



NVDIMM DSM Interface

Revision V1.7

April, 2018

This version of the specification includes the following modifications and additions:

- General:
 - Updated Table 3-C Supported _DSM Return Status Values
 - Updated description for error code 5
- Added new security specific return status values 10, 11
- Get SMART Health Info:
 - Removed PMIC Temperature
 - Renamed Unsafe Shutdown Count (USC) to Dirty Shutdown Count (DSC) to match FIS
- Get FW Info
 - Corrected FW Update Capabilities to be a Bit Field.
- Start FW Update
- Send FW Update Data
 - Removed return status 05 (retry due to operation in progress) – Since these two DSMs start a FW update sequence and move data from host buffer into BIOS mailbox buffer and do not issue mailbox commands to the NVDIMM, they will not return this status.
- New Security DSMs
 - Theory Of Operation
 - Added "Basic Application In-band Get Security State Sequence"
 - Updated "Application In-band Secure Erase & Overwrite NVDIMM Sequence"
 - Get Security State
 - Cleaned up definitions of each security state bit to match FISRemoved Bit[6] - Nonce Set - BIOS Security Nonce has been set on the NVDIMM No host meaning or implication and difficult to test without new injection
 - Set Passphrase
 - Disable Passphrase
 - Unlock Unit
 - Freeze Lock
 - Removed Current Passphrase argument
 - Secure Erase NVDIMM
 - Added details and SW implications when the on-NVDIMM metadata is destroyed



This version of the specification includes the following modifications and additions (Continued):

- *New Security DSMs (Continued)*
 - o *Overwrite NVDIMM*
 - *Added details and SW implications when the on-NVDIMM metadata is destroyed*
 - *Added details and SW implications when the PCD and LSA metadata is destroyed*
 - *Added details and SW implications for changes in the NFIT due to the overwrite*
 - o *Query Overwrite NVDIMM Status*
 - *Added missing Extended Status for sequencing error, failing to call Overwrite NVDIMM before Query Overwrite NVDIMM Status*

Notices

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel 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.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

The products and services described may contain defects or errors known as errata which may cause deviations from published specifications. Current characterized errata are available on request.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.htm.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others

© 2015-2017 Intel Corporation.



Contents

Contents

1	Introduction	5
1.1	Document Scope	5
1.2	Related Documents	5
1.3	Terminology	5
2	_DSM Interface for NVDIMM ACPI0012 Root Device.....	7
3	_DSM Interface for the NVDIMM Device	8
3.1	SMART Health Monitoring & Alarms	11
3.1.1	Get SMART and Health Info (Function Index 1).....	11
3.1.2	Get SMART Threshold (Function Index 2).....	14
3.1.3	Set SMART Threshold (Function Index 17)	16
3.2	Command Effect Log	18
3.2.1	Get Command Effect Log Info (Function Index 7)	18
3.2.2	Get Command Effect Log (Function Index 8).....	19
3.3	Pass-Through Command (Function Index 9)	21
3.4	Enable Latch System Shutdown Status (Function Index 10).....	22
3.5	Get Supported Modes (Function Index 11)	23
3.6	NVDIMM FW Download.....	24
3.6.1	Get FW Info (Function Index 12).....	24
3.6.2	Start FW Update (Function Index 13)	26
3.6.3	Send FW Update Data (Function Index 14)	28
3.6.4	Finish FW Update (Function Index 15)	30
3.6.5	Query Finish FW Update Status (Function Index 16).....	32
3.7	Inject Error (Function Index 18).....	34
3.8	NVDIMM Security Management	37
3.8.1	Theory Of Operation.....	37
3.8.2	Get Security State (Function Index 19)	39
3.8.3	Set Passphrase (Function Index 20).....	41
3.8.4	Disable Passphrase (Function Index 21)	42
3.8.5	Unlock Unit (Function Index 22)	43



3.8.6	Freeze Lock (Function Index 23)	44
3.8.7	Secure Erase NVDIMM (Function Index 24)	45
3.8.8	Overwrite NVDIMM (Function Index 25).....	47
3.8.9	Query Overwrite NVDIMM Status (Function Index 26).....	49
3.9	Deprecated Functions	51
3.9.1	Get Block NVDIMM Flags (Function Index 3).....	51
3.9.2	Get Namespace Label Size (Function Index 4).....	52
3.9.3	Get Namespace Label Data (Function Index 5)	53
3.9.4	Set Namespace Label Data (Function Index 6)	54



1 Introduction

1.1 Document Scope

This document is targeted to writers of BIOS and OS drivers for NVDIMMs whose design adheres to the NFIT Tables in the ACPI V6.2 specification. This document specifically discusses the NVDIMM Device Specific Method (_DSM).

1.2 Related Documents

The related documents are:

- ACPI Specification Version 6.0, 6.1 & 6.2, [6.2 Errata A](http://www.uefi.org/specifications) (<http://www.uefi.org/specifications>)
- UEFI 2.7, [2.7 Errata A](http://www.uefi.org/specifications) - NVDIMM Label Protocol, UEFI 2.7 NVDIMM BTT Layout (<http://www.uefi.org/specifications>)
- This DSM Specification (<http://pmem.io/documents>)

1.3 Terminology

Refer to Table 1-1 for definitions of terms used in this document.

Table 1-1 – Terminology

Term	Description
NFIT	The NVDIMM Firmware Interface Table defines the ACPI 6.2 specified information created by the BIOS to inform the OS about NVDIMMs in the system.
NVDIMM	Non-volatile memory in a DIMM form factor.
NVDIMM Namespace Label	Labels, stored at a known location on NVDIMMs, which define the DIMM's contribution to NVDIMM Namespaces. This is a software mechanism; the DIMM itself just sees the labels as part of the overall data stored on the DIMM. See the ACPI 6.2 NVDIMM Label additions and the UEFI 2.7 NVDIMM Label Protocol additions to describe this in more detail.
NVDIMM Namespace	Similar to an NVMe Namespace or a Logical Unit (LUN) on a SCSI disk, this is a software mechanism for managing ranges of persistence on NVDIMMs. See the ACPI 6.2 NVDIMM Label additions, and the UEFI 2.7 NVDIMM Label Protocol additions to describe this in more detail.
Persistent Memory	Byte-addressable memory that retains its contents after power loss.



Term	Description
SPA	System Physical Address. A physical address on the host operating system.



2 ***_DSM Interface for NVDIMM ACPI0012 Root Device***

All Root ACPI0012 scoped _DSMs ~~are~~ and NVDIMM layout / Label interfaces are now found in the following specifications and have been removed from this document, which will now only document the Intel NVDIMM specific _DSMs.

Please see:

ACPI Specification V6.0 – Initial NVDIMM NFIT additions

ACPI Specification V6.1 – Addition of Common ARS _DSMs, Clear Uncorrectable Error _DSM

ACPI Specification V6.2 – Addition of NVDIMM Label API, ARS Error Injection _DSMs

ACPI Specification V6.2 Errata A – Addition of RAS Capabilities structure as a new NFIT structure

UEFI Specification V2.7 – See additions of NVDIMM Label Protocol and BTT Layout

UEFI Specification V2.7 Errata A – Small fixes to the Label Protocol that remove inconsistencies with existing implementations



3 ***_DSM Interface for the NVDIMM Device***

Platforms that have the _DSM interface implemented, as outlined in this section, can support a NVDIMM region with Region Format Interface Code (RFIC) of 0x0201 or 0x0301.

Note that the _DSM methods defined in this section are required to be implemented under NVDIMM devices that are child devices of the NVDIMM Root Device associated with _HID of ACPI0012 in ACPI name space hierarchy.

The following tables outlines the required Arg1, Arg2 parameters that are to be utilized for this version of the specification. The platform shall support the Arg1 - Revision Id = 1 and Arg1 - Revision Id = 2 Function Indexes simultaneously as outlined below. No other Arg1 - Revision Id values are supported at this time.

Arg0 – UUID - 4309AC30-0D11-11E4-9191-0800200C9A66

Table 3-A Supported Function Index for Arg1 - Revision Id = 1

Arg1 - Revision Id	Arg2 – Function Index	_DSM Function Name
1	0	<i>Query implemented commands per ACPI Specification (returns the list below based on Arg1 - Revision Id = 1).</i>
	1	<i>Get SMART and Health Info</i>
	2	<i>Get SMART Threshold</i>
	3	<i>Get Block NVDIMM Flags</i>
	4	<i>Deprecated - Get Namespace Label Data Size</i>
	5	<i>Deprecated - Get Namespace Label Data</i>
	6	<i>Deprecated - Set Namespace Label Data</i>
	7	<i>Get Command Effect Log Info</i>
	8	<i>Get Command Effect Log</i>
	9	<i>Pass-Through Command</i>
	10	<i>Enable Latch System Shutdown Status</i>

