



LEGAL NOTICE:

© **Copyright 2007 - 2021 NVM Express, Inc. ALL RIGHTS RESERVED.**

This NVM Express revision 1.4 technical proposal is proprietary to the NVM Express, Inc. (also referred to as "Company") and/or its successors and assigns.

NOTICE TO USERS WHO ARE NVM EXPRESS, INC. MEMBERS: Members of NVM Express, Inc. have the right to use and implement this NVM Express revision 1.4 technical proposal subject, however, to the Member's continued compliance with the Company's Intellectual Property Policy and Bylaws and the Member's Participation Agreement.

NOTICE TO NON-MEMBERS OF NVM EXPRESS, INC.: If you are not a Member of NVM Express, Inc. and you have obtained a copy of this document, you only have a right to review this document or make reference to or cite this document. Any such references or citations to this document must acknowledge NVM Express, Inc. copyright ownership of this document. The proper copyright citation or reference is as follows: "© **2007 - 2021 NVM Express, Inc. ALL RIGHTS RESERVED.**" When making any such citations or references to this document you are not permitted to revise, alter, modify, make any derivatives of, or otherwise amend the referenced portion of this document in any way without the prior express written permission of NVM Express, Inc. Nothing contained in this document shall be deemed as granting you any kind of license to implement or use this document or the specification described therein, or any of its contents, either expressly or impliedly, or to any intellectual property owned or controlled by NVM Express, Inc., including, without limitation, any trademarks of NVM Express, Inc.

LEGAL DISCLAIMER:

THIS DOCUMENT AND THE INFORMATION CONTAINED HEREIN IS PROVIDED ON AN "**AS IS**" BASIS. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, NVM EXPRESS, INC. (ALONG WITH THE CONTRIBUTORS TO THIS DOCUMENT) HEREBY DISCLAIM ALL REPRESENTATIONS, WARRANTIES AND/OR COVENANTS, EITHER EXPRESS OR IMPLIED, STATUTORY OR AT COMMON LAW, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, VALIDITY, AND/OR NONINFRINGEMENT.

All product names, trademarks, registered trademarks, and/or servicemarks may be claimed as the property of their respective owners.

NVM Express Workgroup
c/o VTM Group
3855 SW 153rd Drive
Beaverton, OR 97003 USA
info@nvmexpress.org

NVM Express Technical Proposal for New Feature

Technical Proposal ID	4052c – Endurance Group Management
Change Date	2021-01-14
Builds on Specification	NVM Express specification 1.4 NVM Express Management Interface Specification 1.1
Referenced Ratified Technical Proposals	None
Referenced Technical Proposals in Development	TP 4009 – Domains and Partitioning

Technical Proposal Author(s)

Name	Company
Mark Carlson	TMC
Mark Carlson, Paul Suhler	Kioxia
Paul Suhler, Walt Hubis	Micron
Fred Knight, John Meneghini	NetApp
David Black	DellEMC

This technical proposal defines an interface for interoperable management of Endurance Groups and NVM Sets.

Revision History

Revision Date	Change Description
2018-12-17	Initial draft
2018-12-25	Declined to use the wordy “Endurance Group NVM Set Management” terminology. Capability bits in CRATT. Log page to describe current configuration of each Parallel Unit. Command to assign each Parallel Unit to an Endurance Group and NVM Set, and to delete Endurance Groups and NVM Sets. Identify data structure listing Supported Configuration for Endurance Groups. This is not integrated into the command.
2019-01-14	Fixed wording related to number of endurance groups. Improved integration of host managed configuration and NVM subsystem managed configuration into one command. Added Command Dword 11 definition, to be used for NVM subsystem managed configuration.
2019-01-28	Incorporated changes from 2019-01-15 small team meeting. <ul style="list-style-type: none">Changed terminology from “parallel unit” to “media unit”.

	<ul style="list-style-type: none"> • Miscellaneous other editorial changes. • Added discussion by Peter Onufryk.
2019-02-13	<p>Incorporated changes from 2019-02-05 small team meeting.</p> <ul style="list-style-type: none"> • Reorganized data into a number of descriptors and lists. • Moved proposed section 4.TBD into the model section 8.TBD • Added general descriptions of Media Unit structures. • Moved the operation to select one pre-defined configuration to the first in the Endurance Group Management command's operation list.
2019-02-14	<p>Incorporated changes from 2019-02-14 small team meeting.</p> <ul style="list-style-type: none"> • Aligned figure numbers with NVMe 1.3c. • Merged Memory Type Supported field into Media Unit Status log page, and deleted Media Unit Capability List and associated descriptor. • Removed discussion comparing host and NVM subsystem management of media units. • Added notes of areas that will be affected by TP 4009 (Domains and Partitioning).
2019-02-20	<p>Inserted RFC ballot comments from DellEMC and Micron.</p>
2019-02-21	<p>Inserted RFC ballot comments from NetApp. Inserted comments from 2019-02-21 small team meeting.</p>
2019-02-27	<p>Resolved comments from 2019-02-21 small team meeting:</p> <ul style="list-style-type: none"> • Added Media Unit definition • Modified Media Unit Status log page to not use Media Unit Identifier. • Removed RW column from Media Unit Status Descriptor. • Removed R/W column from Endurance Group Configuration Descriptor. • Added the Endurance Group Management command to privileged actions. <p>Inserted RFC ballot comments from Intel.</p>
2019-03-05	<p>Inserted comments from 2019-02-28 small team meeting. Resolved more comments:</p> <ul style="list-style-type: none"> • Suggesting "Endurance Group Management" for the new command name. • Named two forms of management as Media Unit Endurance Group Management and Capacity Endurance Group Management. Added descriptions to 8.TBD. Added a figure of which structures and operations are required for each form. • Deleted Configure Media Units operation in favor of only selecting a complete configuration (Select Capacity Configuration operation). • Deleted the Media Type List, as the only remaining item was the capacity adjustment factor, and that is in the Endurance Group Configuration Descriptor. • Updated bits in CTRATT.
2019-03-18	<p>Inserted comments from 2019-03-07 small team meeting.</p> <ul style="list-style-type: none"> • Delete "entry" in each list item name. • Convert Fig_Mandatoriness figure to two separate lists. • Describe media units as being in a domain, rather than a NVM subsystem. • Added an operation to undo the Select Capacity Configuration operation. • Defined Media Units as belonging to domains, rather than to NVM subsystems. <p>Added an Endurance Group List (CNS value TBD), as suggested by Fred Knight in a direct e-mail. Added a list of items to be resolved before exiting phase 2.</p>
2019-03-27	<p>Resolved comments from 2019-03-20 Technical WG face-to-face meeting.</p> <ul style="list-style-type: none"> • Simplified Endurance Group List to contain only Endurance Group identifiers. • Removed Delete Media Unit Configuration operation, in favor of allowing Select Capacity Configuration operation specifying a Media Unit Configuration Identifier of 0000h. • Modified TP 4018b description of NVM Set List to state that the list contains identifiers of NVM Sets that exist in the NVM Subsystem. • Added to section 8.TBD lists of actions for creation and deletion of Endurance Groups and NVM Sets, including AENs and changes to data structures. <p>Other changes:</p> <ul style="list-style-type: none"> • In Media Unit Status log page, renamed Last Selected Configuration field to Selected configuration, and deleted Configuration Status field. • Clarified behavior when controller does not support NVM Sets. • Added definition of completion queue entry field to return identifier of created Endurance Group or NVM Set.

	<ul style="list-style-type: none"> Explained distinction between wear information reported for Media Units and that reported for the NVM subsystem.
2019-04-02	<p>Fixed typo reported by David May.</p> <p>Resolved comments from 2019-03-28 subteam meeting:</p> <ul style="list-style-type: none"> Added new Insufficient Capacity status applicable to the Create Endurance Group and the Create NVM Set operations. Added initial version of the reference for Capacity Adjustment Factor.
2019-04-10	<p>Changes from 2019-04-04 subteam meeting:</p> <ul style="list-style-type: none"> Corrected base spec revision to 1.4. Changed TBDs from TP 4018b to numbers in NVMe 1.4. Used the Unallocated NVM Capacity field in the Identify Controller data structure to bound the maximum size of an Endurance Group being created. Revised definition of Capacity Adjustment Factor. Added Endurance Estimate field to Endurance Group Configuration Descriptor. Renamed the Identifier field of the command to Endurance Group Management Identifier. Dropped use of “object” in favor of the more-concrete “Endurance Group or NVM Set”. <p>Dixit Prashant’s suggests from e-mail:</p> <ul style="list-style-type: none"> Require Endurance Group List to be ordered by identifier. Define behavior of delete operations specifying nonexistent Endurance Group or NVM Set. <p>Fred Knight’s suggestions from e-mail:</p> <ul style="list-style-type: none"> Fixed effects on Unallocated NVM Capacity of creating and deleting Endurance Groups and NVM Sets. Changed “Object Identifier” to “Element Identifier”. Refer to Media Units as “accessible by the domain”, rather than “in the domain”. Added his comments. <p>Other:</p> <ul style="list-style-type: none"> Added ordering requirements for descriptors in various structures. Deleted “Fig_” from figure numbers. Added note with items to be resolved in Phase 3.
2019-04-15	<p>Changes from 2019-04-11 subteam meeting:</p> <ul style="list-style-type: none"> The MU Supported Configuration List is converted to a log page, and listed for both I/O controllers and Admin controllers Added descriptor length fields to structure headers. Descriptor granularity and alignment is Dword. Add examples of usage of Capacity Adjustment Factor. Updated 5.15.2.5 (NVM Set List) to include TP 4009 text. Globally replaced “zero” with 0h. Renamed Endurance Group Management Identifier field to Element Identifier. A Create NVM Set operation with the Element Identifier = 0h lets the controller select the Endurance Group in which to create the NVM Set. Data structure requirements were moved from command description into the data structure definitions. Adapted the capacity rounding text from namespace creation to Endurance Group and NVM Set creation. Other minor changes. <p>Changes from e-mail exchange with Fred Knight:</p> <ul style="list-style-type: none"> Added Domain Identifier to Create Endurance Group Operation. Added Destination Element Identifier to completion queue entry to indicate where an Endurance Group or NVM Set was created. Necessary because if host does not specify destination, then controller chooses it and thus must report it.

2019-04-18	<p>Changes from e-mail exchange with Fred Knight:</p> <ul style="list-style-type: none"> Added adding Domain Identifier to change list for previous version. Deleted Destination Element Identifier because the host can look up the information elsewhere. <p>Changes from 2019-04-18 subteam meeting:</p> <ul style="list-style-type: none"> Added statement that MU Supported Config List may change, based on usage and MU status. Changed to latest TP 4009 text for 5.15.2.5 (NVM Set List data structure). Made minor changes to the rounding and capacity adjustment factor example. Deleted the Destination Element Identifier field in the CQE. Listed both log pages in figures 420 and 428. Several minor wording changes were made. Added comments to flag areas that may change depending upon the results of offline discussions.
2019-04-25	<p>Changes from 2019-04-25 subteam meeting:</p> <ul style="list-style-type: none"> Dispositioned Mike Allison's comments. Decided to import entity relationships description (i.e., Domains, MUs, Endurance Groups, NVM Sets) which is being written for TP 4009. Import will happen in Phase 3.
2019-04-29	Implemented the dispositions for Mike's comments that were recorded in the previous revision.
2019-05-02	Markups recording proposed changes from 2019-05-02 Technical WG meeting and subteam meeting.
2019-05-03	<p>Changes from 2019-05-02 Technical WG meeting and subteam meeting:</p> <ul style="list-style-type: none"> Added a summary outside of the spec changes which describes the results of the discussion of how to present additional information for each supported configuration. This will serve as a record upon which to base a future TP. Fixed Capacity Configuration Descriptor to indicate 32-byte Endurance Group Configuration Descriptors. Added informative statement to Endurance Group Management command section that for the Create Endurance Group operation, the controller chooses MUs for the created Endurance Group. Fixed capitalization of first word of list items. Corrected byte ranges in log pages. Fixed minor typos.
2019-05-07	<p>Change from Mark Carlson e-mail of 2019-04-29:</p> <ul style="list-style-type: none"> Added comment to specify in Phase 3 that the relationship between Available Spare in Media Units and in Endurance Groups is outside the scope of the specification.
2019-06-05	<p>Phase 3 initial revision</p> <p>Added architecture section from TP 4009. Expanded to address ballot comments from WDC and Microsemi.</p> <p>Added notes from 2019-05-16 subteam meeting:</p> <ul style="list-style-type: none"> Need to reconsider whether Channel Identifier value indicating not indicated should be 0h or FFFFh. A similar question could apply to MU Identifier. Sentence defining effect on commands of deletion of Endurance Group or NVM Set needs a reference to 8.12, as deletions involve namespace deletion. <p>For reference, added Phase 2 exit ballot comments in the Background section.</p>
2019-06-07	<p>Added changes from e-mail:</p> <ul style="list-style-type: none"> Added ENDGID to data for namespace create operation. Added comment questioning whether an Endurance Group (or NVM subsystem) can create some namespace inside NVM Sets and some in Endurance Group but outside NVM Sets. <p>Changes from 2019-06-06 subteam meeting:</p> <ul style="list-style-type: none"> Added comments describing WDC and Microsemi concerns about current representation of configurations. Resolved some comments per meeting discussion.
2019-06-18	<p>Changes from 2019-06-13 subteam meeting:</p> <ul style="list-style-type: none"> Revised 5.TBD text regarding commands submitted to namespaces whose Endurance Group or NVM Set is deleted. Decided that modification of the definition of NVM Set is for an ECN, not this TP.

