



LEGAL NOTICE:

© Copyright 2008 to 2024 NVM Express, Inc. ALL RIGHTS RESERVED.

This 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 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: "© 2008 to 2024 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, Inc.
3855 SW 153rd Drive
Beaverton, OR 97003 USA
info@nvmexpress.org

NVM Express Technical Proposal for New Feature

Technical Proposal ID	TP4146b Flexible Data Placement
Change Date	2024-05-30
Builds on Specification	NVM Express® Base Specification 2.0b NVM Express® NVM Command Set Specification 1.0b NVM Express® Zoned Namespace Command Set Specification 1.1b NVM Express® Management Interface Specification 1.2b
Ratified Technical Proposals Referenced	TP4090 Deallocation Granularity Enhancement TP4095a Namespace Capability Reporting TP4074a Defining Scope for Features TP4115 ZNS Namespace Management Enhancements
Technical Proposals in Development Referenced	

Technical Proposal Author(s)

Name	Company
Bill Martin, Judy Brock, Dan Helmick, Robert Moss, Mike Allison, Benjamin Lim, Jiwon Chang	Samsung
Ross Stenfort, Young Ahn, Wei Zhang, Sumit Gupta	Meta
Amber Huffman, Chris Sabol	Google
Paul Suhler, Mahinder Saluja, John Geldman, Mark Carlson, Fred Knight	KIOXIA
Walt Hubis, Dan Hubbard, Steven Wells	Micron
Santosh Kumar, Andrés Baez	SK hynix / Solidigm
Kwok Kong	Microchip
Erich Haratsch	Marvell
Anu Murthy, Hyunmo Kang, Yeong-Jae Woo	FADU
Mike James, Yoni Shternhell	WDC

This proposal adds an enhancement of the NVM Command Set to enable host guided data placement. The proposal introduces a Reclaim Unit (RU) that consists of physical non-volatile storage with the property that a RU may be erased/reused/repurposed without disturbing any other Reclaim Units. A RU is a physical construct only, and is unrelated to logical addresses or LBAs. A namespace consists of capacity allocated in Reclaim Units. A Reclaim Group (RG) contains one or more RUs. There are one or more Placement Handles that reference RUs which I/O commands can target. I/O commands performed on one RG do not interfere with the performance or reliability of commands executed on another RG.

A Placement Identifier (PID) specifies a specific RG/Placement Handle pair that references an RU available for writing random LBAs (i.e., the Flexible Data Placement of user data to the non-volatile storage allocated to the RU). The written capacity of an RU referenced by a PID increments on each write command specifying to use that PID and the PID can be modified to reference another RU once the capacity of the RU is fully written. The host tracks the set of LBAs written to an RU. This enables the host to deallocate all LBAs associated with a particular RU at the same time to minimize garbage collection by the controller, thereby reducing write amplification. The host is responsible for managing Placement Identifiers and related device resources.

Revision History

Revision Date	Change Description
2022-05-18	Initial version.
2022-05-19	Renaming contents of DSPEC and using Reclaim Unit Handle for the host reference to writing to Reclaim Units
2022-05-24	Deleted text not needed in the ARUTHR field. Updated Figure Z to correct color coding. Changed the name "Direct Placement" to "Data Placement". Incorporated Steven Wells (Micron) editorial comments. Integrated editorial comments from Paul Suhler. Incorporated edits from special 2 hour meeting.
2022-05-25	Updated based on the comments made by that 2 hour special meeting.
2022.05.31	<ul style="list-style-type: none"> Accepted all changes. Moved Reclaim Unit Handle change to Dataset Management command to a new command. Made the Reclaim Unit Nominal Size field to be 8 bytes (Hyunmo Kang (FADU)) Clarified Endurance Groups. Fixed editorial items identified by Yoni Shternhell (WDC). Fixed editorial items from Dan Hubbard (Micron). Updated the Reclaim Unit Handle Descriptor description to describe each entry is processed separately by the controller and that outstanding commands that affect an entry may or may not be included in the returned data. Incorporated aborting the I/O Management Send command if FDP Mode is not enabled or the namespace does not support the new directive. Created a new Reclaim Unit Handle type that is the Host Isolated type but without the required FDP events.
2022-06-01	<ul style="list-style-type: none"> As Identified by Dan Hubbard (Micron): Updated the Reclaim Unit Handle List Data Structure to return the default Placement Identifier for the namespace as it is specific to a RG and RUH.