**Table 3-B Supported Function Index for Arg1 - Revision Id = 2**

Arg1 - Revision Id	Arg2 – Function Index	_DSM Function Name
2	0	<i>Query implemented commands per ACPI Specification (returns the list below based on Arg1 - Revision Id = 2).</i>
	1	<i>Get SMART and Health Info</i>
	2	<i>Get SMART Threshold</i>
	3	<i>Deprecated - Get Block NVDIMM Flags</i>
	4	<i>Deprecated - Get Namespace Label Data Size</i>
	5	<i>Deprecated - Get Namespace Label Data</i>
	6	<i>Deprecated - Set Namespace Label Data</i>
	7	<i>Get Command Effect Log Info</i>
	8	<i>Get Command Effect Log</i>
	9	<i>Pass-Through Command</i>
	10	<i>Enable Latch System Shutdown Status</i>
	11	<i>Get Supported Modes</i>
	12	<i>Get FW Info</i>
	13	<i>Start FW Update</i>
	14	<i>Send FW Update Data</i>
	15	<i>Finish FW Update</i>
	16	<i>Query Finish FW Update Status</i>
	17	<i>Set SMART Threshold</i>
	18	<i>Inject Error</i>
	<u>19</u>	<u><i>Get Security State</i></u>
	<u>20</u>	<u><i>Set Passphrase</i></u>
	<u>21</u>	<u><i>Disable Passphrase</i></u>
	<u>22</u>	<u><i>Unlock Unit</i></u>
	<u>23</u>	<u><i>Freeze Lock</i></u>
	<u>24</u>	<u><i>Secure Erase NVDIMM</i></u>
	<u>25</u>	<u><i>Overwrite NVDIMM</i></u>
	<u>26</u>	<u><i>Query Overwrite NVDIMM Status</i></u>



Arg3 – A package containing parameters for the function specified by the *UUID*, *Revision ID*, and *Function Index*. The layout of the package for each command along with the corresponding output are illustrated in the respective *Function Index* description sections. For DSM functions that take an input argument, *Arg3* is a package containing a Buffer, list of bytes, values. For DSM functions that do not take an input parameter, *Arg3* is an empty package. The output of all functions in the DSM is a Buffer with a list of bytes. The first four bytes provide Status and Extended Status for the DSM function. Depending on the status code, additional bytes may follow the status bytes. If status bytes signal an error condition, the additional bytes are not present, unless some additional information is explicitly defined for the particular error code. If status bytes signal success, all output bytes defined for the function are present.

The following table outlines the returned Status field common to all of the DSMs defined in this specification. The status adopts the following conventions for the *_DSM* function return status codes. This status can always be utilized for the status of each *_DSM* function, whether the specific status value is defined in the output buffer or not.

Table 3-C Supported *_DSM* Return Status Values

Return Status Value - Bytes[1-0]	Return Status <u>Value - Description</u>
0	Success
1	Failure - Function Not Supported
2	Failure - Non-Existing Memory Device
3	Failure - Invalid Input Parameters
4	Failure – HW Error
5	Failure – Retry Suggested - Command Timed Out, Other Command In Progress, Mailbox not Ready <u>Typically an operation is executing and cannot be interrupted. Operations most likely to be executing are: ARS, Overwrite NVDIMM, and Finish FW Update. Software shall wait for those operations to complete utilizing Get ARS Status, Query Overwrite NVDIMM Status, or Query Finish FW Update Status before restarting an ARS, Overwrite NVDIMM, or FW Update sequence respectively.</u>
6	Failure – Unknown Reason
7	Function Specific Error (details in Extended Status Field)
8	Failure – Retry Suggested - Out of Resources
9	Failure – HW Not Ready
<u>10</u>	<u>Failure – Invalid Security State</u>
<u>11</u>	<u>Failure – Invalid Current Passphrase Supplied - Returned by the NVDIMM when the Current Passphrase does not match the saved passphrase. If the NVDIMM is also in the wrong security state, the Invalid Security State status is reported instead of this status.</u>



3.1 SMART Health Monitoring & Alarms

3.1.1 Get SMART and Health Info (Function Index 1)

This command requests the device to return Smart and Health information for the requested device.

Function Input

None

Function Output

The following tables outline the expected output payload for this command.

Table 3-1 Get SMART and Health Info – Output Format

Field	Byte Length	Byte Offset	Description
Status	2	0	Defined above in Table 3-C
Extended Status	2	2	Extended Status Field
Smart and Health Data	128	4	Output formatted as shown in Table 3-2.

Table 3-2 SMART and Health Data – Output Format

Field	Byte Length	Byte Offset	Description
Validity Flags	4	0	Validity Flags – if the corresponding validation flag is not set in this field, it is indication to software that the corresponding field is not valid and must not be interpreted. Bit[0] – if set to 1, indicates that Health Status field is valid Bit[1] – if set to 1, indicates that Spare Blocks Remaining field is valid Bit[2] – if set to 1, indicates that Percentage Used field is valid Bit[3] – if set to 1, indicates that Current NVDIMM Media Temperature field is valid Bit[4] – if set to 1, indicates that Current NVDIMM Controller Temperature field is valid Bit[5] – If set to 1, indicates that <u>Unsafe-Dirty</u> Shutdown Count field is valid Bit[6] – If set to 1, indicates that the AIT DRAM Status field is valid Bit[7] – if set to 1, indicates that Current NVDIMM PMIC Temperature field is valid Bit[8: <u>7</u>] – Reserved, shall return 0. Bit[9] – if set to 1, indicates that Alarm Trips field is valid



			<p>Bit[10] – if set to 1, indicates that Last Shutdown Status field is valid</p> <p>Bit[11] – if set to 1, indicates that Size of Vendor-specific Data field is valid. If this field is not valid, the software will ignore the vendor-specific data fields.</p> <p>Bits[31:12] – Reserved, shall return 0.</p>
Reserved	4	4	Shall return 0.
Health Status	1	8	<p>Health Status (HS): Overall health summary. Normal health is indicated by all HS bits being clear. Only one bit will be set at a time.</p> <p>Bit[0] – if set to 1, indicates Non-Critical condition, maintenance required but no data loss detected</p> <p>Bit[1] – if set to 1, indicates Critical condition, features or performance degraded due to failures but no data loss detected</p> <p>Bit[2] – if set to 1, indicates fatal condition, data loss is detected or is imminent.</p> <p>Bit[7:3] - Reserved, shall return 0.</p>
Spare Blocks Remaining	1	9	<p>Spare Blocks Remaining: Remaining Spare Capacity as % of factory configured space.</p> <p>Valid range 0 to 100.</p> <p>0 = All of the factory configured spare block capacity has been utilized</p> <p>100 = None of the factory configured spare block capacity has been utilized</p>
Percentage Used	1	10	<p>Percentage Used: Device life span as percentage</p> <p>Valid range 0 to 100.</p> <p>100 = the warranted life span of the device has been reached.</p>
Alarm Trips	1	11	<p>Alarm Trips: Bits to signify if values have tripped their respective alarm thresholds</p> <p>Bit[0] - Spare Blocks Remaining Trip - If set then the Spare Blocks Remaining value has gone below the pre-programmed threshold limit</p> <p>Bit[1] – NVDIMM Media Temperature Trip - If set then the NVDIMM Media temperature value has gone above the pre-programmed threshold limit</p> <p>Bit[2] – NVDIMM Controller Temperature Trip - If set then the NVDIMM Controller temperature value has gone above the pre-programmed threshold limit</p> <p>Bits[7:3] - Reserved, shall return 0.</p>



Current NVDIMM Media Temperature	2	12	Current Media Temperature: Current temperature of the NVDIMM Media Bits[14:0] - Temperature in 0.0625 degree Celsius resolution. Bit[15] – Sign bit for temperature (1 = negative, 0 = positive)
Current NVDIMM Controller Temperature	2	14	Current Controller Temperature: Current temperature of the NVDIMM Controller Bits[14:0] - Temperature in 0.0625 degree Celsius resolution. Bit[15] – Sign bit for temperature (1 = negative, 0 = positive)
<u>Unsafe-Dirty</u> Shutdown Count	4	16	<u>Unsafe-Dirty</u> Shutdown Count (<u>USCDSC</u>) – Number of times the NVDIMM Last Shutdown Status (LSS) was non-zero, <u>indicating a dirty shutdown</u> . Incremented anytime Last Shutdown Status (LSS) != 0 & Latch System Shutdown Status is set by host SW (via Enable Latch System Shutdown Status _DSM) . Count wraps back to 0 at overflow.
AIT DRAM Status	1	20	AIT DRAM Status 00 – AIT DRAM is disabled 01 – AIT DRAM is enabled If the AIT DRAM is disabled, it will cause a performance degradation and will trigger a SMART Health Status change to critical state
Current NVDIMM PMIC Temperature	2	21	Current PMIC Temperature: Current temperature of the NVDIMM PMIC Bits[14:0] – Temperature in 0.0625 degree Celsius resolution. Bit[15] – Sign bit for temperature (1 = negative, 0 = positive)
Reserved	<u>10</u>	<u>21</u>	Shall return 0.
Last Shutdown Status	1	31	Last Shutdown Status (LSS): status of last shutdown 00 – Clean shutdown All other Values – Not Clean Shutdown, indicates that there was either a platform or memory device-related failure occurred when saving data targeted for this memory device. <u>Unsafe-Dirty</u> Shutdown Count (<u>USCDSC</u>) above maintains a count of the number of times a non-clean shutdown occurs. Updated when Latch System Shutdown Status is set by host SW (via Enable Latch System Shutdown Status _DSM)
Size of Vendor Specific Data	4	32	Size of Vendor-specific Data. If set to 0, indicates that there is no vendor specific data that follows. Otherwise, indicates size of the Vendor-specific data that follows.
Vendor Specific SMART Data	92	127-36	Vendor-specific SMART Data



3.1.2 Get SMART Threshold (Function Index 2)

This command requests the device to return Smart Threshold values that have been programmed by the platform for the requested device.

