle, during drive boot up sequence.
Details	In rare occurrences, drive boot up sequence may fail resulting in no NVMe enumeration or drive showing as "functional disabled".
Impact	No NVMe enumeration or drive shows as "functional disabled". Drive comes back healthy after power cycle.
Workaround	Power cycle the drive once.
Root Cause	In occurrences where wear leveling operation has been initiated prior to or during a power cycle, it may interfere with saving internal system files to NAND. This contention prevents the drive to move to a full ready state.
Firmware Version Found / Affected	All firmware versions prior to 9CV10450
Status	Closed- Fixed with Firmware 9CV10450

### 2.2.26 NSGSE-254953: PCIe link up sequence may fail during AC/DC cycle

	Description
Problem Statement	PCIe link up sequence may fail during AC/DC cycle.
Details	In rare occurrences, drive PCIe link up sequence can fail resulting in a no drive attached failure.
Impact	No drive attached. Drive comes back healthy after a power cycle.
Workaround	None
Root Cause	Improper firmware operation to read DVSEC (Designed Vendor-Specific Capability) register during PCIe subsystem reset. DVSEC register will be masked during PCIe subsystem reset firmware flow. DVSEC hash check will find an unexpected mismatch, which may result in a drive de-enumeration failure.
Firmware Version Found / Affected	All firmware versions prior to 9CV10450
Status	Closed- Fixed with Firmware 9CV10450

### 2.2.27 NSGSE-267647: Admin Command Timeout During Read Intensive Workloads

	Description
Problem Statement	Drive may encounter admin command timeout during read intensive workloads.
Details	Under pure read workloads, where defrag read disturb relocation operations can trigger more frequently, admin commands to read error log pages may be blocked resulting in error log page timeout.
Impact	Admin command timeout reported. Drive is healthy.
Workaround	Avoid issuing admin commands to read error log pages under pure read workloads.
Root Cause	Priority conflict between defrag operation and admin commands.
Firmware Version Found / Affected	All firmware versions prior to 9CV10490
Status	Closed- Fixed with Firmware 9CV10490

### 2.2.28 NSGSE-270971: Drive May Fail to Enumerate Following Unplanned Power Cycle

	Description
Problem Statement	Drive may fail to enumerate following an unplanned power cycle.
Details	During unplanned power cycle events, drive initialization may fail because firmware is unable to complete erase events.
Impact	No drive attached.
Workaround	Avoid unplanned power cycle.
Root Cause	Incorrect firmware optimization prevents erase events from completing during initialization after an unplanned power cycle.
Firmware Version Found / Affected	9CV10450
Status	Closed- Fixed with Firmware 9CV10490

### 2.2.29 NSGSE-275552: Previous Firmware Fix May Lead to Drive in Disable Logical Mode

	Description
Problem Statement	A previous firmware fix (NSGSE-248503) may lead drives to enter disable logical mode.
Details	Due to the previous fix a band may be left in an improper state following a power cycle. This improper band state may create a contention between two bands leading to the drive entering disable logical mode.
Impact	Drive enters disable logical mode.
Workaround	Avoid power cycling the drive.
Root Cause	A previous fix introduced a potential contention between two bands.
Firmware Version Found / Affected	9CV10450
Status	Closed- Fixed with Firmware 9CV10490

## 2.3 Closed - Firmware Mitigation Implemented

### 2.3.1 NSGSE-273122: Failure of MOS-FET Component Leads to Higher Drive Failure Rate

	Description
Problem Statement	Failure of a MOS-FET component leads to failure in the power delivery circuitry of the drive.
Details	The reliability margin of the Diodes Inc. 5015 MOS-FET was insufficient to account for the oxide degradation due to sustained voltage in the Power Loss Protection circuit. This leads to a higher drive failure rate.
Impact	Drive is NDA and unrecoverable.
Workaround	None
Root Cause	Insufficient voltage margin leads to decreased reliability under sustained voltage.
Firmware Version Found / Affected	All firmware versions prior to 9CV10490
Status	Closed- Mitigated with Firmware 9CV10490



## 2.4 Closed – Fixed with Driver Change

None

## 2.5 Closed – Documentation / Specification Change

None

## 2.6 Closed – No Plan to Fix

### 2.6.1 NSGSE-192480: MBR Set / Get Methods Cause Concurrent IO Failures

	Description
Problem Statement	MBR Set / Get methods cause concurrent IO failures.
Details	Issuing MBR Set / Get commands on Opal drives at the same time as IO commands cause IO drop or not be serviced.
Impact	Drive is healthy. IO commands resume operation once MBR configuration is complete.
Workaround	Avoid running IO concurrent with OPAL MBR Set / Get commands.
Firmware Version Found / Affected	All firmware versions
Status	Closed – No Plan to Fix

### 2.6.2 NSGSE-195579: Error during simultaneous Extended Device Self-Test and Firmware Upgrade Without Reset

	Description
Problem Statement	Device self-test failure after firmware upgrade without reset.
Details	After performing a firmware upgrade without reset the device self-test will report that the operation completed with one or more failed segments (status code 7h - as per NVMe Specification 1.4, Section 5.14.1.6).
Impact	Error during firmware upgrade. Drive is healthy.
Workaround	None
Firmware Version Found / Affected	All firmware versions
Status	Closed – No Plan to Fix