Revision Date	Change Description
2022-06-08	<ul style="list-style-type: none"> Fixed two figures with the same name of Reclaim Unit Handle Descriptor by changing the one defined in the Reclaim Unit Handle Status log page. Corrected the byte definitions in some of the data structures. Changed the Active Reclaim Unit Time Remaining to be an estimate as accuracy is not measurable. Changed the Default Placement Identifier of a namespace to a Reclaim Unit Handle and when more than one Reclaim Group exists, the controller is allowed to pick the Reclaim Group on a write using the default Reclaim Unit Handle. Micron (Dan Hubbard) requested to add the RUHF field to the Identify Flexible Data Placement data structure so the host can just query the Reclaim Unit Handles and usage of that value in the Placement Identifier. Made some editorial changes. Fixed identified editorial issues. Fixed editorial items from Kioxia (Paul Suhler).
2022-06-14	<ul style="list-style-type: none"> Moved the functionality to the Base specification. Removed the Host Isolated with/without events and added a feature to enable the events.
2022-06-15	<ul style="list-style-type: none"> Added the FDP Events Feature to allow the host to enable/disable events associated with a Host Isolated Reclaim Unit Handle. Incorporated comments from Paul Suhler (Kioxia) and Ross Stenfort. This included requirements of ordering in the lists.
2022-06-16	<ul style="list-style-type: none"> Incorporated comments from Paul Suhler (Kioxia). Added comment from Dan Hubbard on the MBE field.
2022-06-29	<ul style="list-style-type: none"> Editorial updates as sent to me by David Black, Paul Suhler, Dan Helmick. Did not get all of them in yet. Removed the scope of the Reclaim Unit Handles to a Reclaim Group be the request of Matias Bjorling this caused the following to change: <ul style="list-style-type: none"> FDP Configurations have to list all Reclaim Unit Handles and indicate the Reclaim Group association Reclaim Group was removed from the Placement Identifier Removed the Reclaim Unit Handle Format field Allowed a dynamic list of Reclaim Unit Handles to be specified by the host in the Namespace Management command
2022-07-13	<ul style="list-style-type: none"> Reverted the Reclaim Unit Handles to be associated with all Reclaim Groups.
2022-07-14	<ul style="list-style-type: none"> Updated the TP to reflect the FDP model accepted in the 7/13/2022 special TP4146 meeting.
2022-07-14	<ul style="list-style-type: none"> Per the Straw Polls in the 7/14 meeting I updated the TP to have the FDP configuration scope to an Endurance Group and not supporting NVM Sets. As reported by Jiwon Chang (Samsung) The Reclaim Unit Handle Statistics log page requires the host to specify a Placement Identifier to identify the Reclaim Group and the Namespace Reclaim Unit Handle.

Revision Date	Change Description
2022-07-18	<ul style="list-style-type: none"> • Editorial fixes especially updating the section and figure numbers references. • Uses a byte offset in the Namespace Management command of the NVM Command Set specification to locate the variable sized Namespace Reclaim Unit Handle Identifier list. • Added examples to figure ZZ (FDP Model) for I/O commands using the Data Placement Directive and not using the Data Placement Directive. • Still to do; power safe, managing FDP events, editorial feedback from Paul Suhler and David Black.
2022-07-19	<ul style="list-style-type: none"> • Made the Endurance Group Event Aggregate log page required to be supported if more than one Endurance Group is supported in the NVM subsystem.
2022-07-20	<ul style="list-style-type: none"> • Allow FDP to use volatile write cache in a backward way • Updated the feature to enable/disable FDP Events to allow the host to enable/disable a set of FDP Event Types. The controller is able to report the list of supported FDP Events and attributes associated with each
2022-07-28	<ul style="list-style-type: none"> • Incorporated feedback from David Black. • Removed the field specifying the maximum Reclaim Units usage in the Namespace Management command of the NVM Command Set Specification. That parameter is independent of FDP and is to be put into a separate TP. • Due to member requests, renaming the "Namespace Reclaim Unit Handle" to the "Placement Handle Identifier" as just having a single word separation from the Reclaim Unit Handle was complicating the understanding. • Removed the number of Reclaim Unit Handles limited to the number of RG. • Increases of reported number of Reclaim Unit Handle Status Descriptors. • Incorporated editorial comments from Paul Suhler.
2022-07-28	<ul style="list-style-type: none"> • Updated the volatile write cache to the decision made in the Technical WG meeting on 7/28/2022. This includes adding a new field in the I/O Command Set Independent Identify Namespace data structure to indicate if the namespace has a volatile write cache present. • Changed "active Reclaim Unit Handle" to "active Reclaim Unit". • Clarified that the Reclaim Group Identifier in the Placement Identifier may be reserved or does not exist depending on the values in the RUHF field and the NRG field. • Minor editorial fixes.