Function Input

None

Function Output

The following tables outline the expected output payload for this command.

Table 3-3 Get SMART Threshold – Output Format

Field	Byte Length	Byte Offset	Description
Status	2	0	Defined above in Table 3-C
Extended Status	2	2	Extended Status Field
Smart Threshold Data	8	4	Output formatted as shown in Table 3-4.



Table 3-4 SMART Threshold Data – Output Format

Field	Byte Length	Byte Offset	Description
Threshold Alarm Enable	2	0	Threshold Alarm Control – If a bit is set to 1, the specific alarm is enabled and the corresponding Alarm Trip bit in the SMART Health Status output payload will be set when a specific threshold outlined below has been reached. Bit[0] - Spare Blocks Remaining Threshold Alarm Enable Bit[1] – NVDIMM Media Temperature Threshold Alarm Enable Bit[2] – NVDIMM Controller Temperature Threshold Alarm Enable Bit[15:3] - Reserved, shall return 0
Spare Blocks Remaining Threshold	1	2	Spare Blocks Remaining Threshold: Remaining Spare Capacity as % of factory configured space. Valid range 0 to 100. If the <i>Spare Blocks Remaining Threshold Alarm Enable</i> bit is set and when the remaining spare block capacity goes below this threshold, the <i>Spare Blocks Remaining Trip</i> bit will be set in the SMART and Health Data structure defined in Table 3-2.
NVDIMM Media Temperature Threshold	2	3	Media Temperature Threshold Bit[14:0] – Temperature in 0.0625 degree Celsius resolution. Bit[15] – Sign bit for temperature (1 = negative, 0 = positive) If the <i>NVDIMM Media Temperature Threshold Alarm Valid</i> bit is enabled and when the <i>NVDIMM Media</i> temperature goes above this value, the <i>NVDIMM Media Temperature Trip</i> bit will be set in the SMART and Health Data structure defined in Table 3-2.
NVDIMM Controller Temperature Threshold	2	5	Controller Temperature Threshold Bit[14:0] - Temperature in 0.0625 degree Celsius resolution. Bit[15] - Sign bit for temperature (1 = negative, 0 = positive) If the <i>NVDIMM Controller Temperature Threshold Alarm Valid</i> bit is enabled and when the NVDIMM Controller temperature goes above this value, the <i>NVDIMM Controller Temperature Trip</i> bit will be set in the SMART and Health Data structure defined in Table 3-2.
Reserved	1	7	Shall return 0.



3.1.3 Set SMART Threshold (Function Index 17)

This command requests the device to simultaneously enable specific SMART Threshold Alarm Triggers and set the SMART Threshold Alarm Trigger values for the device. Parameter values are verified first before any enable/disable state or threshold values are updated.

Function Input

The following tables outline the expected input payload for this command.

Table 3-5 Set SMART Threshold – Input Format

Field	Byte Length	Byte Offset	Description
Threshold Alarm Enable	2	0	Threshold Alarm Control - If a bit is set to 1, the specific alarm is enabled and the corresponding Alarm Trip bit in the SMART Health Status output payload will be set when a specific threshold outlined below has been reached. Bit[0] - Spare Blocks Remaining Threshold Alarm Enable Bit[1] – NVDIMM Media Temperature Threshold Alarm Enable Bit[2] – NVDIMM Controller Temperature Threshold Alarm Enable Bit[15:3] - Reserved, shall be 0
Spare Blocks Remaining Threshold	1	2	Remaining Spare Capacity Alarm - A % of factory configured spare blocks. Values 0 & 100 are not valid and will result in an error. If the <i>Spare Blocks Remaining Threshold Alarm Enable</i> bit is set and when the spare block capacity goes below this threshold, the <i>Spare Blocks Remaining Trip</i> bit will be set in the SMART and Health Data structure defined in Table 3-2. This field is ignored if the <i>Spare Blocks Remaining Threshold Alarm Enable</i> bit above is cleared to 0.
NVDIMM Media Temperature Threshold	2	3	Media Temperature Alarm Bit[14:0] – Temperature in 0.0625 degree Celsius resolution. Bit[15] – Sign bit for temperature (1 = negative, 0 = positive) If the <i>NVDIMM Media Temperature Threshold Alarm Valid</i> bit is enabled and when the <i>NVDIMM Media</i> temperature goes above this value, the <i>NVDIMM Media Temperature Trip</i> bit will be set in the SMART and Health Data structure defined in Table 3-2. This field is ignored if the <i>NVDIMM Media Temperature Threshold Alarm Enable</i> bit above is cleared to 0.
NVDIMM Controller Temperature Threshold	2	5	Control Temperature Alarm Bit[14:0] - Temperature in 0.0625 degree Celsius resolution. Bit[15] - Sign bit for temperature (1 = negative, 0 = positive)



			<p>If the <i>NVDIMM Controller Temperature Threshold Alarm Valid</i> bit is enabled and when the NVDIMM Controller temperature goes above this value, the <i>NVDIMM Controller Temperature Trip</i> bit will be set in the SMART and Health Data structure defined in Table 3-2.</p> <p>This field is ignored if the <i>NVDIMM Controller Temperature Threshold Alarm Enable</i> bit above is cleared to 0.</p>
--	--	--	---

Function Output

The following tables outline the expected output payload for this command.

Table 3-6 Set SMART Threshold – Output Format

Field	Byte Length	Byte Offset	Description
Status	2	0	Defined above in Table 3-C 03 – Invalid Input Parameters Returned If any threshold value requested to be enabled is invalid. No changes are made to any previously set threshold enable/disable state and no changes are made to any previously set threshold values.
Extended Status	2	2	Extended Status Field



3.2 Command Effect Log

3.2.1 Get Command Effect Log Info (Function Index 7)

This command requests the device to return the Command Effect Log Information for the requested device.

Function Input

None

Function Output

The following tables outline the expected output payload for this command.

Table 3-13 Get Command Effect Log Info – Output Format

Field	Byte Length	Byte Offset	Description
Status	2	0	Defined above in Table 3-C
Extended Status	2	2	Extended Status Field
Max Command Effect Log Data Length	4	<u>4</u>	In bytes, Maximum size of the command effect log data buffer supported by the device



3.2.2 Get Command Effect Log (Function Index 8)

This command requests the device to return the Command Effect Log associated with the requested device. If the OpCode is not in the Command Effect log, OSPM may block the Pass-Through Command calls for that OpCode.

Function Input

None

Function Output

The following tables outline the expected output payload for this command.

Table 3-14 Get Command Effect Log – Output Format

Field	Byte Length	Byte Offset	Description
Status	2	0	Defined above in Table 3-C
Extended Status	2	2	Extended Status Field
OpCode Count	2	4	Number of OpCode command effect logs returned
Reserved	2	6	Shall return 0.
Command Effect Data	Max Command Effect Log Data Length from Get Command Effect Log Info	8	The command effect data for each OpCode. The Fields in Table 3-8 are repeated <i>OpCode Count</i> times.

Table 3-15 Command Effect Data – Output Format

Field	Byte Length	Byte Offset	Description
OpCode	4	0	OpCode representing a Vendor-specific command
OpCode Command Effect	4	4	Bit[0] – No Effects (NE) If set to 1, execution of this OpCode does not change DIMM state. If this bit is set, all the following bits must be clear. Bit[1] – Security State Change (SSC) If set to 1, execution of this Opcode results in immediate security state change of the NVDIMM. Bit[2] – DIMM Configuration Change after Reboot (DCC) If set to 1, execution of this Opcode results in change to the configuration of the NVDIMM or data contained within persistent memory regions of the NVDIMM. The change does not take effect until the system re-boots. Bit[3] – Immediate DIMM Configuration Change (IDCC)



			<p>If set to 1, execution of this Opcode results in immediate change to the configuration of the NVDIMM or data contained within persistent memory regions of the NVDIMM.</p> <p>Bit[4] – Quiesce All IO (QIO)</p> <p>If set to 1, execution of this Opcode may disrupt on-going operations of the memory region covered by this NVDIMM. The outstanding IO operations corresponding to this NVDIMM must be quiesced before executing this command; otherwise, undefined system behavior will result. Operations that must be quiesced include cpu load/store/move/flush memory operations, writes to NFIT Flush Hint Addresses, HW Block aperture programming sequences, in progress long operation sequences including ARSs, and NVDIMM controller mailbox commands.</p> <p>Bit[5] - Immediate DIMM Data Change (IDDC)</p> <p>If set to 1, execution of this Opcode results in immediate change to the data written to the NVDIMM.</p> <p>Bit[6] – Test Mode (TM)</p> <p>If set to 1, execution of this Opcode activates a test feature that may disrupt on-going operations. This may result in errors or error recovery operations.</p> <p>Bit[7] – Debug Mode (DM)</p> <p>If set to 1, execution of this Opcode activates a debug feature that is non-disruptive, but may alter performance characteristics of the NVDIMM.</p> <p>Bit[31:8] – Reserved, shall return 0.</p>
--	--	--	---



3.3 Pass-Through Command (Function Index 9)

This command requests the device to execute the vendor specific command contained in the input payload for the requested device.

Function Input

The following tables outline the expected input payload for this command.

Table 3-16 Pass-Through Command – Input Format

Field	Byte Length	Byte Offset	Description
OpCode	4	0	Vendor-specific command OpCode
OpCode Parameters Data Length	4	4	In bytes Length of OpCode parameters data
OpCode Parameters Data	OpCode Parameters Data Length	8	Vendor-specific command input data

Function Output

The following tables outline the expected output payload for this command.