	<ul style="list-style-type: none"> • In MU Status Descriptor, moved Channel Identifier and Domain Identifier to second and third positions. They are fixed values and should precede changeable values. Similarly for MU Assignment Descriptor. • Added text specifying that there if NVM Sets are not supported, then every Endurance Group contains a default NVM Set. • Revised references in 7.TBD.3.
2019-07-03	<p>Strawman proposal for alternate organization:</p> <ul style="list-style-type: none"> • Each Capacity Configuration Descriptor contains one or more Endurance Group Configuration Descriptors. • Each Endurance Group Configuration Descriptor contains zero or more NVM Set Identifiers and one or more Channel Configuration Descriptors. • Each Channel Configuration Descriptor contains one or more Media Unit Configuration Descriptors. <p>Included some modifications by Yoni.</p>
2019-07-15	<p>Changes from 2019-07-11 subteam meeting:</p> <ul style="list-style-type: none"> • Modified identifier fields in MU descriptor • Reordered fields in config descriptors to put the list lengths adjacent to the lists. This will facilitate adding more list types in the future.
2019-07-24	<p>Changes from 2019-07-18 subteam meeting and subsequent discussions:</p> <ul style="list-style-type: none"> • Removed 5.4 Create I/O Submission Queue figure, which was a holdover from this TP being based on 1.3. No changes are required to the description of this command in the 1.4 text. • Reverted to "Media Unit Identifier", removing the virtual and physical versions of the identifier. • Defined the effect on Media Unit Status Descriptor fields by clearing the selected configuration. • Added capacity adjustment factor definition allowing non-integer values. Borrowed the formula from UFS, in which a factor of 1 is represented by 0100h in a 16-bit field. • In the Channel Configuration Descriptor, changed the two-byte Media Unit Identifier field to the Media Unit Configuration Descriptor, containing a length field. While all of these descriptors are currently defined as fixed-length, this will permit extension if needed in the future. • Removed the Error Correction Endurance Group Capacity field from the EG config descriptor, as it wasn't clear that it would actually be useful. • Renamed Media Unit Endurance Group Management to Direct Endurance Group Management.
2019-08-21	<p>Changes from 2019-08-01 subteam meeting, and Mike Allison's comments:</p> <ul style="list-style-type: none"> • Renamed Endurance Group Management command to Capacity Management command, with Fixed Capacity Management and Variable Capacity Management. • Added Domain Identifier to the Log Specific Identifier field of the Get Log Page command. • Added Channel 0 Identifier Offset to the MU Status Descriptor, to allow room for future expansion. • Required a minimum of one configuration, one Channel, and one MU per Channel. • Clarified scopes of identifiers. • Removed requirement that Channel Identifier is less than or equal to EGCHANS. That would have not worked for (e.g.) the vertically-organized case where each Endurance Group has a dedicated Channel. • Marked Supported Capacity Configuration List log page as possibly using selection by UUID.
2019-08-23	<p>Changes from Fred Knight's comments and discussion in an after the 2019-08-22 Technical WG meeting:</p> <ul style="list-style-type: none"> • Specified behavior of Get Log Page commands requesting information in an inaccessible domain. This will be to return no descriptors in the log page, only the header. There will be no use of ANA statuses. • Specified behavior of Get Log Page command if the Domain Identifier is not in the Domain List. • Editorial changes.

2019-08-27	<p>Changes from 2019-08-27 joint meeting of the ANA and Endurance Group Management subteams:</p> <ul style="list-style-type: none"> Modified log page descriptions to clarify behavior for single or multiple domain support and domain identifier specified by the command. Added changes to Figure 494 to state that the new log pages are not available in some ANA states. Modified Capacity Configuration Descriptor to allow reporting zero Endurance Groups.
2019-08-28	<p>Changes from discussion with Fred Knight:</p> <ul style="list-style-type: none"> For Endurance Group and NVM Set creation operations that request too large an entity, Error Information log page may return the largest size that can be created. (5.TBD.2, 5.TBD.3) Clarified wording about ENDGID field values when Endurance Groups are not supported. (8.17) If NVM Set creation is supported, require that TEGCAP and UEGCAP (TP 4009) indicate non-zero values, so the host knows the largest NVM Set it can create. (8.TBD.3)
2019-08-29	Revision for member review. No changes from previous revision. Member review advanced comments have been received but are not included here.
2019-08-29	<p>Revision a:</p> <p>Allow reporting zero Channel Config Descriptors in an Endurance Group Config Descriptor.</p>
2019-09-06	<p>New revision for member review.</p> <p>Resolved technical comments reported after Member Review entry:</p> <ul style="list-style-type: none"> Added bits in CTRATT to indicate support for Delete Endurance Group and Delete NVM Set operations. Incremented the Event Type Revision field of the Change Namespace Event. Selection of a non-zero Configuration Identifier is only permitted if Selected Configuration is 0h, i.e., no Configuration Identifier is presently selected. Deleted "this log may change" statement for the Supported Capacity Configuration List. Allow drive to indicate zero Channels in the Media Unit Status log page.
2019-10-02	Proposed resolution of member review comments.
2019-10-16	Resolution of member review comments, per 2019-10-03 task group meeting. Added one Micron comment.
2019-10-28	<p>Changes from 2019-10-17 task group meeting.</p> <p>Open comments highlight changes to be discussed in 2019-10-31 Technical WG meeting.</p> <p>Replaced capacity organization drawings with ones having better labels.</p>
2019-10-31	<p>Changes from 2019-10-31 Technical WG meeting for Member Review re-entry.</p> <ul style="list-style-type: none"> 5.14.1.TBD0: Changed definition of C0IO field in Figure MUsatD. This is a functional change. 5.TBD.1: Deleted requirements to set fields in Media Unit Status Descriptor. Judged not to be a functional change. 7.TBD.2: Deleted second sentence in first paragraph. Not a functional change.
2019-12-05	Renamed for integration
2020-02-08	Ready for ratification. Minor editorial edits.
2020-02-11	Correct section heading TBDs with Log Identifiers and CNS values.
2020-04-15	<p>Changes since TP 4052 (ratified):</p> <ul style="list-style-type: none"> Added the Capacity Management command to Figure 110 and Figure 116 in NVMe-MI 1.1. Added the Capacity Management command to NVMe Base Specification Figure 462: Command Behavior in the Presence of a Reservation. Updated authors.
2020-06-01	<p>Changes from 4052a Member Review (Samsung comments):</p> <ul style="list-style-type: none"> 8.TBD.2 list item d: Changed "set the Variable Capacity Management bit" to "...Fixed...". 8.TBD.3 next-to-last paragraph: Changed "controller can delete" to "host can delete".
2020-06-15	Integrated into the NVM Express Base Specification and NVMe MI Specification.
2020-06-17	Fixed integration comments.

2020-10-29	Added the Capacity Management command to Figure 139: Opcodes for Admin Commands.
2020-10-30	Changes per the 2020-10-29 Technical WG meeting: <ul style="list-style-type: none"> Figures 128: Changed the name of the status code Namespace Identifier Unavailable (16h) to Identifier Unavailable, and added the Capacity Management command. Figures 266: Changed the name of the status code Namespace Identifier Unavailable (16h) to Identifier Unavailable. 5.TBD.2: Added requirement to return Identifier Unavailable if no Endurance Group Identifier is available. 5.TBD.3: Added requirement to return Identifier Unavailable if no NVM Set Identifier is available. Figure CMStat: Added the Identifier Unavailable status code.
2020-11-05	Changes per the 2020-11-05 Technical WG meeting: <ul style="list-style-type: none"> Do not modify the Namespace Identifier Unavailable status code. Added the Identifier Unavailable status code.
2020-11-12	Changes per the 2020-11-12 Technical WG meeting: <ul style="list-style-type: none"> Removed comments and renamed file for Member Review.
2021-01-11	Integrated into the NVMe Base Specification.
2021-01-14	Accepted all changes and removed all comments for ratification.

Incompatible Changes

None

Description for NVMe Base Specification Changes Document

- Defined the Capacity Management command, to select a capacity configuration or to create and delete Endurance Groups and NVM Sets.
- Added the Identifier Unavailable status code.
- Defined the Media Unit Status log page.
- Defined the Supported Capacity Configuration List log page.
- Added a field to the Change Namespace Event Data Format.
- Extended the Get Log Page command to allow specification of a domain identifier
- Defined the NVM Capacity Model, in section 7 (Controller Architecture).
- Defined the Capacity Management model, in section 8 (Features).
- Defined the Endurance Group List, an Identify data structure.
- References:
 - NVMe Base Specification 1.4 sections 1, 4, 5, 7, and 8.
 - Technical Proposal 4052a

Description for NVMe Management Interface Specification Changes Document

- Added the Capacity Management command to the List of NVMe Admin Commands Supported Using the Out-of-Band Mechanism (Figure 110).
- Added the Capacity Management command to Command Messages Allowed During Sanitize Operation (Figure 116).
- References:
 - NVMe Management Interface Specification 1.1 section 6.
 - Technical Proposal 4052a

Markup Conventions:

Black: Unchanged (however, hot links are removed)
~~Red Strikethrough:~~ Deleted
Blue: New

Blue Highlighted: TBD values, anchors, and links to be inserted in new text.

<Green Bracketed>: Notes to editor

< Note to Editor: The MEGCAP,TEGCAP, and UEGCAP fields referred to in this proposal were defined in TP 4009. >

1 Introduction

1.6 Definitions

...

1.6.TBD0 Channel

A Channel represents a communication path between the controller and one or more Media Units in an NVM subsystem.

1.6.TBD1 Media Unit

A Media Unit represents a component of the underlying media in an NVM subsystem. Endurance Groups are composed of Media Units.

...

4 Data Structures

...

4.6 NVM Sets

...

4.6.1 Status Field Definition

...

4.6.1.2 Status Code (SC)

...

4.6.1.2.2 Command Specific Status Definition

...

Figure 128: Status Code – Command Specific Status Values

Value	Description	Commands Affected
...		
25h	ANA Attach Failed	Namespace Attachment
26h	Insufficient Capacity	Capacity Management
2Dh	Identifier Unavailable	Capacity Management
...		

...

4.9 NVM Sets

...

There is a subset of Admin commands that are NVM Set aware as described in **Figure 134**.

Figure 134: NVM Set Aware Admin Commands

Admin Command	Details
Identify	<ul style="list-style-type: none"> The Identify Namespace data structure includes the associated NVM Set Identifier. The NVM Set List data structure includes attributes for each NVM Set.
Capacity Management	<ul style="list-style-type: none"> The Create NVM Set action returns the NVM Set Identifier of the NVM Set that is created. The Delete NVM Set action includes the NVM Set Identifier of the NVM Set that is to be deleted.
Namespace Management	<ul style="list-style-type: none"> The create action includes the NVM Set Identifier as a host specified field.
Set Features	<ul style="list-style-type: none"> The Read Recovery Level Feature specifies the associated NVM Set Identifier.

...

If NVM Sets are supported, then all controllers in the NVM subsystem shall:

- Indicate support for NVM Sets in the Controller Attributes field in the Identify Controller data structure;
- Support the NVM Set Identifier in all commands that use the NVM Set Identifier;
- Support the NVM Set List for the Identify command;
- Indicate the NVM Set Identifier with which the namespace is associated in the Identify Namespace data structure;
- Support Endurance Groups; and
- For each NVM Set, indicate the associated Endurance Group as an attribute.