### 2.6.3 NSGSE-201063: NVMe-MI Command Following a Firmware Update and NVMe Reset exceeds MCTP command timeout

	Description
Problem Statement	NVMe-MI Command Following a Firmware Update and NVMe Reset exceeds MCTP command timeout
Details	NVMe-MI Command Following a Firmware Update and NVMe Reset exceeds MCTP command timeout of 100ms. Drive does not respond back with More Processing Required.
Impact	NVMe MI commands not completing within 100ms.
Workaround	None
Root Cause	Hardware limitation due to SPI erase and program commands taking too long.
Firmware Version Found / Affected	All firmware versions
Status	Closed – No Plan to Fix

### 2.6.4 NSGSE-202068: Concurrent IO interrupted by Set Method [OPAL Enabled SKU Only]

	Description
Problem Statement	Opal Set Method execution fails concurrent IO traffic.
Details	In band or out of band Opal Set Method commands will fail with concurrent IO Traffic.
Impact	Drive is healthy, but IO interruption during security operation.
Workaround	Interrupt IO traffic before executing Set Method Command.
Firmware Version Found / Affected	All firmware versions
Status	Closed – No Plan to Fix

### 2.6.5 NSGSE-205704: MNS Trim with VSS – NS MGMT Commands Running Long / Never Complete

	Description
Problem Statement	Creating and deleting Namespace with VSS takes longer than default 60s timeout.
Details	Creating and deleting NS with VSS takes longer than 60s and in some cases can take up to 10 minutes.
Impact	User unable to create or delete NS with VSS enabled. Drive is healthy.
Workaround	Increase OS timeout to 10 minutes while creating or deleting NS.
Firmware Version Found / Affected	All firmware versions
Status	Closed – No Plan to Fix

### 2.6.6 NSGSE-212430: Level 0 Discovery Over NVMe-MI May Fail During Execution of Concurrent Commands [OPAL Enabled SKU Only]

	Description
Problem Statement	Level 0 Discovery Over NVMe-MI May Fail during Execution of Concurrent Commands
Details	<p>Level 0 Discovery over NVMe-MI may fail when the following commands are executed concurrently:</p> <ul style="list-style-type: none"> <li>• Sanitize over NVMe</li> <li>• OPAL Level 0 discovery over NVMe MI</li> </ul> <p>TCG OPAL commands over VDM</p>
Impact	Drive is healthy.
Workaround	Retry Opal Level 0 discovery command.
Root Cause	Sequence of concurrent commands listed is not allowed per NVMe spec. In particular, Security Send/Receive is not allowed when Sanitize is in progress.
Firmware Version Found / Affected	All firmware versions
Status	Closed – No Plan to Fix

### 2.6.7 NSGSE-214347: Namespace Attach Operation May Fail When Issued Through NVMe-MI

	Description
Problem Statement	Namespace attach operation may fail when issued through NVMe-MI
Details	<p>Sequence of commands causing Namespace attach command to not complete when issued over NVMe-MI</p> <ol style="list-style-type: none"> <li>1. Firmware update (either over NVMe or NVMe-MI)</li> <li>2. Namespace create (over NVMe-MI)</li> </ol> <p>Namespace attach (over NVMe-MI) --- FAIL</p>
Impact	Failure to successfully complete namespace attach operations over NVMe-MI.
Workaround	Re-try Namespace attach command after a Reset.
Root Cause	Architectural limitation.
Firmware Version Found / Affected	All firmware versions
Status	Closed – No Plan to Fix

### 2.6.8 NSGSE-275020: Drive May Fail to Be Detected Following a Power Loss During Firmware Upgrade

	Description
Problem Statement	Drive may fail to be detected following a power loss during firmware upgrade.
Details	When a drive loses power during firmware upgrade, it may become undetected by the operating system once it regains power.
Impact	Drive not detected by operating system, drive is healthy.
Workaround	Issue a PCIe rescan to bring the drive back up.
Root Cause	Architectural limitation.
Firmware Version Found / Affected	All firmware versions
Status	Closed – No Plan to Fix

## 2.7 Closed – Non-Solidigm SSD Sighting

None

## 3.0 Open Sightings

### 3.1 Open – Intend to Fix

### 3.2 Open – Under Investigation

### 3.3 3rd Party Sightings

#### 3.3.1 NSGSE-205946: FW Upgrade/Downgrade with IMAS or NVMe-cli will cause error message in dmesg

	Description
Problem Statement	An error message is entered in dmesg log when an admin command is issued during reset process.
Details	“Identify namespace failed” or “Identify descriptors failed” error may occur due to timing of the first admin command sent from the tool during reset.
Impact	Error message logged in dmesg log. Drive is healthy.
Workaround	None
Root Cause	Error will be seen if admin command gets into submission queue after its already being deleted as part of reset.
Firmware Version Found / Affected	N/A
Status	Closed – Non-Solidigm SSD Sighting