Table 3-17 Pass-Through Command – Output Format

Field	Byte Length	Byte Offset	Description
Status	2	0	Defined above in Table 3-C
Extended Status	2	2	Extended Status Field
Output Data Length	4	4	In bytes. If <i>Status</i> is not <i>Success</i> , output data length returned is 0.
Output Data	Output Data Length	8	The <i>Output Data</i> is valid only when the <i>Output Data Length</i> is non-zero.



3.4 Enable Latch System Shutdown Status (Function Index 10)

DSM command to allow a SW agent enable the latching of SMART LSS & SMART ~~Unsafe~~ Dirty Shutdown Count state of each NVDIMM. By default the NVDIMM powers up assuming that this latch is disabled. When the latch is disabled the NVDIMM will report the previously saved value for the SMART LSS and SMART ~~D~~USC values. Those values will not change again until the next power down sequence following the enable of the latch utilizing this DSM.

Function Input

The following tables outline the expected input payload for this command.

Table 3-18 Enable Latch System Shutdown Status – Input Format

Field	Byte Length	Byte Offset	Description
Latch System Shutdown Status	1	0	Enable System Shutdown Status –Enables latching of SMART Last Shutdown Status (LSS) & SMART Unsafe Dirty Shutdown Count in NVDIMM on the next power down event. 01 – Enable the latch. Update SMART LSS & SMART Unsafe Dirty Shutdown Count on next power-down, power-up sequence All other values are reserved.

Function Output

The following tables outline the expected output payload for this command.

Table 3-19 Enable Latch System Shutdown Status – Output Format

Field	Byte Length	Byte Offset	Description
Status	2	0	Defined above in Table 3-C
Extended Status	2	2	Extended Status Field



3.5 Get Supported Modes (Function Index 11)

This command requests the platform to return details about the supported Modes of the NVDIMM Interface implementation.

Function Input

None

Function Output

The following tables outline the expected output payload for this command.

Table 3-20 Get Supported Modes – Output Format

Field	Byte Length	Byte Offset	Description
Status	2	0	Defined above in Table 3-C
Extended Status	2	2	Extended Status Field
Supported Modes	2	4	The list of the DIMMs capabilities: Bit[0] – Memory Mode supported Bit[1] – PMEM Mode supported Bit[2] – Block Aperture Mode supported Bit[15:3] – Reserved, shall return 0.



3.6 NVDIMM FW Download

3.5.13.6.1 Get FW Info (Function Index 12)

This command returns information for the limits utilized for Send FW Update Data function, the running FW image revision, the running FW image Firmware Interface Specification (FIS) version, and the Updated FW Image, if one exists.

Function Input

None

Function Output

The following tables outline the expected output payload for this command.

Table 3-21 Get FW Info – Output Format

Field	Byte Length	Byte Offset	Description
Status	2	0	Defined above in Table 3-C
Extended Status	2	2	Extended Status Field
Size of FW Update Image Storage Area	4	4	In bytes, Total size of the FW Update Image Storage Area supported by the platform.
Max Send FW Update Data Length	4	8	In bytes, Maximum Length value that can be utilized with each Send FW Update Data command.
Query Finish FW Update Status Polling Interval	4	12	Polling interval in uSecs describing how often software should issue a Query Finish FW Update Status polling command to check for Finish FW Update completion.
Max Time to Query Finish FW Update Status	4	16	Maximum time in uSec software should have to poll for Query Finish FW Update Status on a single NVDIMM.
FW Update Capabilities	1	20	Flags further defining the FW Update capabilities or features of the NVDIMM Bit[0] – FW Update Requires System Cold Re-Boot – If set the NVDIMM requires a system cold-boot for the new FW image to become the new executing FW image. This assumes that the FW update sequence has completed successfully.



			Bit[7:1] - Reserved All other values are reserved and Returned read as 0.
Reserved	3	21	Read as 0
Running FW Interface Version	4	24	The current running FW Interface Specification (FIS) revision using the product specific format. -Implementations that do not report a full 4 bytes of Running FW Interface Version information shall fill unused MSB bytes with 0's. Note: This is for informational purposes only and shall not be utilized to determine the command set that is supported by the NVDIMM.
Running FW Revision	8	28	Contains the revision information of the currently running NVDIMM firmware using the product specific format. -Implementations that do not report a full 8 bytes of Running FW Revision information shall fill unused MSB bytes with 0's. -Larger version value indicate newer FW revision.
Updated FW Revision	8	36	Upon successful completion of the Finish FW Update command this field contains the revision information of the updated NVDIMM firmware using the product specific format. -This revision becomes valid after successful completion of a Send FW Update Data & Finish FW Update sequence. This field then becomes invalid after a cold system boot and this revision shall be reported as all 0's at that time. -If no FW image has been sent or an image has been sent but the update has not been finished, or the Finish FW Update fails, then this revision shall be reported as all 0's. -Implementations that do not report a full 8 bytes of Updated FW Revision information shall fill unused MSB bytes with 0's. -Larger version value indicate newer FW revision.



3.5.23.6.2 Start FW Update (Function Index 13)

This command requests the NVDIMM device to start a FW download sequence. The FW download sequence consists of a single Start FW Update, followed by one or more Send FW Update Data commands and completes with a single Finish FW Update command followed by one or more Query Finish FW Update Status to poll for Finish FW Update completion.

Function Input

None

Function Output

The following tables outline the expected output payload for this command.

Table 3-22 Start FW Update – Output Format

Field	Byte Length	Byte Offset	Description
Status	2	0	Defined above in Table 3-C 00 – Success. FW Update Context field is valid. 05 – Failure – Retry Suggested – ARS in progress. Software shall wait for ARS completion utilizing Get ARS Status, before starting a FW update sequence. 07 – Function Specific Status (see Extended Status below) 08 - Failure – Retry Suggested - Out of Resources. Software may need to complete other outstanding FW update sequences, potentially for other NVDIMM devices before retrying the Start FW Update command. It is also possible to abort other FW update sequences in progress to recover internal platform resources, using the Control Flags in the Finish FW Update input payload.
Extended Status	2	2	Extended Status Field 01 – FW Update already in progress for this NVDIMM device. The FW Update Context field returned is valid and indicates the context for the currently executing FW Update on the NVDIMM device. Software must complete the current FW update sequence with one of the two methods: -Sending a Finish FW Update command and possibly a system cold re-boot before another FW update sequence can be started on the same NVDIMM -Using the returned FW Update Context to abort the existing FW Update that is in progress by calling Finish FW Update with the Control Flag set to Abort Existing FW Update Sequence 02 – FW Update already occurred – A successful FW update sequence has already occurred and another Start FW Update command is being attempted without a system cold-boot.



NVDIMM DSM Interface – V1.7

FW Update Context	4	4	Upon successful completion of the Start FW Update command this field contains a platform implementation specific value that must be passed as an input parameter to Send FW Update Data and Finish FW Update commands.
-------------------	---	---	--



3.5.33.6.3 Send FW Update Data (Function Index 14)

This command requests the device to update the FW image in the NVDIMMs FW Update Image Storage Area as part of a FW download sequence. The FW download sequence consists of a single Start FW Update, followed by one or more Send FW Update Data commands and completes with a single Finish FW Update command followed by one or more Query Finish FW Update Status to poll for Finish FW Update completion.

The Offset and Length fields allow software to divide the FW image in to pieces based on the Max Send FW Update Data Length reported in the Get FW Info output payload. There is no ordering restriction regarding how the pieces of the FW image are sent to the NVDIMMs FW Update Image Storage Area.

No validation of the FW image occurs until the FW download sequence is complete. The FW image is considered complete and its validity is verified only after the Finish FW Update command has completed.

If software is aborting a FW Update sequence that is already in progress it can call Finish FW Update directly without issuing any Send FW Update Data commands. See the Control Flags in the Finish FW Update command for details on aborting an outstanding FW Update sequence.

Function Input

The following tables outline the expected input payload for this command.

Table 3-23 Send FW Update Data – Input Format

Field	Byte Length	Byte Offset	Description
FW Update Context	4	0	Platform specific FW update sequence context provided by the platform as part of the Start FW Update output payload.
Offset	4	4	In bytes Indicates the byte offset in the NVDIMMs FW Update Image Storage Area where this portion of the FW Image data will be written
Length	4	8	In bytes Indicates the number of bytes to be written starting at the Offset specified above
FW Image Data	Length	12	FW Image data to be written at the starting Offset for Length bytes



Function Output

The following tables outline the expected output payload for this command.

Table 3-24 Send FW Update Data – Output Format

Field	Byte Length	Byte Offset	Description
Status	2	0	Defined above in Table 3-C 03 – Invalid Input Parameters - Offset + Length is > Size of FW Update Image Storage Area reported in the Get FW Info command - Length is > Max Send FW Update Data Length reported in the Get FW Info command <u>- Length does not match the size of the ACPI input package contained in Arg3</u> 07 – Function Specific Status (see Extended Status below) 08 - Failure – Out of Resources. Software may need to complete other outstanding FW update sequences, potentially for other NVDIMM devices before retrying the Start FW Update command. It is also possible to abort other FW update sequences in progress to recover internal platform resources, using the Control Flags in the Finish FW Update input payload.
Extended Status	2	2	Extended Status Field 01 – FW Update Context invalid



3.5.43.6.4 Finish FW Update (Function Index 15)

This command requests the NVDIMM device to begin the process of finishing a FW download sequence. The FW download sequence consists of a single Start FW Update, followed by one or more Send FW Update Data commands and completes with a single Finish FW Update command followed by one or more Query Finish FW Update Status to poll for Finish FW Update completion.