If support for NVM Sets is not reported (i.e., the NVM Sets bit is cleared to '0' in the CTRATT field; refer to [Figure 247](#)), then the NVM Set Identifier field shall be cleared to 0h in all commands and data structures that support an NVM Set Identifier field.

4.10 Namespace List

...

5 Admin Command Set

...

Figure 1: Opcodes for Admin Commands

Opcode by Field			Combined Opcode ¹	Namespace Identifier Used ²	Command
(07)	(06:02)	(01:00)			
Generic Command	Function	Data Transfer ³			
0b	010 00b	00b	20h	No	Capacity Management

...

5.14 Get Log Page command

The Get Log Page command returns a data buffer containing the log page requested. [The Get Log Page command may be impacted by the ANA state \(refer to section 8.20.4\).](#)

The Get Log Page command uses the Data Pointer, Command Dword 10, Command Dword 11, Command Dword 12, Command Dword 13, and Command Dword 14 fields. All other command specific fields are reserved.

There are mandatory and optional Log Identifiers defined in [Figure 193](#) and [Figure 192](#). If a Get Log Page command is processed that specifies a Log Identifier that is not supported, then the controller should abort the command with status Invalid Field in Command.

The controller indicates support for the Log Page Offset and extended Number of Dwords (32 bits rather than 12 bits) in the Log Page Attributes field of the Identify Controller data structure. If extended data is not supported, then bits 27:16 of the Number of Dwords Lower field specify the Number of Dwords to transfer.

...

Figure 287: Get Log Page – Command Dword 11

Bits	Description										
31:16	Log Specific Identifier: This field specifies an identifier that is required for a particular log page. The log pages that require a log specific identifier are indicated in the table below.										
	<table><tr><th>Log Page</th><th>Definition</th></tr><tr><td>Endurance Group Information</td><td>Endurance Group Identifier (refer to section 8.17)</td></tr><tr><td>Predictable Latency Per NVM Set</td><td>NVM Set Identifier (refer to section 4.9)</td></tr><tr><td>Media Unit Status</td><td>Domain Identifier (refer to section TP4009_7.NEW)¹</td></tr><tr><td>Supported Capacity Configuration List</td><td>Domain Identifier (refer to section TP4009_7.NEW)¹</td></tr></table>	Log Page	Definition	Endurance Group Information	Endurance Group Identifier (refer to section 8.17)	Predictable Latency Per NVM Set	NVM Set Identifier (refer to section 4.9)	Media Unit Status	Domain Identifier (refer to section TP4009_7.NEW) ¹	Supported Capacity Configuration List	Domain Identifier (refer to section TP4009_7.NEW) ¹
	Log Page	Definition									
	Endurance Group Information	Endurance Group Identifier (refer to section 8.17)									
	Predictable Latency Per NVM Set	NVM Set Identifier (refer to section 4.9)									
	Media Unit Status	Domain Identifier (refer to section TP4009_7.NEW) ¹									
Supported Capacity Configuration List	Domain Identifier (refer to section TP4009_7.NEW) ¹										
15:00	Number of Dwords (NUMDU): This field specifies the upper 16 bits of the number of dwords to return.										
NOTES:											
1. If the NVM subsystem does not support multiple domains, then this field is reserved. If this field specifies a non-zero Domain Identifier that is not reported in the Domain List (refer to section TP4009_5.15.2.99), then the controller shall abort the command with Invalid Field in Command.											

...

5.14.1 Log Specific Information

...

< **Note to Editor:** The following figure is modified by TP 4009 to show that the Scope for the three pages below is “Domain / NVM subsystem”. >

Figure 193: Get Log Page – Log Page Identifiers

Log Identifier	Scope	Log Page Name	Reference Section
...			
0Fh	NVM subsystem	Endurance Group Event Aggregate	5.14.1.15
10h	NVM subsystem	Media Unit Status	5.14.1.TBD0
11h	NVM subsystem ⁵	Supported Capacity Configuration List	5.14.1.TBD1
10h 12h to 6Fh	Reserved		
...			
KEY: Namespace = The log page contains information about a specific namespace. Controller = The log page contains information about the controller that is processing the command. NVM subsystem = The log page contains information about the NVM subsystem.			
NOTES: 1. For namespace identifiers of 0h or FFFFFFFFh. 2. For namespace identifiers other than 0h or FFFFFFFFh. 3. Bit 0 is cleared to ‘0’ in the DSTO field in the Identify Controller data structure (refer to Figure 247). 4. Bit 0 is set to ‘1’ in the DSTO field in the Identify Controller data structure. 5. Selection of a UUID may be supported. Refer to section 8.24 .			

...

5.14.1.11 Predictable Latency Event Aggregate Log Page (Log Identifier 0Bh)

This log page indicates if a Predictable Latency Event (refer to section 8.18) has occurred for a particular NVM Set. If a Predictable Latency Event has occurred, the details of the particular event are included in the Predictable Latency Per NVM Set log page for that NVM Set. An asynchronous event is generated when an entry for an NVM Set is newly added to this log page.

This log page shall not contain an entry (i.e., an NVM Set Identifier) that is cleared to 0h.

If there is an enabled Predictable Latency Event pending for an NVM Set, then the Predictable Latency Event Aggregate log page includes an entry for that NVM Set. The log page is an ordered list by NVM Set Identifier. For example, if Predictable Latency Events are pending for NVM Set 27, 13, and 17, then the log page shall have entries in numerical order of 13, 17, and 27. A particular NVM Set is removed from this log page after the Get Log Page is completed successfully with the Retain Asynchronous Event bit cleared to '0' for the Predictable Latency Per NVM Set log page for that NVM Set.

...

5.14.1.13 Persistent Event Log (Log Identifier 0Dh)

...

5.14.1.13.1 Persistent Event Log Events

...

5.14.1.13.1.6 Change Namespace Event (Event Type 06h)

...

The Changed Namespace Event shall set the Persistent Event Log Event Format Header:

- Event Type field to 06h; and
- Event Type Revision Field to ~~04h~~ 02h.

Figure 222: Change Namespace Event Data Format (Event Type 06h)

Bytes	Value
...	
41:40	NVM Set Identifier (NVMSETID): For a create operation, contains the NVMSETID value from the Identify Namespace data structure in the Namespace Management command (refer to Figure 265). For a delete operation that specifies a single namespace this field contains the value from the NVMSETID field of the Identify Namespace data (refer to Figure 245) for the namespace being deleted. For a delete operation that specifies all namespaces this field is reserved.
43:42	Reserved Endurance Group Identifier (ENDGID): For a create operation, contains the ENDGID value from the Identify Namespace data structure in the Namespace Management command (refer to section 5.20). For a delete operation that specifies a single namespace, this field contains the value from the ENDGID field of the Identify Namespace data structure (refer to Figure 245) for the namespace being deleted. For a delete operation that specifies all namespaces this field is reserved.
47:44	Namespace ID (NSID): For a create operation, contains the NSID for the namespace that was created. For a delete operation, contains the NSID from CDW1.NSID (i.e., the NSID for the namespace being deleted or FFFFFFFFh for a delete operation specifying all namespaces).

...

5.14.1.TBD0 Media Unit Status (Log Identifier 10h)

This log page is used to describe the configuration and wear of Media Units (refer to section 8.TBD). The log page contains one Media Unit Status Descriptor for each Media Unit accessible by the specified domain. Each Media Unit Status Descriptor (refer to Figure MUsatD) indicates the configuration of the Media Unit (e.g., to which Endurance Group the Media Unit is assigned, to which NVM Set the Media Unit is assigned, to

which Channels the Media Unit is attached) and indications of wear (e.g., the Available Spare field and the Percentage Used field). The indications of wear change as the Media Unit is written and read.

If the NVM subsystem supports multiple domains, then the controller reports the Media Unit Status log page for the domain specified in the Log Specific Identifier field (refer to [Figure 187](#)), if accessible. If the information is not accessible, then the log page is not available (refer to section 8.20.4). If the Log Specific Identifier field is cleared to 0h, then the specified domain is the domain containing the controller that is processing the command.

Media Unit Identifier values (refer to [Figure MUMStatD](#)) begin with 0h and increase sequentially. If the NVM subsystem supports multiple domains, then the Media Unit Identifier values are unique within the specified domain. If the NVM subsystem does not support multiple domains, then the Media Unit Identifier values are unique within the NVM subsystem.

Media Unit Status Descriptors are listed in ascending order by Media Unit Identifier.

Requirements for supporting the Media Unit Status log page are defined in [section 8.TBD](#).

Figure MUMLog: Get Log Page – Media Unit Status

Bytes	Description
01:00	Number of Media Unit Status Descriptors (NMU): This field indicates the number of Media Unit Status Descriptors in the log page. If this field is cleared to 0h, then no Media Unit Status Descriptors are reported.
03:02	Number of Channels (CCHANS): This field indicates the number of Channels accessible by the controller. A value of 0h indicates that the number of Channels accessible by the controller is not reported.
05:04	Selected Configuration: This field indicates the Configuration Identifier selected by the most recent successful completion of the Capacity Management command Select Capacity Configuration operation. If a Select Capacity Configuration operation has not been completed, this field may indicate a non-zero value (i.e., a configuration was selected by default). If a Configuration Identifier is not selected, then this field shall be cleared to 0h.
15:6	Reserved
NOTE 1	Media Unit Status Descriptor 0: This field contains the first Media Unit Status Descriptor (refer to Figure MUMStatD), if any.
NOTE 1	Media Unit Status Descriptor 1: This field contains the second Media Unit Status Descriptor, if any.
...	
NOTE 1	Media Unit Status Descriptor NMU-1: This field contains the last Media Unit Status Descriptor, if any.
NOTES: 1. Media Unit Status Descriptors may be different lengths.	

The Media Unit Status Descriptor is defined in [Figure MUMStatD](#).

If the Selected Configuration field is cleared to 0h, then the following fields in each Media Unit Status Descriptor shall be cleared to 0h:

- a) Endurance Group Identifier (ENDGID);
- b) NVM Set Identifier (NVMSETID);
- c) Capacity Adjustment Factor; and
- d) Number of Channels (MUCS).

Channel Identifier values begin with 0h and increase sequentially. If the NVM subsystem supports multiple domains, then Channel Identifier values are unique within the specified domain. If the NVM subsystem does not support multiple domains, then Channel Identifier values are unique within the NVM subsystem.

In the Media Unit Status Descriptor, Channel Identifiers are listed in ascending order by value, and each Channel Identifier shall appear only once.

Figure MUsatD: Media Unit Status Descriptor

Bytes	Description
01:00	Media Unit Identifier: This field indicates the identifier of the Media Unit.
03:02	Domain Identifier: This field indicates the identifier of the domain that contains this Media Unit. Refer to section TP4009_7.NEW2.3 . A value of 0h indicates that a Domain Identifier is not reported. If multiple domains are not supported, then this field shall be cleared to 0h.
05:04	Endurance Group Identifier (ENDGID): This field indicates the identifier of the Endurance Group that contains this Media Unit. Refer to section 8.17 . The value shall be less than or equal to the value of the Endurance Group Identifier Maximum field (refer to Figure 247). A value of 0h indicates that this Media Unit is not part of an Endurance Group.
07:06	NVM Set Identifier (NVMSETID): This field indicates the identifier of the NVM Set that contains this Media Unit. Refer to section 4.9 . This field shall indicate a value less than or equal to the value of the NVM Set Identifier Maximum field (refer to Figure 247). If the controller does not support NVM Sets, then this field shall be cleared to 0h.
09:08	Capacity Adjustment Factor: This field indicates the capacity adjustment factor (refer to section 8.TBD.1) for this Endurance Group. A value of FFFFh indicates that value and all higher values. A value of 0h indicates that the Capacity Adjustment Factor is not reported. All Media Unit Status Descriptors which indicate the same Endurance Group Identifier shall indicate the same value in their Capacity Adjustment Factor fields.
10	Available Spare: Contains a normalized percentage (0 to 100%) of the remaining spare capacity available for the Media Unit. The relationship between this value and the value in the Available Spare field in the Endurance Group Information log page (refer to section 5.14.1.9) is outside the scope of this specification.
11	Percentage Used: Contains a vendor specific estimate of the percentage of life used for the Media Unit based on the actual usage and the manufacturer's prediction of NVM life. A value of 100 indicates that the estimated endurance of the NVM in the Media Unit has been consumed, but may not indicate an NVM failure. The value is allowed to exceed 100. Percentages greater than 254 shall be represented as 255. This value shall be updated once per power-on hour when the controller is not in a sleep state. Refer to the JEDEC JESD218A standard for SSD device life and endurance measurement techniques. The relationship between this value and the value in the Percentage Used field in the Endurance Group Information log page is outside the scope of this specification.
12	Number of Channels (MUCS): This field indicates the number of Channels attached to this Media Unit. If this field is cleared to 0h, then no Channel Identifiers are reported for this Media Unit.
13	Channel Identifiers Offset (CIO): This field indicates the offset of the Channel 0 Identifier field from the beginning of the Media Unit Status Descriptor. This field shall be a non-zero value and a multiple of 16.
CIO-1:14	Reserved
Channel Identifiers	
CIO+1 : CIO	Channel 0 Identifier: This field contains the identifier for the first Channel attached to this Media Unit, if any.
CIO+3 : CIO+2	Channel 1 Identifier: This field contains the identifier for the second Channel attached to this Media Unit, if reported, if any.
	...
(MUCS*2)+CIO-1 : (MUCS*2)+CIO-2	Channel MUCS-1 Identifier: This field contains the identifier for the last Channel attached to this Media Unit, if any.