Revision Date	Change Description
2022-08-09	<ul style="list-style-type: none"> • Incorporated feedback from Paul Suhler (Kioxia). • Replaced the Reclaim Unit Handle Status log page with the I/O Management Receive command. • The FDP Configuration List number of entries are static so the CNS for reading the current FDP configuration is not required because the host can get the index from the Get Feature and look up the index. • Changed the volatile write cache field in the I/O Command Set Independent Identify Namespace data structure to be only 1 bit as Christoph Hellwig indicated this was better for the LINUX code. • Fixed the number of Placement Handle Identifiers to a maximum of 128. • On a Namespace Management command that does not specify any Placement Handle Identifiers, then the controller shall select a Reclaim Unit Handle that was not used by a namespace create with a Namespace Management command that did specify a Placement Handle Identifiers. • Removed the CNS value to return the list if Placement Handle Identifiers for the namespace as the host can use the I/O Management Receive command to obtain the information. • Added a log page to obtain the list of Reclaim Unit Handles and the association to an existing namespaces in the Endurance Group. • Updated the Directive enable to abort enabling the Data Placement Directive if the NSID is set to FFFFFFFFh since a specific Endurance Group cannot be specified by the host.
2022-08-10	<ul style="list-style-type: none"> • Incorporate changes from Dan Helmick (Samsung). • Added the RUH and the RG to the FDP events as the controller would provide that information over the Placement Identifier. • Fixed the numbering of the data structure for Flexible Data Placement Events. • Incorporated feedback from Judy Brock (Samsung). • Renamed the “Reclaim Unit Handle Format” field to “Placement Identifier Format” and made the entry a 0's based number. • Changed the FDP Events feature to Endurance Group scope.

Revision Date	Change Description
2022-08-15	<ul style="list-style-type: none"> • Editorial fixes • Moved text from the Volatile Write Cache Present bit in the Identify Controller data structure to the FDP Volatile Write Cache bit in the FDP Configuration Descriptor. • Updated the definition of Volatile Write Cache Not Present (VWCNP) bit to specify it overrides the Volatile Write Cache Present bit in the Identify Controller data structure. • Limited the number of entries in the Invalid Placement Handle Identifier List of the Namespace Management command which changed the new error code. • All FDP Event field that are allowed to be indicated as valid or not, shall be cleared to 0h by the controller and the host should ignore the fields.
2022-08-16	<ul style="list-style-type: none"> • Corrected the list of log pages in the NVMe-MI specification • Incorporated comments from Robert Moss (Samsung). • Updated the fields in FDP Statistics log page to indicate the scope is an Endurance Group. • Added more naming suggestions. • Incorporated editorial comments from Paul Suhler. • Clarified the use of “Reclaim Group Identifier” to just “Reclaim Group”. • Updated the EARUTR field to indicate the time returned is in reference to the time the Get Log Page command was processed. • Corrected the use of comparing an identifier value to the number of identifiers to use “greater than or equal to”
2022-08-18	<ul style="list-style-type: none"> • Added suggested TBD values • Changed “Host Isolated” to “Initially Isolated” • Changed “Totally Isolated” to “Persistently Isolated” • Removed the use of “Mode” in “Flexible Data Placement Mode”
2022-08-19	<ul style="list-style-type: none"> • Add a Size field to the FDP Configuration log page to assist the host in determining the buffer size of a variable sized log page. • Updated the padding for the FDP Configuration Descriptor to be bounded to the “next” 8-byte boundary (i.e., 1 to 7 bytes in size). • Updated the I/O Management Receive command for the Reclaim Unit Handle Status to indicate if the host reads more than the size of the Reclaim Unit Handle Status data structure, then the controller returns zeroes.