Upon successful completion of this command, the NVDIMM has begun the process of finishing the FW update process. This consists of decrypting the FW image header, verifying header information including checksum, and saving the FW image in the internal NVDIMM FW Image Storage Area. This can take seconds to complete, requiring the use of the Query Finish FW Update Status so that applications can poll for Update FW completion without waiting for the update to be completed by the NVDIMM.

Software must issue the Query Update FW Status command to poll for Update FW completion. The Update FW image sequence is not complete until the query command returns proper status indicating the Update FW process is complete.

The Control Flags allow software to abort an existing FW Download instead of completing the sequence. Aborting a FW download sequence results in no change to the NVDIMM FW image. If aborting a FW Update sequence, software does not send the Query Finish FW Update command.

Function Input

The following tables outline the expected input payload for this command.

Table 3-25 Finish FW Update – Input Format

Field	Byte Length	Byte Offset	Description
Control Flags	1	0	Finish FW Update Control Flags 00 – Finish the FW Update sequence. Once software instructs the platform to finish the FW Update, it is not possible to abort the Finish FW Update sequence at a later date. Software needs to wait for the FW Update to complete using the Query Finish FW Update Status. 01 – Abort Existing FW Update Sequence. The FW Update Context describes an existing FW Download sequence that shall be aborted without updating the FW image on the NVDIMM. When aborting an active FW Update sequence, software does not call Query Finish FW Update Status. All other values are reserved.
Reserved	3	1	Must be 0
FW Update Context	4	4	Platform specific FW update sequence context provided by the platform as part of the Start FW Update output payload.



Function Output

The following tables outline the expected output payload for this command.

Table 3-26 Finish FW Update – Output Format

Field	Byte Length	Byte Offset	Description
Status	2	0	<p>Defined above in Table 3-C</p> <p>00 – Success – The Finish FW Update sequence has started. Software shall call Query Finish FW Update Status command to poll for FW Update sequence completion</p> <p><u>05 - Failure – Retry Suggested - Command Timed Out, Other Command In Progress, Mailbox not Ready</u></p> <p>05 – Failure – ARS in progress. Software shall wait for ARS completion utilizing Get ARS Status, before retrying a FW update sequence.</p> <p>07 – Function Specific Status (see Extended Status below)</p> <p>08 - Failure – Out of Resources. Software may need to complete other outstanding FW update sequences, potentially for other NVDIMM devices before retrying the Start FW Update command. It is also possible to abort other FW update sequences in progress to recover internal platform resources, using the Control Flags in the Finish FW Update input payload.</p> <p><u>09 - Failure – HW Not Ready</u></p>
Extended Status	2	2	<p>Extended Status Field - Any non-zero value returned here means the FW Update sequence is not active. Software does not need to call Query Finish FW Update Status for any of these cases.</p> <p>01 – FW Update Context invalid</p> <p>02 – FW Update already occurred – A successful FW update sequence has already occurred and another Finish FW Update command is being attempted without a system cold-boot.</p> <p>03 – Current updated FW Image failed authentication checks – fallback to prior FW image</p> <p>04 – FW update sequence successfully aborted. Only returned if the caller requested a FW Update sequence to be aborted by setting Control Flags to Abort Existing FW Update Sequence.</p>



3.5.53.6.5 Query Finish FW Update Status (Function Index 16)

This command allows software to poll for completion of the FW download sequence. The FW download sequence consists of a single Start FW Update, followed by one or more Send FW Update Data commands and completes with a single Finish FW Update command followed by one or more Query Finish FW Update Status to poll for Finish FW Update completion.

Finish FW Update consists of decrypting the FW image header, verifying header information including checksum, and saving the FW image in the internal FW Image Storage Area. This can take seconds to complete requiring the use of the Query Finish FW Update Status so that applications can poll for completion without the BIOS blocking in SMM waiting for the update to be completed by the NVDIMM. The Query Finish FW Update Status Polling Interval returned in the Get FW Info command specifies what frequency software should utilize when polling for Finish FW Update completion using the Query Finish FW Update Status command.

Upon successful completion of this command, the updated FW image will become the new executing FW image on the next system cold re-boot, replacing the currently executing FW image.

Sending a Finish FW Update followed by one or more Query Finish FW Update Status commands completes the FW download sequence and requests the NVDIMM to verify the Updated FW Image and report the revision information for the Updated FW Image. If no updated FW image is sent or the updated FW image is incomplete, Query Finish FW Update Status command will return an appropriate error and the Updated FW Image Revision will be reported as all 0's.

Only a single FW Update sequence can be handled per NVDIMM per system cold-boot sequence. Once successful status is returned for Query Finish FW Update Status, the system must go through a cold-boot cycle before another FW Update sequence can be executed on that same NVDIMM. Multiple NVDIMMs can have FW images updated and utilize a single system cold-boot to activate the new FW image on all NVDIMMs.

**Function Input**

The following tables outline the expected input payload for this command.

Table 3-27 Query Finish FW Update Status – Input Format

Field	Byte Length	Byte Offset	Description
FW Update Context	4	0	Platform specific FW update sequence context provided by the platform as part of the Start FW Update output payload.

Function Output

The following tables outline the expected output payload for this command.

Table 3-28 Query Finish FW Update Status – Output Format

Field	Byte Length	Byte Offset	Description
Status	2	0	Defined above in Table 3-C 00 – Success – The Update FW sequence has completed successfully. Authentication checks passed. Updated FW Revision field is valid. The Updated FW Image will be loaded on the next system cold-boot. 07 – Function Specific Status (see Extended Status below) 08 - The Finish FW Update sequence timed out
Extended Status	2	2	Extended Status Field 01 – FW Update Context invalid 02 – FW Update in progress 03 – Current updated FW Image failed authentication checks – fallback to prior FW image 04 – Sequencing Error – Query Finish FW Update Status called without first calling Finish FW Update
Updated FW Revision	8	4	Upon successful completion of the Finish FW Update command this field contains the revision information of the updated NVDIMM firmware using the product specific format. -This becomes valid after successful completion of a Send FW Update Data & Finish FW Update sequence. This field then becomes invalid after a cold system re-boot. -If no FW image has been updated or the updated FW image is invalid, or the Finish FW Update fails, then this revision shall be reported as all 0's. -Implementations that do not report a full 8 bytes of Updated -FW Revision information shall fill unused MSB bytes with 0's. -Larger version value indicates newer FW revision.



3.63.7 Inject Error (Function Index 18)

Inject NVDIMM specific errors not covered by the ACPI ARS Error Inject function. None of the injected errors are persistent across power cycles or re-boots unless otherwise stated below. An error will stay injected until disabled using this command or the system is restarted, unless otherwise stated below.

Function Input

The following tables outline the expected input payload for this command.

Table 3-29 Inject Error - Input Format

Field	Byte Length	Byte Offset	Description
Error Inject Validity Flags	8	0	Valid Fields – if the corresponding validation flag is not set in this field, it is indication to software that the corresponding field is not valid and must not be interpreted. Bit[0] – if set to 1, indicates that all Media Temperature Error Inject fields are valid Bit[1] – if set to 1, indicates that all Spare Blocks Remaining Trigger fields are valid Bit[2] – if set to 1, indicates that all Fatal Error Trigger fields are valid Bit[3] – if set to 1, indicates that all Unsafe-Dirty Shutdown Error Trigger fields are valid Bit[63:4] – Reserved, shall be 0
Media Temperature Error Inject	3	8	Media Temperature Error Inject fields - This will override the NVDIMM from reading the actual temperature of the media device and spoof a media temperature reading of the injected value instead. Byte[0] Bit[0] – Enable If 0, injecting Media Temperature Errors is disabled. If 1, the Media Temperature specified will be injected. Bit[7:1] - Reserved, shall be 0. Byte[2:1] - Media Temperature to Inject Bit[14:0] – Temperature in Celsius with 0.0625 resolution Bit[15] – Sign Bit, if 1 the Temperature is negative, if 0 the temperature is positive Note: Although actions taken due to the Media Temperature injected may cause adverse effects on the NVDIMM, including IO throttling, the media temperature injected is an artificial temperature and will not cause harm to the NVDIMM. If the



			critical shutdown temperature, or higher, is injected, the NVDIMM may shutdown in order to preserve the part and data.
Spare Blocks Remaining Inject	2	11	<p>Spare Blocks Remaining Trigger - This will spoof the NVDIMM to trigger either:</p> <ul style="list-style-type: none"> -User Configured Spare Blocks Remaining Alarm for a previously set value using the Set SMART Threshold function -SMART Health Change Notification for Health Status Non-Critical or Critical <p>Byte[0]</p> <p>Bit[0] – Enable</p> <p>If 0, injecting Spare Blocks Remaining is disabled</p> <p>If 1, the Spare Blocks Remaining will be injected</p> <p>Bit[7:1] – Reserved, shall be 0.</p> <p>Byte[1] – Spare Blocks Remaining to inject. Valid values are 0-99. All other values are reserved and will result in returned Status of Invalid Input Parameters.</p> <p>Note: For this trigger to inject a User Configured Spare Block Alarm Threshold Trigger requires the Spare Block Alarm Threshold to be previously enabled using the Set SMART Threshold function. If the Spare Block Alarm Threshold has not been enabled, this function will inject SMART Health Change notification ACPI Notification 0x81 as follows:</p> <p>Spare Blocks Remaining of 1% - Causes Health Status to change to Non-Critical</p> <p>Spare Blocks Remaining of 0% - Causes Health Status to change to Critical</p>
Fatal Error Inject	1	13	<p>Fatal Error Trigger – This trigger will spoof the NVDIMM to trigger a fatal NVDIMM error. Injecting this error will result in a change to the SMART Health Info – Health Status of fatal.</p> <p>Bit[0] – Enable</p> <p>If 0, injecting Fatal Error Trigger is disabled</p> <p>If 1, a Fatal Error Trigger will be injected</p> <p>Bit[7:1] – Reserved, shall be 0</p>
Unsafe-Dirty Shutdown Error Inject	1	14	<p>Unsafe-Dirty Shutdown Error Trigger – This trigger will spoof an ADR or system shutdown failure on the next power down as follows:</p> <ul style="list-style-type: none"> -Enable SMART Last Shutdown Status (LSS) and Unsafe-Dirty Shutdown Count (DUSC) increment via the Enable Latch System Shutdown Status DSM with Bit[0] - Enable System Shutdown Status set -Power down the system – The device spoofs a failure and