...

5.14.1.TBD1 Supported Capacity Configuration List (Log Identifier 11h)

This log page is used to provide a list of Supported Capacity Configuration Descriptors (refer to [Figure SCCLog](#)). Each entry in the list defines a different configuration of Endurance Groups supported by the specified domain.

If the NVM subsystem supports multiple domains, then the controller reports the Supported Capacity Configuration List log page for the domain specified in the Log Specific Identifier field (refer to [Figure 187](#)), if accessible. If the information is not accessible, then the log page is not available (refer to [section 8.20.4](#)). If the Log Specific Identifier field is cleared to 0h, then the specified domain is the domain containing the controller that is processing the command.

If the NVM subsystem supports multiple domains, then Capacity Configuration Identifier values are unique within the specified domain. If the NVM subsystem does not support multiple domains, then Capacity Configuration Identifier values are unique within the NVM subsystem.

Capacity Configuration Descriptors are listed in ascending order by Capacity Configuration Identifier, and each Capacity Configuration Identifier shall appear only once.

Figure SCCLog: Get Log Page – Supported Capacity Configuration List

Bytes	Description
0	Number of Supported Capacity Configurations (SCCN): This field indicates the number of Capacity Configuration Descriptors in the list. If this field is cleared to 0h, then no Capacity Configuration Descriptors are reported.
15:1	Reserved
NOTE 1	Capacity Configuration 0 Descriptor: This field indicates the first Capacity Configuration Descriptor (refer to Figure CapCfD) in the list, if any.
NOTE 1	Capacity Configuration 1 Descriptor: This field indicates the second Capacity Configuration Descriptor in the list, if any.
...	
NOTE 1	Capacity Configuration SCCN-1 Descriptor: This field indicates the last Capacity Configuration Descriptor in the list, if any.
NOTES: 1. Capacity Configuration Descriptors may be different lengths.	

The Capacity Configuration Descriptor (refer to [Figure CapCfD](#)) indicates the details of a particular configuration of Endurance Groups and contains one Endurance Group Configuration Descriptor for each Endurance Group accessible by the controller processing the command.

In the Capacity Configuration Descriptor, Endurance Group Configuration Descriptors are listed in ascending order by Endurance Group Identifier, and each Endurance Group Identifier shall appear only once.

Figure CapCfD: Capacity Configuration Descriptor

Bytes	Description
1:0	Capacity Configuration Identifier: This field indicates the identifier for this Capacity Configuration.
3:2	Domain Identifier: This field indicates the identifier of the domain (refer to section TP4009_7.NEW2.3) containing the Endurance Group configurations described by this Capacity Configuration Descriptor. A value of 0h indicates that a Domain Identifier is not reported. If multiple domains are not supported, then this field shall be cleared to 0h.
5:4	Number of Endurance Group Configuration Descriptors (EGCN): This field indicates the number of Endurance Group Configuration Descriptors in the list. If this field is cleared to 0h, then no Endurance Group Configuration Descriptors are reported.
31:6	Reserved
NOTE 1	Endurance Group Configuration 0 Descriptor: This field indicates the first Endurance Group Configuration Descriptor (refer to Figure EGCfD) in the list, if any.
NOTE 1	Endurance Group Configuration 1 Descriptor: This field indicates the second Endurance Group Configuration Descriptor in the list, if any.
...	
NOTE 1	Endurance Group Configuration EGCN-1 Descriptor: This field indicates the last Endurance Group Configuration Descriptor in the list, if any.
NOTES: 1. Endurance Group Configuration Descriptors may be different lengths.	

The Endurance Group Configuration Descriptor is defined in [Figure EGCfgD](#).

In the Endurance Group Configuration Descriptor, NVM Set Identifiers are listed in ascending order by value, and each NVM Set Identifier shall appear only once.

In the Endurance Group Configuration Descriptor, Channel Configuration Descriptors are listed in ascending order by Channel Identifier value, and each Channel Identifier shall appear only once.

Figure EGCfgD: Endurance Group Configuration Descriptor

Bytes	Description
1:0	Endurance Group Identifier (ENDGID): This field indicates the identifier of the Endurance Group (refer to section 8.17) described by this Endurance Group Configuration Descriptor. This field shall indicate a value greater than or equal to 1h and less than or equal to the value of the Endurance Group Identifier Maximum field (refer to Figure 247).
3:2	Capacity Adjustment Factor: This field indicates the capacity adjustment factor (refer to section 8.TBD.1) for this Endurance Group. A value of FFFFh indicates that value and all higher values. A value of 0h indicates that the Capacity Adjustment Factor is not reported.
15:4	Reserved
31:16	Total Endurance Group Capacity (TEGCAP): This field indicates the total NVM capacity in this Endurance Group. The value is in bytes. If this field is cleared to 0h, the NVM subsystem does not report the total NVM capacity in this Endurance Group.
47:32	Spare Endurance Group Capacity (SEGCAP): This field indicates the spare NVM capacity in this Endurance Group. The value is in bytes. If this field is cleared to 0h, the NVM subsystem does not report the unallocated NVM capacity in this Endurance Group.
63:48	Endurance Estimate: This field is an estimate of the total number of data bytes that may be written to the Endurance Group over the lifetime of the Endurance Group assuming a write amplification of 1 (i.e., no increase in the number of write operations performed by the device beyond the number of write operations requested by a host). This value is reported in billions (i.e., a value of 1h corresponds to 1,000,000,000 bytes written) and is rounded up (e.g., a value of 1h indicates the number of bytes written is from 1 to 1,000,000,000, 2h indicates the number of bytes written is from 1,000,000,001 to 2,000,000,000). A value of FFFFFFFF_FFFFFFFF_FFFFFFFF_FFFFFFFFh means that value and all higher values. A value of 0h indicates that the Endurance Estimate is not reported. The relationship between this value and the value in the Endurance Estimate field in the Endurance Group Information log page (refer to section 5.14.1.9) is outside the scope of this specification.
79:64	Reserved
NVM Set Identifiers	
81:80	Number of NVM Sets (EGSETS): This field indicates the number of NVM Set Identifiers in this Endurance Group Configuration Descriptor. A value of 0h indicates that no NVM Set Identifiers are reported for this Endurance Group.
83:82	NVM Set 0 Identifier: This field indicates the identifier of the first NVM Set assigned to this Endurance Group, if reported. Refer to section 4.9 . This field shall indicate a value greater than or equal to 0h and less than or equal to the value of the NVM Set Identifier Maximum field in the Identify Controller data structure.
85:84	NVM Set 1 Identifier: This field indicates the identifier of the second NVM Set assigned to this Endurance Group, if reported. This field shall indicate a value greater than or equal to 0h and less than or equal to the value of the NVM Set Identifier Maximum field in the Identify Controller data structure.
	...

Bytes	Description
(EGSETS*2)+81 : (EGSETS*2)+80	NVM Set EGSETS-1 Identifier: This field indicates the identifier of the last NVM Set assigned to this Endurance Group, if reported. This field shall indicate a value greater than or equal to 0h and less than or equal to the value of the NVM Set Identifier Maximum field in the Identify Controller data structure.
Channel Configuration Descriptors	
(EGSETS*2)+83 : (EGSETS*2)+82	Number of Channels (EGCHANS): This field indicates the number of Channel Configuration Descriptors in this Endurance Group Configuration Descriptor. If this field is cleared to 0h, then no Channel Configuration Descriptors are reported for this Endurance Group.
NOTE 1	Channel 0 Configuration Descriptor: This field contains the Channel Configuration Descriptor (refer to Figure ChCfgDescr) for the first Channel in this Endurance Group, if any.
NOTE 1	Channel 1 Configuration Descriptor: This field contains the Channel Configuration Descriptor for the second Channel in this Endurance Group, if any.
	...
NOTE 1	Channel EGCHANS-1 Configuration Descriptor: This field contains the Channel Configuration Descriptor for the last Channel in this Endurance Group, if any.
NOTES: 1. Channel Configuration Descriptors may be different lengths.	

The Channel Configuration Descriptor (refer to [Figure ChCfgDescr](#)) lists the Media Units attached to a Channel. Media Unit Configuration Descriptors (refer to [Figure MUCfgDescr](#)) are listed in ascending order by Media Unit Identifier, and each Media Unit Identifier shall appear only once.

Figure ChCfgDescr: Channel Configuration Descriptor

Byte	Description
1:0	Channel Identifier: This field indicates the identifier of this Channel. A value of FFFFh indicates that the Channel Identifier is not specified.
3:2	Number of Channel Media Units (CHMUS): This field indicates the number of Media Units that are attached to this Channel. If this field is cleared to 0h, then no Media Unit Configuration Descriptors are reported for this Channel.
NOTE 1	Media Unit 0 Configuration Descriptor: This field contains the Media Unit Configuration Descriptor (refer to Figure MUCfgDescr) for the first Media Unit attached to this Channel, if any.
NOTE 1	Media Unit 1 Configuration Descriptor: This field contains the Media Unit Configuration Descriptor for the second Media Unit attached to this Channel, if reported.
	...
NOTE 1	Media Unit CHMUS-1 Configuration Descriptor: This field contains the Media Unit Configuration Descriptor for the last Media Unit attached to this Channel, if any.
NOTES: 1. Media Unit Configuration Descriptors may be different lengths.	

The Media Unit Configuration Descriptor is defined in [Figure MUCfgDescr](#).

Figure MUCfgDescr: Media Unit Configuration Descriptor

Byte	Description
1:0	Media Unit Identifier: This field indicates the identifier of this Media Unit.
5:2	Reserved
7:6	Media Unit Descriptor Length (MUDL): This field contains the length in bytes of the descriptor information that follows. The total length of the Media Unit Configuration Descriptor in bytes is the value in this field plus 8. In controllers compliant to this revision of this specification, this field shall be cleared to 0h.

...

5.15 Identify Command

5.15.1 Identify Command Overview

...

Figure 242: Identify – Command Dword 11

Bits	Description						
31:16	Reserved						
15:00	<p>NVM Set Identifier (NVMSETID): This field specifies the identifier of the NVM Set. This field is used for Identify operations with a CNS value of 04h. This field should be cleared to 0h for Identify operations with other CNS values.</p> <p>CNS Specific Identifier: This field specifies an identifier that is required for a particular CNS value. The CNS values that require a CNS specific identifier are indicated in the table below.</p> <table><tr><th>CNS Value</th><th>Definition</th></tr><tr><td>NVM Set List (04h)</td><td>NVM Set Identifier (NVMSETID) (refer to section 4.9)</td></tr><tr><td>Endurance Group List (19h)</td><td>Endurance Group Identifier (ENDGID) (refer to section 8.17)</td></tr></table>	CNS Value	Definition	NVM Set List (04h)	NVM Set Identifier (NVMSETID) (refer to section 4.9)	Endurance Group List (19h)	Endurance Group Identifier (ENDGID) (refer to section 8.17)
CNS Value	Definition						
NVM Set List (04h)	NVM Set Identifier (NVMSETID) (refer to section 4.9)						
Endurance Group List (19h)	Endurance Group Identifier (ENDGID) (refer to section 8.17)						

...