Revision Date	Change Description
2022-08-22	<ul style="list-style-type: none"> Editorial fixed from Dan Hubbard (Micron) comments: Added the aborting of a Directive Send command if the Data Placement is enabled on a namespace not in and Endurance Group with FDP enabled. Specified that the PIV and LV bits are cleared to '0' is an FDP event is not associated to a specific Reclaim Unit Handle. Added FDP Event Feature support required if FDP is supported. Editorial fixes from Paul Suhler (Micron). Clarified the text associated with the controller selecting the Reclaim Unit Handle on the Namespace Management command when the NPHID field is cleared to 0h. Updated the NVMe-MI log page and feature support requirements per Management Interface Task Group review on 8/22/2022. Pre Christoph Hellwig made the following updates <ul style="list-style-type: none"> Changed the Reclaim Unit Handle Attributes in the Reclaim Unit Handle Usage log page to be an enumeration as opposed to a bit set. Moved the FDPS bit into the CTRATT field as there is no need to reserve a whole byte for a single bit.
2022-08-22	<ul style="list-style-type: none"> Modified the Return Parameters Data Structure to allow each directive except the Identify Directive to state whether the host directed state is preserved across Controller Level Resets. The Data Placement Directive is set to '1' is supported. The Streams directive is cleared to '0' to maintain current definition and backwards compatibility. This request is from Kioxia and Micron. Both Meta and Google prefer not having to issue a Directive Send command to each namespace after a Controller Level Reset.
2022-08-23	<ul style="list-style-type: none"> Updated the Placement Identifier data structure to make it clear that PIF field is a 0's based value.
2022-08-25	<ul style="list-style-type: none"> Per Request from Christoph Hellwig, I changed the PIF field to reflect the number of bits for the Reclaim Group Identifier as it makes the definition less complicated. Editorial fixes identified by Fred Knight. Per the Technical WG review: <ul style="list-style-type: none"> Split the Placement Identifier format to two separate diagrams Updated the text about the consequences on parallel I/O commands that utilize a Placement Identifier that is being modified by an I/O Management Send command Removed redline tracking of changes prior to the meeting.

Revision Date	Change Description
2022-08-31	<ul style="list-style-type: none"> Added a comment to clarify the VWCNP bit is per namespace and the ignoring of the Volatile Write Cache Present bit is also per namespace. Corrected the reserved size in the header for the FDP Events log page.
2022-09-01	<ul style="list-style-type: none"> Added comments to FDP Configuration to move elements around that will be fixed in Phase 3.
2022-09-09	<ul style="list-style-type: none"> Per Jesper Devantier (Samsung) corrected the byte alignment of 16b field on 16b boundaries in the FDP Statistics log page.
2022-09-19	<ul style="list-style-type: none"> Changed “I/O write command” to “write command”. Clarified that the VWCNP bit is for a specific namespace. Clarified that changing the value of the Flexible Data Placement feature affects fields associated with the Format Index in all Identify Namespace data structures. Updated the FDPVWC bit and moved the last paragraph to section 7.1 per comment by David Black. Moved the RUH type values to the FDP configuration log page section and updated the text associated with the types. Clarified the CNS values data structures affected by changing the Flexible Data Placement Feature value. Removed the word “Mode” as the WG decided to get rid of it. Editorial items from David Black. Editorial items from Paul Suhler.
2022-09-20	<ul style="list-style-type: none"> Editorial items from Ross Stenfort. From Paul Suhler: <ul style="list-style-type: none"> Fixed naming of NPHI to NPHID field. Added Flexible Data Placement Events to the Features supported tables.
2022-09-21	<ul style="list-style-type: none"> Changes due to comments from Solidigm: <ul style="list-style-type: none"> Updated Figure FDPEX3 and text to make it cleared that the placement of user data for a namespace is directed by a host in write commands. Updated FDP Overview to reflect that the capacity of user data for a namespace is allocated to an RU on a write command. Tied the text to figure FDPEX3 to align to text that states the same. Request to rename “Placement Handle Identifier” to “Placement Handle” to separate it from “Placement Identifier: Changes due to comments from Paul Suhler (Kioxia) <ul style="list-style-type: none"> editorial fixes Add a Get Features request to the host actions to determine the FDP events supported.
2022-09-22	<ul style="list-style-type: none"> Updated Figure ZZ with the new name Placement Handle.