			<p>latches SMART LSS, increments SMART DUSC</p> <p>-Power the system up – SMART Health Change is reported with non-zero LSS ad incremented DUSC</p> <p>Bit[0] – Enable</p> <p> If 0, injecting ADR Failure is disabled</p> <p> If 1, an ADR Failure will be injected</p> <p>Bit[7:1] – Reserved, shall be 0</p>
--	--	--	---

Function Output

The following tables outline the expected output payload for this command.

Table 3-30 Inject Error Data – Output Format

Field	Byte Length	Byte Offset	Description
Status	2	0	Defined above in Table 3-C 03 – Invalid Input Parameters Returned If any Error Inject parameter value requested is invalid. No changes are made to any previous enable/disable Error Injection state and no changes are made to any previously set Error Inject values.
Extended Status	2	2	Extended Status Field 01 – Platform not enabled for error injection. Error Injection must be enabled on the platform before attempting to inject NVDIMM specific errors.



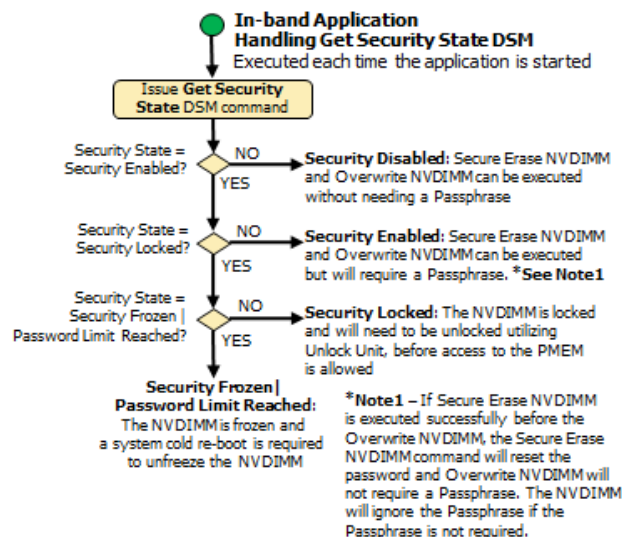
3.73.8 NVDIMM Security Management

3.7.13.8.1 Theory Of Operation

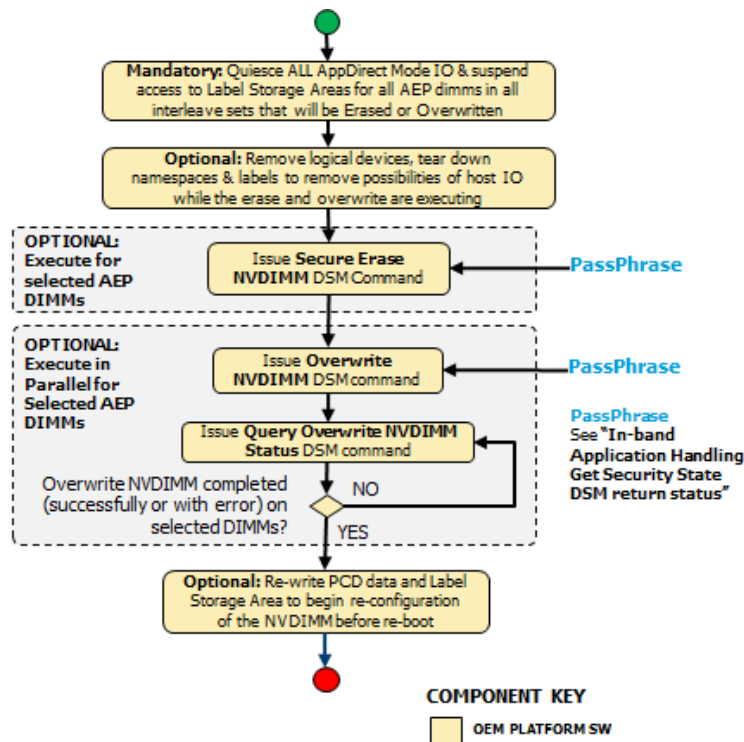
The following sequences outline the supported DSM based Secure Erase NVDIMM and Overwrite NVDIMM execution flow that utilizes an in-band application. Note the following notes regarding the in-band DSM based security implementation:

- The NVDIMM implements a security model similar to the legacy SATA/SCSI ATA security model utilized with HDDs and SSDs.
- To support this, the NVDIMM requires a pass phrase to enable security on the NVDIMM, disable of previously enabled NVDIMM security, Secure Erase NVDIMM and Overwrite NVDIMM requests.
- Removing the logical devices from access by the running OS while the Secure Erase and Overwrite NVDIMM operations are executed is recommended to remove any interactions with host IO while the erase or overwrite are executing.
- Speculative reads from the mapped in PMEM will pollute cpu caches with all 1's data for a locked NVDIMM. It is imperative that the system either be restarted before first read access, OR cpu caches are invalidated before the first read access is allowed, after unlocking the NVDIMM. The DSM commands do NOT invalidate cpu caches.
- The NVDIMM allows access to the Label Storage Area and PMEM after the Overwrite NVDIMM completes and before the system has been restarted with a cold system re-boot. This allows optional re-configuration of the NVDIMM, including the initial re-write of the Label Storage Area to occur before the first reboot in the configuration process.

The following figures shows the basic handling of the Get Security State DSM each time the application is executed, the Passphrase requirements when executing a Secure Erase NVDIMM or Overwrite NVDIMM command, and the basic Secure Erase & Overwrite NVDIMM execution flow.



Application In-band Get Security State Sequence



Application In-band Secure Erase & Overwrite NVDIMM Sequence



3.8.2 Get Security State (Function Index 19)

Retrieve the current security state of the NVDIMM.

Function Input

None

Function Output

The following tables outline the expected output payload for this command.

Table 3-31 Get Security State – Output Format

Field	Byte Length	Byte Offset	Description
Status	2	0	Defined above in Table 3-C 00 – Success. The security command executed successfully and the returned Security State is valid. 05 - Failure – Retry Suggested - Command Timed Out, Other Command In Progress, Mailbox not Ready 09 – Failure – HW Not Ready
Extended Status	2	2	Extended Status Field Undefined
Reserved	4	4	Returned as 0
Security State	1	8	Current NVDIMM Security State NO BITS SET: Security Supported, Security Disabled Bit[0] – Reserved - Returned as 0 Bit[1] – Security Enabled - When Set to 1, security is enabled on the NVDIMM. Security Disabled - When Clear to 0, the PMEM can be accessed without unlocking the NVDIMM. Overwrite NVDIMM, Secure Erase NVDIMM and security commands that normally require a Passphrase, can be executed without a Passphrase. Bit[2] – Security Locked - When set to 1, the NVDIMM is locked. Access to the Label Storage Area and the PMEM is only allowed by supplying the correct Passphrase. Changing the configuration is not allowed. To enable media access the NVDIMM media is unlocked by executing a Unlock Unit command and supplying the required Passphrase Bit[3] – Security Frozen - When set to 1, the current Security State is frozen. The NVDIMM may freeze its internal security state to protect itself from a denial-of-service attack. Once frozen, all security DSMs, except Get Security State and Query Overwrite NVDIMM Status, will be rejected, while other management interfaces and DSMs are not affected. If this occurs, a cold system re-boot, that includes power cycling of the NVDIMM, is required



			<p><u>before the security state of the NVDIMM can be changed utilizing these DSMs. The NVDIMM is not Freeze Locked by default.</u></p> <p><u>Bit[4] – Password Limit Reached</u> - When set to 1, the NVDIMM is frozen due to password limit reached. While this is reported as a different state than the frozen state (Bit[3]), the NVDIMM will freeze its internal security state when the password limit has been reached, to protect itself from a denial-of-service attack. Once frozen, all security DSMs, except Get Security State and Query Overwrite NVDIMM Status, will be rejected, while other management interface and DSMs are not affected. If this occurs, a cold system re-boot, that includes power cycling of the NVDIMM, is required before the security state of the NVDIMM can be changed utilizing these DSMs..</p> <p><u>Bit[5] – Security Not Supported</u> - When set to 1, the NVDIMM security feature is not supported by this HW</p> <p><u>Bit[6] – Reserved</u> - Returned as 0</p> <p><u>Bit[7] – Reserved</u> - Returned as 0</p>
<u>Reserved</u>	<u>3</u>	<u>9</u>	<u>Returned as 0</u>



3.8.3 Set Passphrase (Function Index 20)

Allows the caller to enable security on the NVDIMM.