Figure 244: Identify – CNS Values

CNS Value	O/M ¹	Definition	NSID ²	CNTID ³	Reference Section
Active Namespace Management					
...					
Controller and Namespace Management					
...					
17h	O	A UUID List (refer to Figure 257) is returned to the host.	N	N	5.15.2.15
19h	O ⁶	Endurance Group List.	N	N	5.15.2.TBD1
18h 1Ah to 1Fh		Reserved			
Future Definition					
20h to FFh		Reserved			
NOTES:					
1. O/M definition: O = Optional, M = Mandatory.					
2. The CDW1.NSID field is used: Y = Yes, N = No.					
3. The CDW10.CNTID field is used: Y = Yes, N = No.					
4. Mandatory for controllers that support the Namespace Management capability (refer to section 8.12).					
5. Mandatory for controllers that support Virtualization Enhancements (refer to section 8.5).					
6. Mandatory for controllers that support Variable Capacity Management (refer to section 8.TBD.3).					

5.15.2 Identify Data Structures

...

5.15.2.2 Identify Controller data structure (CNS 01h)

...

Figure 247: Identify – Identify Controller Data Structure

Bytes	O/M ¹	Description
Controller Capabilities and Features		
...		
99:96	M	<p>Controller Attributes (CTRATT): This field indicates attributes of the controller.</p> <p>Bits 31:40 15 are reserved.</p> <p>Bit 14 (Delete NVM Set): If set to '1', then the controller supports the Delete NVM Set operation (refer to section 8.TBD.3). If cleared to '0', then the controller does not support the Delete NVM Set operation.</p> <p>Bit 13 (Delete Endurance Group): If set to '1', then the controller supports the Delete Endurance Group operation (refer to section 8.TBD.3). If cleared to '0', then the controller does not support the Delete Endurance Group operation.</p> <p>Bit 12 (Variable Capacity Management): If set to '1', then the controller supports Variable Capacity Management (refer to section 8.TBD.3). If cleared to '0', then the controller does not support Variable Capacity Management.</p> <p>Bit 11 (Fixed Capacity Management): If set to '1', then the controller supports Fixed Capacity Management (refer to section 8.TBD.2). If cleared to '0', then the controller does not support Fixed Capacity Management.</p> <p>Bit 9 (UUID List): If set to '1', then the controller supports reporting of a UUID List (refer to Figure 257). If cleared to '0', then the controller does not support reporting of a UUID List (refer to section 8.24).</p> <p>...</p> <p>Bit 0 if set to '1', then the controller supports a 128-bit Host Identifier. Bit 0 if cleared to '0', then the controller does not support a 128-bit Host Identifier.</p>
...		

...

< **Note to Editor:** The change in red in the following paragraph was proposed in TP 4009 and is included here for reference. >

5.15.2.5 NVM Set List (CNS 04h)

[Figure 250](#) defines an NVM Set List. The data structure is an ordered list by NVM Set Identifier, starting with the first NVM Set Identifier supported by the NVM subsystem that is equal to or greater than the NVM Set Identifier indicated in CDW11.NVMSETID and are accessible by the controller processing the command. The NVM Set List describes the attributes for each NVM Set in the list based on the NVM Set Attributes Entry in [Figure 250](#).

The NVM Set List shall not contain an entry cleared to 0h.

...

[5.15.2.TBD1](#) Endurance Group List (CNS 19h)

An Endurance Group List (refer to [Figure EGIL](#)) of up to 2,047 Endurance Group Identifiers in increasing order is returned containing an Endurance Group Identifier greater than or equal to the value specified in the Endurance Group Identifier (CDW11.ENDGID) field. The list contains Endurance Group Identifiers of Endurance Groups that are accessible by the controller processing the command. If the value specified in the Endurance Group Identifier is greater than ENDGIDMAX, then the controller shall complete the command with a status of Successful Completion and return an Endurance Group List containing no Endurance Group Identifiers.

Figure EGIL: Endurance Group List

Bytes	Description
01:00	Number of Identifiers (N): This field contains the number of Endurance Group Identifiers in the list. There may be up to 2,047 identifiers in the list. If this field is cleared to 0h, then no Endurance Group Identifiers are in the list.
03:02	Identifier 0: This field contains the first Endurance Group Identifier in the list, if any.
05:04	Identifier 1: This field contains the second Endurance Group Identifier in the list, if any.
...	...
(N*2+1):(N*2)	Identifier N-1: This field contains the last Endurance Group Identifier in the list.

...

5.15.3 Namespace Management command

...

Figure 4: Namespace Management – Host Software Specified Fields

Bytes	Description	Host Specified
07:00	Namespace Size (NSZE)	Yes
15:08	Namespace Capacity (NCAP)	Yes
25:16	Reserved	
26	Formatted LBA Size (FLBAS)	Yes
28:27	Reserved	
29	End-to-end Data Protection Type Settings (DPS)	Yes
30	Namespace Multi-path I/O and Namespace Sharing Capabilities (NMIC)	Yes
91:31	Reserved	
95:92	ANA Group Identifier (ANAGRPID) ¹	Yes
99:96	Reserved	
101:100	NVM Set Identifier (NVMSETID)	Yes
103:102	Endurance Group Identifier (ENDGID)	Yes
383:102104	Reserved	
Notes:		
1. A value of 0h specifies that the controller determines the value to use (refer to section 8.12).		

< **Note to Editor:** The ENDGID field must be in bytes 103:102 to match Figure 245. >

...

5.TBD Capacity Management command

Host software uses the Capacity Management command to configure Endurance Groups and NVM Sets in an NVM Subsystem by either selecting one of a set of supported configurations (i.e., Fixed Capacity Management; refer to [section 8.TBD.2](#)) or by specifying the capacity of the Endurance Group or NVM Set to be created (i.e., Variable Capacity Management; refer to [section 8.TBD.3](#)). The Capacity Management command specifies the operations defined in [Figure CMC10](#).

The Capacity Management command uses the Command Dword 10, Command Dword 11, and Command Dword 12 fields. All other command specific fields are reserved.

For requirements to support the operations in [Figure CMC10](#), refer to [section 8.TBD](#).

Figure CMC10: Capacity Management – Command Dword 10

Bits	Description
31:16	Element Identifier: This field contains a value specific to the value of the Operation field.
15:04	Reserved

Figure CMC10: Capacity Management – Command Dword 10

Bits	Description		
03:00	Operation: Specifies the operation to be performed by the controller:		
	Value	Description	Element Identifier Field
	0h	Select Capacity Configuration: Endurance Groups and NVM Sets are configured as indicated by the Capacity Configuration Descriptor specified by this command.	Capacity Configuration Identifier (refer to Figure CapCfgD).
	1h	Create Endurance Group: An Endurance Group is created (refer to section 8.TBD.3) with the capacity specified by the Capacity Lower field (refer to Figure CMC11) and the Capacity Upper field (refer to Figure CMC12).	Domain Identifier (refer to section TP4009_7.NEW2.3) of the domain in which the Endurance Group is to be created. A value of 0h specifies that the controller selects the domain in which the Endurance Group is created.
	2h	Delete Endurance Group: The Endurance Group specified by this command is deleted (refer to section 8.TBD.3). All namespaces and NVM Sets contained by the Endurance Group are deleted.	Endurance Group Identifier of the Endurance Group to be deleted.
	3h	Create NVM Set: An NVM Set is created (refer to section 8.TBD.3) in the specified Endurance Group with the capacity specified by the Capacity Lower field and the Capacity Upper field. If the controller does not support NVM Sets, then this operation is not supported.	Endurance Group Identifier of the Endurance Group in which the NVM Set is to be created. A value of 0h specifies that the controller selects the Endurance Group in which the NVM Set is created.
	4h	Delete NVM Set: The NVM Set specified by this command is deleted (refer to section 8.TBD.3). All namespaces in the NVM Set are deleted. If the controller does not support NVM Sets, then this operation is not supported.	NVM Set identifier of the NVM Set to be deleted.
	5h to Fh	Reserved	
NOTES: 1. If the Element Identifier field specifies a non-zero value which does not correspond to an existing capacity entity, then the controller shall abort the command with status Invalid Field in Command.			

Figure CMC11: Capacity Management – Command Dword 11

Bits	Description
31:00	Capacity Lower: This field specifies the lower 32 bits of the capacity in bytes of the Endurance Group or NVM Set to be created.

Figure CMC12: Capacity Management – Command Dword 12

Bits	Description
31:00	Capacity Upper: This field specifies the upper 32 bits of the capacity in bytes of the Endurance Group or NVM Set to be created.

If the Operation field specifies the Create Endurance Group operation or the Create NVM Set operation, then the Capacity Upper and Capacity Lower fields specify the capacity of the Endurance Group or NVM Set to be created. If the Operation field specifies any other operation, then the Capacity Upper field and the Capacity Lower field are reserved.

5.TBD.1 Media Unit Configuration Selection

If:

- a) the Operation field specifies the Select Capacity Configuration operation;
- b) the Element Identifier field specifies a Capacity Configuration Descriptor in the Supported Capacity Configuration List; and
- c) the Selected Configuration field of the Media Unit Status log page (refer to **Figure MULog**) is cleared to 0h,

then the controller shall perform all of the following actions in sequence:

- 1) Create an Endurance Group for each Endurance Group Configuration Descriptor in the selected Capacity Configuration Descriptor.
- 2) Create an NVM Set for each NVM Set Identifier specified in each Endurance Group Configuration Descriptor, if any.

If the Operation field specifies the Select Capacity Configuration operation and the Element Identifier field is cleared to '0', then the controller shall clear the configuration by performing all of the following actions in sequence:

- 1) Delete all namespaces in the domain that contains the controller processing the command, as described in **section 8.12**.
- 2) Delete all NVM Sets in the domain that contains the controller processing the command.
- 3) Delete all Endurance Groups in the domain that contains the controller processing the command.

If the Operation field specifies the Select Capacity Configuration operation and the Element Identifier field specifies the value reported in the Selected Configuration field of the Media Unit Status log page (i.e., the currently-selected configuration), then the controller shall complete the command without error and shall make no change to the capacity configuration.

If the Operation field specifies the Select Capacity Configuration operation and:

- a) the Element Identifier field does not specify either a value of 0h or the Capacity Configuration Identifier of a Capacity Configuration Descriptor in the Capacity Configuration List; or
- b) the Selected Configuration field of the Media Unit Status log page (refer to **Figure MULog**) is not cleared to 0h,

then the controller shall abort the command with a status of Invalid Field in Command and no changes shall be made to the configuration of any Media Unit.

5.TBD.2 Endurance Group Operations

If the Operation field specifies the Create Endurance Group operation, the controller shall select a non-zero Endurance Group Identifier not assigned to an existing Endurance Group in the specified domain (refer to **Figure CMC10**) and indicate that value in Dword 0 of the completion queue entry (refer to **Figure CMCQ0**). If a non-zero unassigned Endurance Group Identifier is unavailable, then the controller shall abort the command with a status code of Identifier Unavailable.

If the Operation field specifies the Create Endurance Group operation and Media Units are supported, then the controller selects Media Units from the specified domain to be allocated to the Endurance Group.

If the Operation field specifies the Create Endurance Group operation and Media Units are not supported, then the controller selects NVM capacity from the specified domain to be allocated to the Endurance Group.

If the Operation field specifies the Create Endurance Group operation and the Capacity Lower and Capacity Upper fields specify a value that requires allocation of NVM capacity that is greater than the value in:

- a) the Max Endurance Group Capacity (MEGCAP) field in the Identify Controller data structure;
- b) the Unallocated NVM Capacity (UNVMCAP) field in the Identify Controller data structure; or
- c) the Max Endurance Group Capacity (MEGCAP) field in the Domain Attributes Entry for the domain in which the Endurance Group is being created,

then the controller:

- a) shall abort the command with Insufficient Capacity status; and
- b) if the Error Information log page is supported, shall indicate in the Command Specific Information field the total amount of NVM capacity in bytes of the largest Endurance Group that is able to be created.

If the Operation field specifies the Delete Endurance Group operation and the Element Identifier field specifies a value of 0h or the identifier of an Endurance Group that does not exist, then the controller shall abort the command with Invalid Field In Command status.