Revision Date	Change Description
2022-09-22	<ul style="list-style-type: none"> Changes due to NVM Technical WG Review <ul style="list-style-type: none"> Updated section 2.3.1 with statement about the placement of namespace user data.
2022-09-22	<ul style="list-style-type: none"> Made the file ready for 30 day member review.
2022-09-29	<ul style="list-style-type: none"> Chris Sabol (Google) <ul style="list-style-type: none"> The 16-byte fields in the FDP Statistics log page need a 16-byte roll over value. Dan Hubbard (Micron) <ul style="list-style-type: none"> Editorial comments Reworded the Controller Specified bit and Host Specified bit to use the Reclaim Unit Handle Attributes field as that field was changed from a bit vector to an enumerated value. Fixed reserved bits in CDW12 of the FDP Events feature. Removed the step to initialize the Data Placement Directive after a CLR as the setting is now persistent. Changed “Invalid Reclaim Unit Handle Identifier” to “Invalid Placement Handle List” Paul Suhler (Kioxia) <ul style="list-style-type: none"> Added comment to the RAUMW field to be addressed.
2022-09-30	<ul style="list-style-type: none"> Jiwon Chang (Samsung) <ul style="list-style-type: none"> Clarified that for the FDP Events feature, the Reclaim Unit Handle is specified by the NSID field and the Placement Handle. Added a requirement to the FDP Events feature that a NSID value of FFFFFFFFh is not supported. Paul Suhler (Kioxia) <ul style="list-style-type: none"> Clarified “active” in the EARUTR field.
2022-10-04	<ul style="list-style-type: none"> Edward Hsieh (Silicon Motion) <ul style="list-style-type: none"> Fix note1 in figure 105 of the NVM Command Set Specification to use the text from version 1.0b of the specification, Paul Suhler (Kioxia) <ul style="list-style-type: none"> Clarified what an Invalid Placement Identifier is for the Invalid Placement Identifier FDP Event.

Revision Date	Change Description
2022-10-10	<ul style="list-style-type: none"> Robert Moss (Samsung) <ul style="list-style-type: none"> Definitions in Figure FDPRUHD need to indicate the user data is to remain in the same RG. Sumanth Venkata (Samsung) <ul style="list-style-type: none"> Editorial comments Edward Hsieh (Silicon Motion) <ul style="list-style-type: none"> Figure FDPConfig should use N instead of N-1 as N is 0's based. Corrected the byte offsets in Figure FDPConfig, Figure FeatDS, FeatGF, and PIDL Editorial fixes Clarified the size of the Placement Handle list is defined by I/O Command Set specifications in Figure ZZ Paul Suhler <ul style="list-style-type: none"> Added comment about handling a bad RG or Placement Handle Changed "associated to" to "associated with" Editorial changes Added a statement that a sanitize operation may modify Reclaim Unit Handles to reference a different Reclaim Unit.
2022-10-11	<ul style="list-style-type: none"> Added a limit on the number of Placement Identifiers in the I/O Management Send command. Addressed all previous comments
2022-10-13	<ul style="list-style-type: none"> Added a requirement that namespaces that share the same RUH have to be formatted with the same Format Index.
2022-10-14	<ul style="list-style-type: none"> Maxing MAXPIDs 0's based to match I/O Management command reported by Edward Hsieh (Silicon Motion). Updated Figure FigLogDED to reflect that the Reclaim Group Identifier and the Reclaim Unit Handle Identifier are the controller selected values if the Invalid Placement Identifier FDP Event is logged.
2022-10-15	<ul style="list-style-type: none"> Editorial items reported by Paul Suhler (Kioxia)
2022-10-23	<ul style="list-style-type: none"> Editorial items from Hobin Lee (Samsung) Editorial items by author
2022-10-27	<ul style="list-style-type: none"> Updated Host Bytes with Metadata Written field to specify writes to the Endurance Group. Judy Brock (Samsung) Editorial feedback from Dan Hubbard (Micron)
2022-10-31	<ul style="list-style-type: none"> Updated the Figure FDPConfig (i.e., FDP Configuration Descriptor) byte offsets per Jiwon Chang (Samsung)