Function Input

The following tables outline the expected input payload for this command.

Table 3-31 Set Passphrase – Input Format

<u>Field</u>	<u>Byte Length</u>	<u>Byte Offset</u>	<u>Description</u>
<u>Current Passphrase</u>	<u>32</u>	<u>0</u>	The end-user passphrase supplied when the user enables security on the NVDIMM using this command. -If the NVDIMM is in the Security Enabled state, this must match the passphrase that was supplied in the New Passphrase field on last successful invocation of the Set Passphrase command. -If the NVDIMM is in the Security Disabled state this field is ignored
<u>New Passphrase</u>	<u>32</u>	<u>32</u>	The new end-user passphrase. Required field utilized as part of setting a new passphrase on the NVDIMM.

Function Output

The following tables outline the expected output payload for this command.

Table 3-32 Set Passphrase – Output Format

<u>Field</u>	<u>Byte Length</u>	<u>Byte Offset</u>	<u>Description</u>
<u>Status</u>	<u>2</u>	<u>0</u>	Defined above in Table 3-C 00 – Success. The security command executed successfully. 07 – Function Specific Status (see Extended Status below) 10 - Failure – Invalid Security State 11 - Failure – Invalid Current Passphrase Supplied
<u>Extended Status</u>	<u>2</u>	<u>2</u>	Extended Status Field Undefined



3.8.4 Disable Passphrase (Function Index 21)

Disable security on the NVDIMM. If security is not enabled via Set Passphrase, it is not required to disable security.

Function Input

The following tables outline the expected input payload for this command.

Table 3-33 Disable Passphrase – Input Format

<u>Field</u>	<u>Byte Length</u>	<u>Byte Offset</u>	<u>Description</u>
<u>Current Passphrase</u>	<u>32</u>	<u>0</u>	<u>The end-user passphrase supplied when the user disables security on the NVDIMM using this command.</u> <u>-This field is required if the NVDIMM returns Security Enabled state for Get Security State command. See Theory Of Operation, Application In-band Get Security State Sequence for this logic.</u>

Function Output

The following tables outline the expected output payload for this command.

Table 3-34 Disable Passphrase – Output Format

<u>Field</u>	<u>Byte Length</u>	<u>Byte Offset</u>	<u>Description</u>
<u>Status</u>	<u>2</u>	<u>0</u>	<u>Defined above in Table 3-C</u> <u>00 – Success. The security command executed successfully.</u> <u>10 - Failure – Invalid Security State</u> <u>11 - Failure – Invalid Current Passphrase Supplied</u>
<u>Extended Status</u>	<u>2</u>	<u>2</u>	<u>Extended Status Field</u> <u>Undefined</u>



3.8.5 Unlock Unit (Function Index 22)

Unlock the NVDIMM and allow configuration changes. To make changes to the configuration of the NVDIMM, it is required to unlock the NVDIMM using this command.

NOTE: NVDIMM users wishing to utilize the Clear Uncorrectable Error DSM found in the ACPI specification must unlock the NVDIMM first utilizing this command.

Function Input

The following tables outline the expected input payload for this command.

Table 3-35 Unlock Unit – Input Format

<u>Field</u>	<u>Byte Length</u>	<u>Byte Offset</u>	<u>Description</u>
<u>Current Passphrase</u>	<u>32</u>	<u>0</u>	The end-user passphrase supplied when the user unlocks the NVDIMM using this command. -This field is required if the NVDIMM returns Security Locked state for Get Security State command. See Theory Of Operation, Application In-band Get Security State Sequence for this logic.

Function Output

The following tables outline the expected output payload for this command.

Table 3-36 Unlock Unit – Output Format

<u>Field</u>	<u>Byte Length</u>	<u>Byte Offset</u>	<u>Description</u>
<u>Status</u>	<u>2</u>	<u>0</u>	Defined above in Table 3-C 00 – Success. The security command executed successfully. 10 - Failure – Invalid Security State 11 - Failure – Invalid Current Passphrase Supplied
<u>Extended Status</u>	<u>2</u>	<u>2</u>	Extended Status Field Undefined



3.8.6 Freeze Lock (Function Index 23)

Lock the current security state of the NVDIMM, requiring a cold system re-boot before further attempts to change the state of the NVDIMM. After successful execution of this command, Get Security State shall report **Security Frozen** state.

Function Input

None

Function Output

The following tables outline the expected output payload for this command.

Table 3-37 Freeze Lock – Output Format

<u>Field</u>	<u>Byte Length</u>	<u>Byte Offset</u>	<u>Description</u>
<u>Status</u>	<u>2</u>	<u>0</u>	Defined above in Table 3-C <u>00 – Success. The security command executed successfully.</u> <u>10 - Failure – Invalid Security State</u>
<u>Extended Status</u>	<u>2</u>	<u>2</u>	<u>Extended Status Field</u> <u>Undefined</u>



3.8.7 Secure Erase NVDIMM (Function Index 24)

Instruct the NVDIMM to generate a new encryption key for the PMEM which crypto-scrambles/randomizes all existing user data in the PMEM. The Label Storage Area is unaffected and the existing configuration is preserved. The erase typically takes seconds or less to complete.

NOTE: There are significant responsibilities placed on the caller utilizing Secure Erase. See the additional notes at the end of this section.

Function Input

The following tables outline the expected input payload for this command.

Table 3-38 Secure Erase NVDIMM – Input Format

Field	Byte Length	Byte Offset	Description
Current Passphrase	32	0	The end-user passphrase supplied when the user secure erases the NVDIMM using this command. -This field is required if the NVDIMM returns Security Enabled state for Get Security State command. See Theory Of Operation, Application In-band Get Security State Sequence for this logic.

Function Output

The following tables outline the expected output payload for this command.

Table 3-39 Secure Erase NVDIMM – Output Format

Field	Byte Length	Byte Offset	Description
Status	2	0	Defined above in Table 3-C 00 – Success. The security command executed successfully. 10 - Failure – Invalid Security State 11 - Failure – Invalid Current Passphrase Supplied
Extended Status	2	2	Extended Status Field Undefined



Further application responsibilities when utilizing Secure Erase NVDIMM:

- Host IO to the PMEM shall be quiesced and new IO blocked while this security command is executing for the NVDIMM
- If the application will also execute an Overwrite NVDIMM command, the Secure Erase NVDIMM shall be executed before the Overwrite NVDIMM
- If this is not followed with an Overwrite NVDIMM sequence, cpu caches must be invalidated before the NVDIMM is returned to operation.
- While the Label Storage Area is preserved, any address abstraction info blocks in PMEM are destroyed.



3.8.8 Overwrite NVDIMM (Function Index 25)

The Overwrite NVDIMM operation overwrites the entire PMEM of the NVDIMM with encrypted 0's including retired ECC blocks, NVM die that were previous swapped out as part of a die sparing operation, and all user addressable PMEM locations. After the overwrite process has completed, the Label Storage Area is overwritten with 0's and all configuration data is lost.

This operation can take a significant amount of time requiring the use of the Query Overwrite NVDIMM Status to poll for the completion of the Overwrite operation. SW shall call Query Overwrite NVDIMM Status at least once to verify completion of the operation. These interfaces allow parallel Overwrite NVDIMM operations to be started on all NVDIMMs and parallel polling for completions. Since this operation takes a long amount of time to complete, it is recommended to overwrite all of the NVDIMMs in the same session.

The Overwrite NVDIMM operation may not be supported for all NVDIMM configurations. If an overwrite operation is requested for an unsupported overwrite configuration, the Overwrite NVDIMM request shall be failed with an explicit extended status, Failure – Unsupported Overwrite Configuration.

NOTE: If Secure Erase NVDIMM is executed successfully before the Overwrite NVDIMM, the Secure Erase NVDIMM command will reset the password and Overwrite NVDIMM will not require a Passphrase. The NVDIMM will ignore the Passphrase argument in the input payload if the Passphrase is not required.

NOTE: There are responsibilities placed on the caller utilizing Overwrite NVDIMM. See the additional notes at the end of this section.

Function Input

The following tables outline the expected input payload for this command.

Table 3-40 Overwrite NVDIMM – Input Format

<u>Field</u>	<u>Byte Length</u>	<u>Byte Offset</u>	<u>Description</u>
<u>Current Passphrase</u>	<u>32</u>	<u>0</u>	The end-user passphrase supplied when the user overwrites the NVDIMM using this command. -This field is required if the NVDIMM returns Security Enabled state for Get Security State command. See Theory Of Operation, Application In-band Get Security State Sequence for this logic.

**Function Output**

The following tables outline the expected output payload for this command.