5.TBD.3 NVM Set Operations

If the Operation field specifies the Create NVM Set operation, the controller shall select a non-zero NVM Set Identifier not assigned to an existing NVM Set in the specified Endurance Group (refer to **Figure CMC10**) and indicate that value in the completion queue entry. If a non-zero unassigned NVM Set Identifier is unavailable, then the controller shall abort the command with a status code of Identifier Unavailable.

If the Operation field specifies the Create NVM Set operation and the Capacity Lower and Capacity Upper fields specify a value that requires allocation of NVM capacity that is greater than the value in the Unallocated Endurance Group Capacity (UEGCAP) field in the Endurance Group Information log page (refer to **Figure 204**) for the specified Endurance Group, then the controller:

- a) if the Error Information log page is supported, shall indicate in the Command Specific Information field the total amount of NVM capacity in bytes of the largest NVM Set that is able to be created; and
- b) shall abort the command with Insufficient Capacity status.

If the Operation field specifies the Create NVM Set operation and the Element Identifier field is cleared to 0h, then the controller shall choose an existing Endurance Group in an existing domain in which to create the NVM Set.

If the Operation field specifies the Create NVM Set operation and Media Units are supported, then the controller selects Media Units from the Endurance Group to be allocated to the NVM Set.

If the Operation field specifies the Create NVM Set operation and Media Units are not supported, then the controller selects NVM capacity from the Endurance Group to be allocated to the NVM Set.

If the Operation field specifies the Delete NVM Set operation and the Element Identifier field specifies a value of 0h or the identifier of an NVM Set that does not exist, then the controller shall abort the command with Invalid Field In Command status.

If the controller does not support NVM Sets and the Operation field specifies either the Create NVM Set operation or the Delete NVM Set operation, then the controller shall abort the command with Invalid Field In Command status.

5.TBD.4 Command Completion

Upon completion of the Capacity Management command, the controller posts a completion queue entry to the Admin Completion Queue. Capacity Management command specific status values are defined in **Figure CMStat**.

Figure CMStat: Capacity Management – Command Specific Status Values

Value	Description
26h	Insufficient Capacity: The requested operation requires more free space than is currently available. The Command Specific Information field of the Error Information log page (refer to Figure 193) specifies the total amount of NVM capacity in bytes required to create the Endurance Group or NVM Set.
2Dh	Identifier Unavailable: The number of Endurance Groups or NVM Sets supported has been exceeded.

Dword 0 of the completion queue entry contains the identifier of the Endurance Group or NVM Set created, if any. Completion queue entry Dword 0 is defined in **Figure CMCQ0**.

Figure CMCQ0: Capacity Management – Completion Queue Entry Dword 0

Bits	Description
31:16	Reserved
15:00	<p>Created Element Identifier: This field indicates the identifier of the NVM Set that was created if the Create NVM Set operation was used.</p> <p>This field indicates the identifier of the Endurance Group Identifier if the Create Endurance Group operation was used.</p> <p>This field is reserved for all other operations.</p>

...

7 Controller Architecture

7.1 Introduction

...

7.1.1 I/O Controller

...

Figure 417: I/O Controller – Admin Command Support

Command	Command Support Requirements ¹
...	
Asynchronous Event Request	M
Capacity Management	O
Namespace Management	O

...

Figure 420: I/O Controller – Log Page Support

Log Page Name	Command Support Requirements ¹
...	
Endurance Group Event Aggregate	O
Media Unit Status	M ²
Supported Capacity Configuration List	M ²
Notes: 1. O = Optional, M = Mandatory, P = Prohibited 2. Optional for controllers that do not support Fixed Capacity Management (refer to section 8.TBD).	

...

7.1.2 Administrative Controller

...

Figure 5: Administrative Controller – Admin Command Support

Command	Command Support Requirements ¹
...	
Identify	M
Capacity Management	O

...

Figure 428: Administrative – Controller Log Page Support

Log Page Name	Command Support Requirements ¹
...	
Endurance Group Event Aggregate	O
Media Unit Status	P
Supported Capacity Configuration List	P
Notes:	
1. O = Optional, M = Mandatory, P = Prohibited	

...

7.14 Privileged Actions

Privileged actions are actions (e.g., command, register write) that affect or have the potential to affect the state of the entire NVM subsystem and not only the controller and/or namespace with which the action is associated.

Admin commands that are privileged include Namespace Management, Namespace Attachment, Virtualization Management, Format NVM, **and** Sanitize, **and** Capacity Management. A privileged register action is NVM subsystem reset. Vendor specific commands and registers may also be privileged.

7.TBD NVM Capacity Model

7.TBD.1 Overview

NVM subsystems may report capacity-related information for multiple entities within the NVM subsystem. This capacity reporting model includes capacity reporting for the NVM subsystem, the domain (refer to [section TP4009_7.NEW2.1](#)), the Endurance Group (refer to [section 8.17](#)), the NVM Set (refer to [section 4.9](#)), the namespace (refer to [section 1.6.26](#)), and the Media Unit (refer to [1.6.TBD](#)). Some, all, or none of this reporting may be supported by an NVM subsystem.

[Figure FigRels](#) shows the hierarchical relationships of the entities within an NVM subsystem which are used to manage NVM capacity.

The capacity in an NVM Set is able to be allocated to one or more namespaces, and each namespace exists entirely in that NVM Set (refer to [section 4.9](#)). Not all of the capacity in the NVM Set is required to be allocated to a namespace.

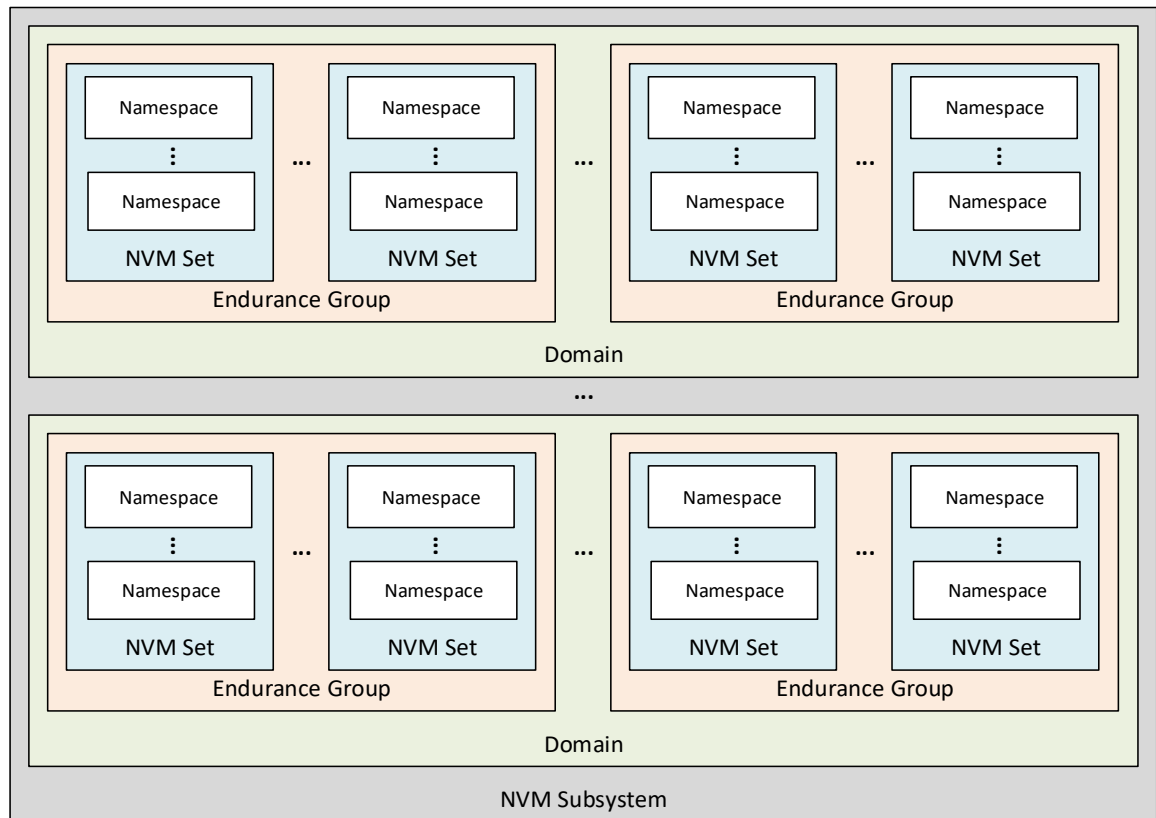
If the controller supports NVM Sets, then the capacity in an Endurance Group is able to be allocated to one or more NVM Sets, each NVM Set exists entirely in that Endurance Group (refer to [section 8.17](#)). Not all of the capacity in an Endurance Group is required to be allocated to an NVM Set.

If the controller supports Endurance Groups and does not indicate support for NVM Sets, then in all data structures that contain an NVMSETID field, the NVMSETID field shall be cleared to 0h.

If the controller does not support Endurance Groups, then in all data structures that contain an ENDGID field, the ENDGID field shall be cleared to 0h.

If the controller supports Endurance Groups, then the capacity in a domain is able to be allocated to one or more Endurance Groups, and each Endurance Group exists entirely in that domain (refer to [section TP4009_7.NEW](#)). Not all of the capacity in a domain is required to be allocated to an Endurance Group.

Figure FigRels: NVM Capacity Entity Hierarchy



NVM subsystems may report the composition of Endurance Groups and NVM Sets in terms of Media Units. Each Media Unit is allocated to exactly one Endurance Group. If NVM Sets are supported, each Media Unit is allocated to exactly one NVM Set. Data is transferred to and from Media Units via Channels. Each Media Unit is connected to one or more Channels. Each Channel is connected to one or more Media Units.

A host uses Capacity Management (refer to [section 8.TBD](#)) to allocate:

- Domain capacity to Endurance Groups;
- Endurance Group capacity to NVM Sets;
- Media Units to Endurance Groups; and
- Media Units to NVM Sets,

as part of creating those entities.

A host uses the Namespace Management create operation (refer to [section 8.12](#)) to allocate capacity to namespaces.

7.TBD.2 Media Unit Organization Examples

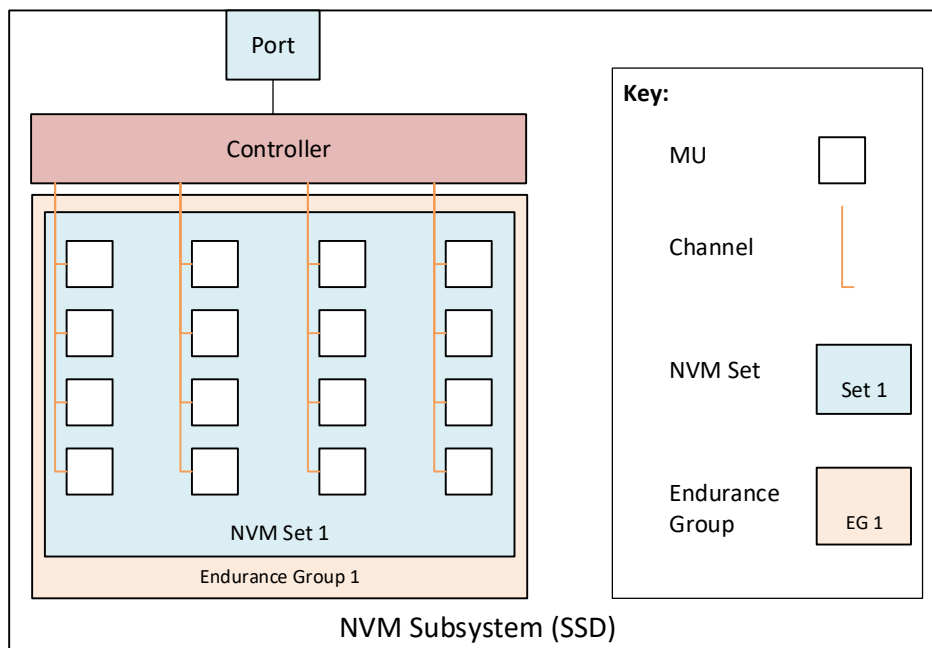
Allocation of Media Units is used to organize the physical NVM resources in an NVM subsystem to meet particular performance goals.

The following examples show an NVM subsystem with all resources in a single domain. The domain has four Channels, with four Media Units attached to each Channel.

7.TBD.2.1 Simple NVM Subsystem

In the example shown in [Figure FigSimpleOrg](#), endurance is managed across all media units. The performance goal is maximum bandwidth, which is achieved by allowing each read or write operation to simultaneously access all Media Units. All Media Units are in the same Endurance Group and in the same NVM Set.