Revision Date	Change Description
2022-11-03	<ul style="list-style-type: none"> Changed black text to blue in section 5.27.1.TBD1 from Paul Suhler (Kioxia) Editorial changes by the author. Includes marking deleted text that moved to be purple instead of red (to make sure the reader understands the text was not actually deleted from the specification). Updated the MAXPIDS field requirement to be less than (NRG*NRUH) to account for the field being a 0's based number. Updated the definition of the ERUTL field in the FDP configuration Descriptor to match the definition on the EARUTR field in the I/O Management Receive command. These need to have the same definition.
2022-11-04	<ul style="list-style-type: none"> Moved data structures specifying LBAs to the NVM Express NVM Command Set Specification. Updated the references to Figure PIDRG to also reference Figure PIDRG_NORGI. Paul Suhler (Kioxia) <ul style="list-style-type: none"> Editorial fixes Update the definition of Persistently Isolated to remove the phrase "moved by the controller".
2022-11-07	<ul style="list-style-type: none"> Paul Suhler (Kioxia) <ul style="list-style-type: none"> Added introduction to section 4.1.4.TBD in the NVM Command Set Specification.
2022-11-08	<ul style="list-style-type: none"> Updated the HBMW and MBMW fields to indicate that it is about processing write commands that have a Directive Type field and a Directive Specific field defined by the command. This TP adds these fields to the Write Zeroes command and the Write Uncorrectable commands.
2022-11-09	<ul style="list-style-type: none"> Diffed the member review version with the latest member review comment resolution to make sure redlines are correct for discussion.
2022-11-15	<ul style="list-style-type: none"> Formatted tables Removed empty tables that moved.
2022-11-16	<ul style="list-style-type: none"> Updated the HBMW field and MBMW field to refer to applicable I/O Command Set for the list of I/O commands. Then updated the NVM Command Set specification to list the commands.
2022-11-17	<ul style="list-style-type: none"> Edits during NVM Express Technical WG meeting were approved for integration.
2022-11-27	<ul style="list-style-type: none"> Integrated
2022-11-29	<ul style="list-style-type: none"> Editorial feedback from Mike Allison and Paul Suhler.
2022-11-30	<ul style="list-style-type: none"> Editorial feedback from Mike Allison and Paul Suhler.