Table 3-41 Overwrite NVDIMM – Output Format

<u>Field</u>	<u>Byte Length</u>	<u>Byte Offset</u>	<u>Description</u>
<u>Status</u>	<u>2</u>	<u>0</u>	Defined above in Table 3-C <u>00 – Success. The security command executed successfully.</u> <u>This status indicates that the overwrite operation was started.</u> <u>Utilize Query Overwrite NVDIMM Status security command to</u> <u>periodically poll for command completion.</u> <u>07 – Function Specific Status (see Extended Status below)</u> <u>10 - Failure – Invalid Security State</u> <u>11 - Failure – Invalid Current Passphrase Supplied</u>
<u>Extended Status</u>	<u>2</u>	<u>2</u>	<u>Extended Status Field</u> <u>01 – Failure – Unsupported Overwrite Configuration</u>

Further application responsibilities when utilizing Overwrite NVDIMM:

- At a minimum, host IO to PMEM of the NVDIMM shall be quiesced and new IO blocked while this security command is executing for the NVDIMM
- If the application will also execute a Secure Erase NVDIMM command, the Overwrite NVDIMM SHALL be executed AFTER the Secure Erase NVDIMM
- After SW has issued the Overwrite NVDIMM DSM request and it has returned successful status, SW is required to poll for Overwrite operation completion by calling Query Overwrite NVDIMM Status command periodically until the command completes
- The NVDIMM's internal configuration information is overwritten and all partition info & interleave set configuration info is lost. The NVDIMM may lose its relationship to its current interleave set after the completion of the Overwrite NVDIMM command.
- The Label Storage Area is overwritten and all label info blocks and labels are destroyed.



3.8.9 Query Overwrite NVDIMM Status (Function Index 26)

After SW has issued the Overwrite NVDIMM DSM request and it has returned successful status, SW is required to poll for Overwrite operation completion by calling this command periodically until successful status or failed status is reported. As long as “Overwrite NVDIMM In Progress” is reported for extended status, SW shall continue to poll for completion. Once the Query Overwrite NVDIMM Status operation completes with success status, optionally, it is possible to begin a new configuration sequence to write new NVDIMM configuration information, including the Label Storage Area data, before the first re-boot in the configuration process.

NOTE: There are responsibilities placed on the caller utilizing Query Overwrite NVDIMM Status. See the additional notes at the end of this section.

Function Input

None

Function Output

The following tables outline the expected output payload for this command.

Table 3-42 Query Overwrite NVDIMM Status – Output Format

<u>Field</u>	<u>Byte Length</u>	<u>Byte Offset</u>	<u>Description</u>
<u>Status</u>	<u>2</u>	<u>0</u>	Defined above in Table 3-C <u>00 – Success. The security command executed successfully.</u> <u>This status indicates the Overwrite NVDIMM operation completed successfully.</u> <u>07 – Function Specific Status (see Extended Status below)</u>
<u>Extended Status</u>	<u>2</u>	<u>2</u>	<u>Extended Status Field</u> <u>01 – Overwrite NVDIMM in Progress. Software should continue to poll for overwrite completion until a success or failure status is returned</u> <u>02 – Sequencing Error – Query Overwrite NVDIMM Status called without first calling Overwrite NVDIMM</u>



Further application responsibilities when utilizing Query Overwrite NVDIMM Status:

- Host IO to the PMEM of the NVDIMM shall be quiesced and new IO blocked while this security command continues to report “Overwrite NVDIMM In Progress”.
- If the Overwrite NVDIMM fails to complete with success the state of the NVDIMM is indeterminate and it is recommended that the Overwrite be executed again.
- The Overwrite NVDIMM operation takes a significant amount of time to complete: 15min - 128GB, 30min - 256GB, 60min – 512GB, etc. There for, it is recommended that SW poll for completion with an interval of 10-60 seconds.. These values can also be used as a guide as to when SW may want to give up on polling for successful completion. Implementations may wish to issue the first query with a small timeout to make sure the Overwrite NVDIMM is still executing before moving to a longer polling frequencies for the rest of the query calls.

Note: It is possible to issue other DSMs while the Overwrite is executing so SW can monitor the DIMM health and temperature (for example) while the Overwrite DIMM executes. Other operations (ARS, FW Update) are not allowed and those DSMs will be rejected.



3.83.9 Deprecated Functions

~~3.8.1~~3.9.1 Get Block NVDIMM Flags (Function Index 3)

This function that is only applicable if block mode is enabled in the NVDIMM (i.e., the Number of Block Control Windows field set is set to a non-zero value in the NVDIMM Control Region Structure). Used by the NVDIMM to report specific features or alternative sequences that need to be implemented by SW drivers.

Warning: This function has been deprecated. It is included here for backwards compatibility with existing Arg1 - Revision Id = 1 implementations.

Function Input

None

Function Output

The following tables outline the expected output payload for this command.

Table 3-7 Get Block NVDIMM Flags - Output Format

Field	Byte Length	Byte Offset	Description
Status	2	0	Defined above in Table 3-C
Extended Status	2	2	Extended Status Field
NVDIMM Flags	4	4	Byte[0] Bit[0] – Block Data Window Invalidation Required – If this bit is set to 1, indicates that the NVDIMM requires the driver to flush previous data from cache lines that will be moved through the Block Data Window, before re-using the Block Data Window for read. If set to '0', flushing of previous data from cache lines that will be moved through the Block Data Window are handled by the platform or VMM. Typical usage of this flag is in a virtualized environment. Bit[1] – Command Register in Block Control Window Latch – If this bit is set to 1, indicates that after a write to the Command Register in Block Control Windows, the NVDIMM requires the software to read the same Command Register to ensure that the command is latched before reading contents from Block Data Window. If this bit is set to 0, software is allowed to read the contents of the Block Data Window immediately after writing to the Command Register of Block Control Window. Bits[7:2] – Reserved, shall return 0 Byte[3:1] – Reserved, shall return 0



3-8-23.9.2 Get Namespace Label Size (Function Index 4)

This command requests the device to return the size of the Namespace Label storage area for the requested device.

Warning: This function has been deprecated in preference to the ACPI 6.2 _LSI (Label Storage Information) NVDIMM Device Interface and is only supported with Arg1 – Revision Id = 1. It is included here for backwards compatibility with existing Arg1 - Revision Id = 1 implementations.

Function Input

None

Function Output

The following tables outline the expected output payload for this command. See **updated/new additions & clarifications** below for this existing LSM.

Table 3-8 Get Namespace Label Size – Output Format

Field	Byte Length	Byte Offset	Description
Status	2	0	Defined above in Table 3-C
Extended Status	2	2	01 – Extended Success Status - Locked Persistent Memory Region – The PMEM is currently in a locked state. This DSM is expected to continue to report a valid namespace label size, returns status success (0) and reports this extended status if the persistent memory region of the NVDIMMs are in a state that requires one or more security keys to be applied before the region is accessible.
Size of Namespace Label Area	4	4	Size returned in bytes
Max Namespace Label Data Length	4	8	In bytes, Maximum size of the namespace label data length supported by the platform in <i>Get/Set Namespace Label Data</i> functions



~~3.8.3~~3.9.3 Get Namespace Label Data (Function Index 5)

This command requests the device to return Namespace Label storage area data based on the requested buffer offset and length for the requested device.

Warning: This function has been deprecated in preference to the ACPI 6.2 _LSR (Label Storage Read) NVDIMM Device Interface and is only supported with Arg1 – Revision Id = 1. It is included here for backwards compatibility with existing Arg1 - Revision Id = 1 implementations.

Function Input

The following tables outline the expected input payload for this command.

Table 3-9 Get Namespace Label Data – Input Format

Field	Byte Length	Byte Offset	Description
Offset	4	0	In bytes Indicates the offset in the namespace label data area, to which the namespace label data is to be read from the target NVDIMM
Length	4	4	In bytes

Function Output

The following tables outline the expected output payload for this command.

Table 3-10 Get Namespace Label Data – Output Format

Field	Byte Length	Byte Offset	Description
Status	2	0	Defined above in Table 3-C 03 – Invalid Input Parameters - Offset + Length is > size of Namespace Label Data Area (Max Namespace Label Data Length from GetNamespaceLabelDataSize LSM) - Length is > maximum amount of data the OSPM can transfer in a single request
Extended Status	2	2	Extended Status Field
Namespace Label Data	Varies	4	The size of the output is equal to input's <i>Length</i> if <i>Status</i> is Success; otherwise, the contents of rest of the output buffer are not valid.



3.8.43.9.4 Set Namespace Label Data (Function Index 6)

This command requests the device to update Namespace Label Data area data based on the requested buffer offset and length for the requested device.

Warning: This function has been deprecated in preference to the ACPI 6.2 _LSW (Label Storage Write) NVDIMM Device Interface and is only supported with Arg1 – Revision Id = 1. It is included here for backwards compatibility with existing Arg1 - Revision Id = 1 implementations.

Function Input

The following tables outline the expected input payload for this command.

Table 3-11 Set Namespace Label Data – Input Format

Field	Byte Length	Byte Offset	Description
Offset	4	0	In bytes Indicates the offset in the namespace label data area, to which the <i>Namespace Label Data</i> is to be written to the target NVDIMM
Length	4	4	In bytes
Namespace Label Data	Varies	8	Namespace label data. Size of the namespace label data is as indicated by <i>Length</i> field above.

Function Output

The following tables outline the expected output payload for this command.

Table 3-12 Set Namespace Label Data – Output Format

Field	Byte Length	Byte Offset	Description
Status	2	0	Defined above in Table 3-C 03 – Invalid Input Parameters - Offset + Length is > size of Namespace Label Data Area (Max Namespace Label Data Length from GetNamespaceLabelDataSize LSM) - Length is > maximum amount of data the OSPM can transfer in a single request
Extended Status	2	2	Extended Status Field