Figure FigSimpleOrg: Simple NVM Subsystem

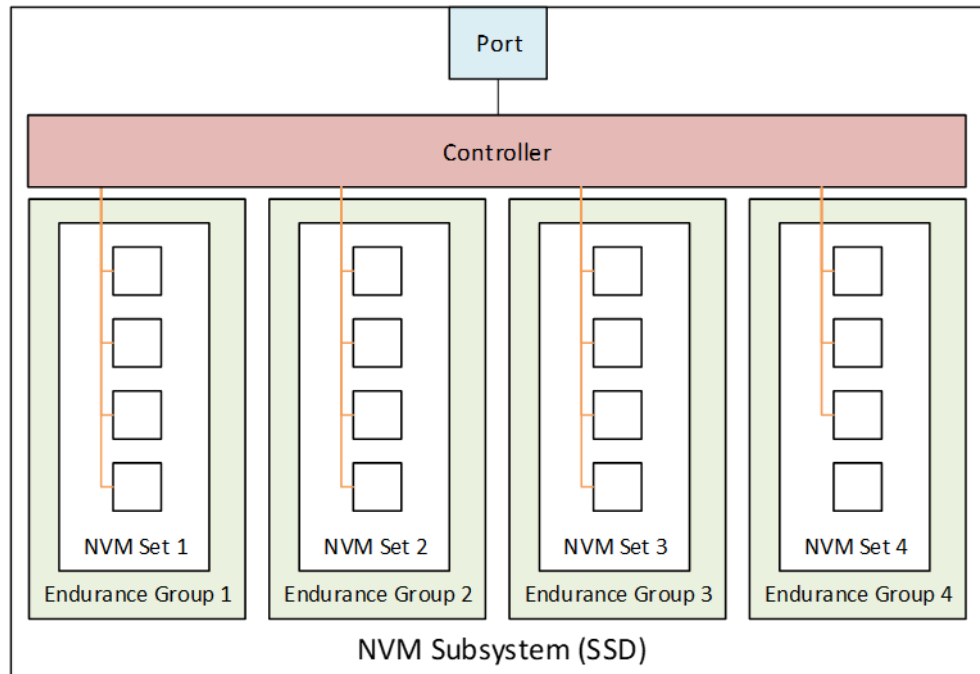


The Capacity Configuration Descriptor for this example contains one Endurance Group Configuration Descriptor. The Endurance Group Configuration Descriptor contains one NVM Set Identifier and four Channel Configuration Descriptors. Each Channel Configuration Descriptor contains four Media Unit Configuration Descriptors.

7.TBD.2.2 Vertically-Organized NVM Subsystem

In the example shown in [Figure FigVerticalOrg](#), the performance goal is isolation among four NVM Sets at the cost of bandwidth. Endurance is managed separately for each set. Media Units sharing a Channel are allocated to the same Endurance Group. All Media Units in an Endurance Group are allocated to the same NVM Set. The bandwidth for any NVM Set is likely to be less than or equal to the bandwidth of the Channel of that NVM Set.

Figure FigVerticalOrg: Vertically-Organized NVM Subsystem

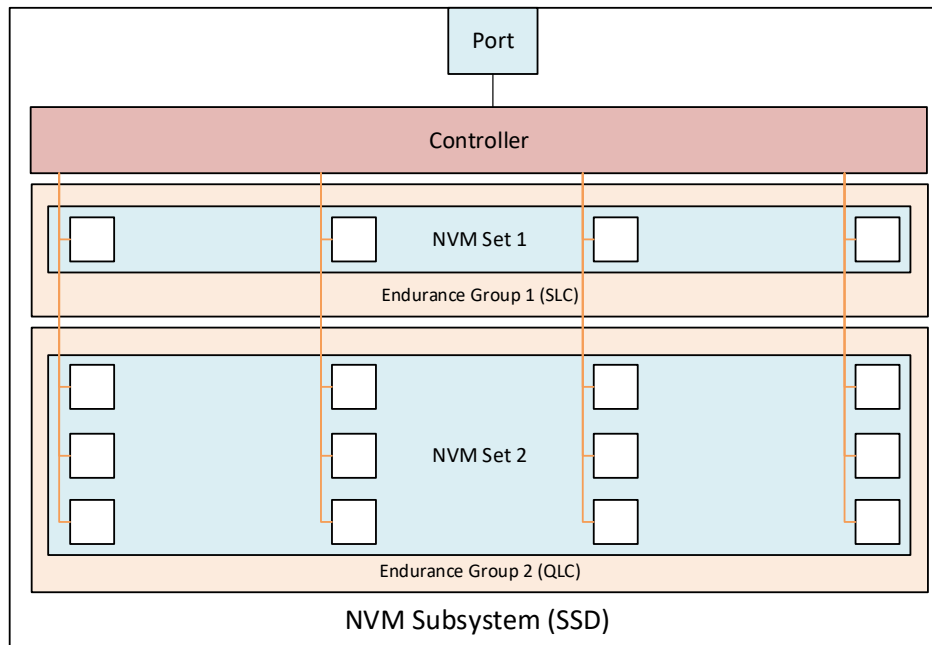


The Capacity Configuration Descriptor for this example contains four Endurance Group Configuration Descriptors. Each Endurance Group Configuration Descriptor contains one NVM Set Identifier and one Channel Configuration Descriptor. Each Channel Configuration Descriptor contains four Media Unit Configuration Descriptors.

7.TBD.2.3 Horizontally-Organized Dual NAND NVM Subsystem

In the example shown in [Figure FigDualOrg](#), the Media Units are NAND which is capable of being operated as QLC or at a lower density. The performance goal is for maximum bandwidth to a small NVM Set operating as SLC and to a larger NVM Set operating as QLC.

Figure FigDualOrg: Horizontally-Organized Dual NAND NVM Subsystem



The Capacity Configuration Descriptor would contain two Endurance Group Configuration Descriptors. The first Endurance Group Configuration Descriptor for this example:

- indicates a Capacity Adjustment Factor of approximately 1,600;
- contains one NVM Set Identifier; and
- contains four Channel Configuration Descriptors. Each Channel Configuration Descriptor contains one Media Unit Configuration Descriptor.

The second Endurance Group Configuration Descriptor for this example:

- indicates a Capacity Adjustment Factor of 100;
- contains one NVM Set Identifier; and
- contains four Channel Configuration Descriptors. Each Channel Configuration Descriptor contains three Media Unit Configuration Descriptors.

7.TBD.3 Capacity Reporting

For an NVM subsystem that does not support multiple domains, the capacity information reported in the Identify Controller data structure (i.e., the TNVMCAP field and the UNVMCAP field in [Figure 247](#)) describes the capacity for the NVM subsystem. If the MEGCAP field is non-zero, that field indicates the largest entity (e.g., Endurance Group, NVM Set, namespace) that may be created in the NVM subsystem.

For an NVM subsystem that supports multiple domains, the capacity information reported in the Identify Controller data structure describes the capacity accessible by the controller processing the Identify command. The host may use the Identify command to access the Domain List data structure (refer to [section TP4009_5.15.2.99](#)) to determine the domains that are accessible by the controller and the capacity information for each of those domains. If the Max Endurance Group Domain Capacity field is non-zero, then the field describes the largest entity (e.g., Endurance Group, namespace) that may be created by this controller in the domain described by that Domain Attributes Entry.

For an NVM subsystem that supports Endurance Groups (refer to [section 8.17](#)), the host may use the Identify command to access the Endurance Group List data structure (refer to [section 5.15.2.TBD1](#)) to determine the Endurance Groups that are accessible by the controller. To determine the capacity information for each

Endurance Group, the host uses the Get Log Page command to access the Endurance Group Information log page (refer to section 5.14.1.9).

For an NVM subsystem that supports NVM Sets (refer to section 4.9), the host may use the Identify command to access the NVM Set List data structure (refer to section 5.15.2.5) to determine the NVM Sets that are accessible by the controller and the capacity information for each of those NVM Sets.

For the management of Endurance Groups, NVM Sets, and namespaces, Figure CAP_USE_1 describes the effects of the support of NVM Sets, Endurance Groups, and domains on which capacity information is used for each management operation.

Figure CAP_USE_1: Capacity Information Field Usage

Entity being created / deleted	NVM Sets supported	Endurance Groups supported	Domains supported	Capacity information used ¹
Endurance Group ⁷	n/a	Yes	No	NVM Subsystem ³
			Yes	Domain ⁴
NVM Set ⁷	Yes	Yes ²	n/a	Endurance Group ⁵
Namespace ⁸	No	No	No	NVM Subsystem ³
			Yes	Domain ⁴
		Yes	n/a	Endurance Group ⁵
	Yes	Yes ²	n/a	NVM Set ⁶
NOTES: 1. This information described in this column is used by the host for creating the entity (e.g., to determine if there is sufficient available capacity) and this information is altered by the controller as a result of the creation or deletion of the entity (e.g., unallocated capacity decreased as a result of entity creation, or unallocated capacity increased as a result of entity deletion). 2. NVM Set support requires support for Endurance Groups as described in section 4.9. 3. Capacity information in the Identify Controller data structure (i.e., TNVMCAP field, UNVMCAP field, and MEGCAP fields (refer to Figure 247)). 4. Capacity information in the Domain Attributes Entry (i.e., Total Domain Capacity field, Unallocated Domain Capacity field, and Max Endurance Group Domain Capacity field (refer to TP4009 Figure 9999)). 5. Capacity information in the Endurance Group Information log page (i.e., TEGCAP field, UEGCAP fields (refer to Figure 204)). 6. Capacity information in the NVM Set Attributes Entry (i.e., Total NVM Set Capacity field, and Unallocated NVM Set Capacity field (refer to Figure 251)). 7. Endurance Groups and NVM Sets are created and deleted using the Capacity Management command (refer to section 5.TBD). 8. Namespaces are created and deleted using the Namespace Management command (refer to section 8.12). Namespaces are deleted using the Capacity Management command.				

...

8 Features

...

8.8 Reservations (Optional)

...

8.8.3 Reservation Types

...

Figure 6: Command Behavior in the Presence of a Reservation

NVMe Command	Write Exclusive Reservation		Exclusive Access Reservation		Write Exclusive Registrants Only or Write Exclusive All Registrants Reservation		Exclusive Access Registrants Only or Exclusive Access All Registrants Reservation	
	Non-Registrant	Registrant	Non-Registrant	Registrant	Non-Registrant	Registrant	Non-Registrant	Registrant
...								
NVM Write Command Group								
Capacity Management (Admin)								
Dataset Management								
Flush								
Format NVM (Admin)								
Namespace Attachment (Admin)								
Namespace Management (Admin)	C	C	C	C	C	A	C	A
Sanitize (Admin)								
Security Send (Admin)								
Write								
Write Uncorrectable								
Write Zeroes								
...								
All Other Commands Group								
All other commands ¹	A	A	A	A	A	A	A	A
Key: A definition: A=Allowed, command processed normally by the controller C definition: C=Conflict, command aborted by the controller with status Reservation Conflict Notes: 1. The behavior of a vendor specific command is vendor specific.								

...

8.12 Namespace Management (Optional)

...

The total and unallocated NVM capacity for the NVM subsystem is reported in the Identify Controller data structure (refer to [Figure 247](#)). For controllers that support NVM Sets, the total and unallocated NVM capacity for each NVM Set is reported as part of the NVM Set Attributes Entry (refer to [Figure 251](#)). For each namespace, the NVM Set [and the Endurance Group](#) in which the namespace is allocated ~~is~~ [are](#) reported in the Identify Namespace data structure. The NVM Set to be used for a namespace is based on the value in the NVM Set Identifier field in a create operation. If the NVM Set Identifier field is cleared to 0h in a create operation, then the controller shall choose the NVM Set from which to allocate the namespace.

If the NVM Set Identifier field and the Endurance Group Identifier field are both cleared to 0h in a create operation, then the controller shall choose the Endurance Group and the NVM Set from which to allocate the namespace.

If the NVM Set Identifier field is cleared to 0h and the Endurance Group Identifier field is set to a non-zero value in a create operation, then the controller shall choose the NVM Set in the specified Endurance Group from which to allocate the namespace.

If the NVM Set Identifier field is set to a non-zero value and the Endurance Group Identifier field is cleared to 0h in a create operation, then the controller shall abort the command with a status of Invalid Field in Command.

If the NVM Set Identifier field and the Endurance Group Identifier field are both set to non-zero values in a create operation and the specified NVM Set exists in the specified Endurance Group, then the controller shall allocate the namespace in the specified NVM Set.

If the NVM Set Identifier field and the Endurance Group Identifier field are both set to non-zero values in a create operation and the specified NVM Set does not exist in the specified Endurance Group, then the controller shall abort the command with a status of Invalid Field in Command.

...

8.17 Endurance Groups (Optional)

...

If Endurance Groups are supported, then the NVM subsystem and all controllers shall:

- Indicate support for Endurance Groups in the Controller Attributes field in the Identify Controller data structure;
- Indicate the Endurance Group Identifier with which the namespace is associated in the Identify Namespace data structure; and
- Support the Endurance Group Information log page.

If Endurance Groups are not supported and the host sends a command in which an Endurance Group Identifier field is defined (e.g., Get Log Page), then that field is not used.

If Endurance Groups are not supported and the controller returns information to the host that contains an Endurance Group Identifier field, then that field shall be cleared to 0h.

8.17.1 Configuring and Managing Events

...

8.20 Asymmetric Namespace Access Reporting (Optional)

...

8.20.4 Asymmetric Namespace Access States Command Processing Effects

...

Figure 494: ANA effects on Command Processing

Command	ANA State	Effects on command processing
...		

Figure 494: ANA effects on Command Processing

Command	ANA State	Effects on command processing
Get Log Page	ANA Inaccessible, ANA Persistent Loss, or ANA Change	<p>The following log pages are affected:</p> <ul style="list-style-type: none"> a) Error Information (i.e., 01h): The log page may contain entries only for namespaces whose relationship to the controller processing the command is in the ANA Optimized state (refer to section 8.20.3.1) or the ANA Non-optimized state (refer to section 8.20.3.2). <p>The following log pages are not available ¹:</p> <ul style="list-style-type: none"> a) Media Unit Status log page (refer to section 5.14.1.TBD0); and b) Supported Capacity Configuration List log page (refer to section 5.14.1.TBD1).
...		
<p>NOTES:</p> <ol style="list-style-type: none"> 1. If the ANA state is ANA Inaccessible State, then commands that use feature identifiers or log pages that are not available shall fail with a status code of Asymmetric Access Inaccessible. If the ANA state is ANA Persistent Loss State, then commands that use feature identifiers or log pages that are not available shall fail with a status code of Asymmetric Access Persistent Loss. If the ANA state is ANA Change State, then commands that use feature identifiers or log pages that are not available shall fail with a status code of Asymmetric Access Transition. 		

8.TBD Capacity Management (Optional)

8.TBD.1 Overview

Capacity Management is a capability for organizing physical media into Endurance Groups and NVM Sets. There are two different forms of Capacity Management, Fixed Capacity Management and Variable Capacity Management. A controller that supports Capacity Management shall support at least one form.

The host uses Fixed Capacity Management to create Endurance Groups and NVM Sets by selecting a configuration which explicitly allocates Media Units (refer to [section 8.TBD.2](#)) to Endurance Groups and NVM Sets.

The host uses Variable Capacity Management to:

- create a single Endurance Group by specifying the desired capacity;
- create a single NVM Set by specifying the desired capacity;
- delete a single Endurance Group; and
- delete a single NVM set.

The Capacity Adjustment Factor is the ratio between the capacity consumed by an Endurance Group from the Unallocated NVM Capacity field in the Identify Controller data structure and the total NVM capacity in that Endurance Group, i.e.:

$$\text{Capacity Adjustment Factor} = \text{INTEGER} \left(\frac{\text{Capacity Consumed}}{\text{Endurance Group Capacity}} * 100 \right)$$

(E.g., the value 200 indicates that creation of an Endurance Group with a total NVM capacity of 5 GiB consumes 10 GiB of the Unallocated NVM Capacity indicated by the controller).

If an Endurance Group is created, then the controller performs the following actions as an atomic operation:

- a) the value indicated by the Unallocated NVM Capacity (UNVMCAP) field of the Identify Controller data structure is changed based on the requested capacity, the Capacity Adjustment Factor of the created Endurance Group, and the granularity at which the controller allocates NVM capacity; and
- b) the Endurance Group Identifier is added to the Endurance Group List.

If an Endurance Group is deleted, then the controller performs the following actions in sequence:

- 1) the Endurance Group Identifier is removed from the Endurance Group List;
- 2) if the Media Unit Status log page is supported, then the Endurance Group Identifier field is cleared to 0h in all Media Unit Status Descriptors, if any, that indicate the deleted Endurance Group;
- 3) every NVM Set in the Endurance Group is deleted; and
- 4) the value indicated by the Unallocated NVM Capacity (UNVMCAP) field of the Identify Controller data structure is increased by the value that was indicated by the Total Endurance Group Capacity (TEGCAP) field of the Endurance Group Information log page of the deleted Endurance Group.

If any of the entities modified by the above sequence are accessed after the sequence begins and before it completes, then the results are indeterminate.

If an NVM Set is created, then the controller performs the following actions as an atomic operation:

- a) the NVM Set Identifier is added to the NVM Set List; and
- b) the Unallocated Endurance Group Capacity indicated by the Endurance Group Information log page (refer to **Figure 204**) is decreased by the amount of capacity allocated to the NVM Set; the controller may allocate NVM capacity in units such that the requested size for an NVM Set may be rounded up to the next unit boundary.

If an NVM Set is deleted, then the controller performs the following actions in sequence:

- 1) the NVM Set Identifier is removed from the NVM Set List;
- 2) if the Media Unit Status log page is supported, then the NVM Set Identifier field is cleared to 0h in all Media Unit Status Descriptors, if any, that indicated the deleted NVM Set;
- 3) for each namespace in the deleted NVM Set:
 - a. all commands targeting the namespace are handled as described for namespace deletion in **section 8.12**;
 - b. the namespace identifier is removed from the Allocated Namespace ID list;
 - c. the namespace is deleted;
 - d. the namespace identifier is added to the Changed Namespace List; and
 - e. a Namespace Attribute Changed event is generated for hosts other than the host which issued the Capacity Management command;

and

- 4) the Unallocated Endurance Group Capacity indicated by the Endurance Group Information log page is increased by the amount of capacity formerly allocated to the NVM Set.

If any of the entities modified by the above sequence are accessed after the sequence begins and before it completes, then the results are indeterminate.

If an NVM Set is created or deleted, the value indicated by the Unallocated NVM Capacity (UNVMCAP) field of the Identify Controller data structure is not changed.

8.TBD.2 Fixed Capacity Management

A Media Unit represents a component of the underlying media in a domain. An implementation may choose to represent each die as a separate Media Unit; however, this is not required. A Media Unit is the smallest media component for which the controller reports wear information (refer to the Available Spare field and the Percentage Used field in the Media Unit Status Descriptor, **Figure MUnitD**).

Two or more I/O operations to a Media Unit at the same time may interfere with each other as they contend for resources internal to or shared by that Media Unit.

A controller supporting Fixed Capacity Management:

- a) shall support the Media Unit Status log page (refer to [section 5.14.1.TBD0](#));
- b) shall support Endurance Groups (refer to [section 8.17](#));
- c) may support NVM Sets (refer to [section 4.9](#));
- d) shall set the Fixed Capacity Management bit to '1' in the CTRATT field of the Identify Controller data structure (refer to [Figure 247](#));
- e) shall support the Supported Capacity Configuration List log page (refer to [section 5.15.1.TBD](#)); and
- f) shall support the Select Capacity Configuration operation of the Capacity Management command (refer to [section 5.TBD](#)).

Media Units are accessed by way of Channels that represent communication pathways that may be a source of contention. This information is reported to facilitate the composition of NVM Sets that minimize interference between independent writers competing for this type of resource.

The host allocates the Media Units in a domain to Endurance Groups and NVM Sets by:

- 1) reading the Supported Capacity Configuration List log page (refer to [Figure SCCLog](#)) and selecting the desired configuration; and
- 2) issuing a Capacity Management command specifying the Select Capacity Configuration operation and the Capacity Configuration Identifier of the desired configuration.

Following successful completion of the command, each Media Unit is allocated to one Endurance Group and to one NVM Set. The resulting configuration of Media Units is reported by the Media Unit Status log page (refer to [section 5.14.1.TBD0](#)).

8.TBD.3 Variable Capacity Management

Variable Capacity Management allows the dynamic creation and deletion of Endurance Groups and NVM Sets.

A controller supporting Variable Capacity Management:

- a) may support the Media Unit Status log page;
- b) shall support Endurance Groups;
- c) may support NVM Sets;
- d) shall set the Variable Capacity Management bit to '1' in the CTRATT field of the Identify Controller data structure (refer to [Figure 247](#));
- e) shall support the Create Endurance Group operation of the Capacity Management command;
- f) may support the Delete Endurance Group operation of the Capacity Management command; and
- g) if NVM sets are supported:
 - a. shall support the Create NVM Set operation of the Capacity Management command;
 - b. shall report non-zero values in the Total Endurance Group Capacity field and the Unallocated Endurance Group Capacity field in the Endurance Group Information log page (refer to [Figure 204](#)); and
 - c. may support the Delete NVM Set operation of the Capacity Management command.

If a controller supports the Delete Endurance Group operation of the Capacity Management command, then it shall set the Delete Endurance Group bit to '1' in the CTRATT field of the Identify Controller data structure.

If a controller supports the Delete NVM Set operation of the Capacity Management command, then it shall set the Delete NVM Set bit to '1' in the CTRATT field of the Identify Controller data structure.

A typical sequence of operations for allocating capacity is:

- 1) determine the available capacities in each domain (refer to [section TP4009_7.NEW](#));
- 2) create Endurance Group with desired capacity (refer to [section 5.TBD](#));
- 3) create NVM Set with desired capacity in the Endurance Group (refer to [section 5.TBD](#)); and
- 4) create namespace with desired capacity in the NVM Set (refer to [section 5.20](#)).

A typical sequence of operations for deallocating capacity is:

- 1) delete namespace, if any, (refer to [section 5.20](#));
- 2) delete NVM Set, if any, (refer to [section 5.TBD](#)); and

3) delete Endurance Group (refer to [section 5.TBD](#)).

If there is insufficient unallocated capacity in an Endurance Group for the controller to create an NVM Set, then the host can delete one or more NVM Sets in that Endurance Group and create the new NVM Set using some or all of the available capacity.

If there is insufficient unallocated capacity in a domain for the controller to create an Endurance Group, then the host can delete one or more Endurance Groups in that domain and create a new Endurance Group using some or all of the available capacity.

Add the following changes to NVMe-MI 1.1:

6 NVMe Express Admin Command Set

...

< Note to editor: Add the Capacity Management command to Figure 110. >

Figure 495: List of NVMe Admin Commands Supported using the Out-of-Band Mechanism

Command	NVMe Storage Device O/M/P ¹	NVMe Enclosure O/M/P ¹
...		
Capacity Management	O	P
...		
NOTES: 1. O/M/P definition: O = Optional, M = Mandatory, P = Prohibited from being supported. An NVMe Enclosure that is also an NVMe Storage Device (i.e., implements namespaces) shall implement mandatory commands required by either an NVMe Storage Device or an NVMe Enclosure and may implement optional commands allowed by either an NVMe Storage Device or an NVMe Enclosure. Mandatory commands shall be supported if the NVMe Controller specified by the Controller ID field supports the command.		

...

< Note to editor: Add the Capacity Management command to Figure 116 as shown. >

Figure 496: Command Messages Allowed During Sanitize Operation

Command Set	Command Message	Allowed During Sanitize Operation ¹
...		
NVMe Admin Command Set ²	Capacity Management	Same restrictions as defined by the NVM Express specification
	Device Self-test	
	Firmware Activate/Commit	
	Firmware Image Download	
	Format NVM	
	Get Features	
	Get Log Page	
	Identify	
	Namespace Attachment	
	Namespace Management	
	Sanitize	
	Security Receive/Send	
	Security Send	
	Set Features	
	Vendor Specific	
	Virtualization Management	
...		
NOTES: 1. Refer to the NVM Express specification for the definition of a sanitize operation. 2. NVMe Admin Commands that are prohibited via the out-of-band mechanism (refer to Figure 495) are not listed since they are always prohibited including during a sanitize operation.		

< end of changes >