Revision Date	Change Description
12-12-2022	<ul style="list-style-type: none"> Clarified the DTYPE values that affect the controller utilizing Placement Handle 0. Removed Directives NVM Subsystem Reset requirements as CLR requirements were added. Clarified reference to CNS 00h in the Enabling Flexible Data Placement section.
1-13-2023	<ul style="list-style-type: none"> Added a requirement that the Streams Directive is not allowed to be supported if the Data Placement Directive is supported.
1-19-2023	<ul style="list-style-type: none"> Added that a Directive Send command is abort if the command causes Streams Directive and Data Placement Directive to be enabled on a namespace at the same time.
1-27-2023	<ul style="list-style-type: none"> Updated Placement Handle 0 Associated RUH to be optional. Made the reporting of the Error Information Entry for enabling Streams Directive and Data Placement Directive optional.
2-2-2023	<ul style="list-style-type: none"> Editorial changes in NVMe Technical WG review. Ready for 30 day member review.
3-21-2023	<ul style="list-style-type: none"> Removed redundant and incorrect text in section 8.7.3. Corrected name of Directive Send command. Fixed capitalizations Changed the Get Log Page command to the I/O Management Receive command in the definition of the EARUTR field.
3-23-2023	<ul style="list-style-type: none"> Change “specifies” to “indicates” for data reported by the controller.
5-30-2023	<ul style="list-style-type: none"> Integrated 4146a changes
5-31-2023	<ul style="list-style-type: none"> Incorporate Bug 9 Correct namespace for NSID usage
1-3-2024	<ul style="list-style-type: none"> Corrected the field name in the NVM Command Set specification section 4.1.6. Incorporated Bug 274 correcting a misleading “i.e.”
1/17/2024	<ul style="list-style-type: none"> Corrected text that should have referenced the Placement Handle 0 Associated RUH field. Rephrased text that implied an Endurance Group created by a Namespace Management command to reflect a namespace created by that command in the specified Endurance Group.
2/8/2024	<ul style="list-style-type: none"> Updated the Placement Identifier Valid (PIV) bit to state that the bit is identifying if there is a reported value in the Placement Identifier (PID) field.

Revision Date	Change Description
03/08/2024, 03/20/2024 and 03/29/2024 (J Brock)	<ul style="list-style-type: none"> Clarified that the NOET field should be ignored on a Get Features for Flexible Data Placement Events FID 1Eh since Dword 0 of the CQE identifies the # of supported FDP Event Types in the data buffer for a Get Features Added reference in Figure FeatDS and Figure SFDPED to Figure FigLogDED which contains the actual FDP Event Type value definitions. <p><u>Changes approved in TPAR4188</u></p> <ul style="list-style-type: none"> Mandated that all namespaces in an EG with FDP enabled shall be associated with the NVM Command Set (e.g., ZNS namespaces are prohibited in an EG with FDP enabled) Prohibited the enablement of Streams Directive on any namespace that is contained in an EG with FDP enabled.
04/10/2024	<ul style="list-style-type: none"> Minor editorial change – converted hard-to-parse sentence to list format
04/11/2024	<ul style="list-style-type: none"> Changed two occurrences of “valid” to “reported” This version approved in WG for re-entering 30 day member review
04/23/2024	<ul style="list-style-type: none"> Integrate feedback from Kioxia – clarified which bits in CQE Dword 0 (the least significant 8 bits) indicate the value of NOET (i.e., the number of supported FDP Event Types that are contained in the data buffer) for a Get Features command for Flexible Data Placement Events FID 1Eh. Integrate feedback from Solidigm – one editorial comment
05/23/2024	<ul style="list-style-type: none"> Integrate additional feedback from Solidigm – two additional editorial comments (change 2 occurrences of “Identifier” to “Handle”) Integrate additional feedback from Kioxia – added CQE Dword 0 figure (Figure FeatCQ)
05/30/2024	<ul style="list-style-type: none"> Accept all changes; approved for integration
6/7/2024	<ul style="list-style-type: none"> Editorial changes. Clarified the Enable Directive may not be enabled due to restrictions.

Description of Specification Changes

Description for Changes Document for NVM Express® Base Specification 2.0b

New Features/Feature Enhancements/Required Changes:

- Flexible Data Placement (optional)
 - Description of change.
 - Added definitions of new terms.
 - Added the I/O Management Receive command to retrieve Reclaim Unit Handle Status descriptors.
 - Added the I/O Management Send command to provide control over Placement Identifiers.
 - Added the Flexible Data Placement feature to enable or disable FDP and to select the FDP configuration to apply to an Endurance Group.
 - Added new log pages:
 - FDP Statistics
 - FDP Events
 - FDP Configurations
 - Reclaim Unit Handle Usage
 - Added the new Flexible Data Placement model section:
 - Reclaim Group and Reclaim Unit concept and usage
 - FDP operations
 - Added the Flexible Data Placement feature to figures listing feature support and feature identifiers.
 - Added a new Data Placement Directive and allow each directive to specify if the enablement is persistent across a Controller Level Reset.
 - **New requirement**
 - Added a new Data Placement Directive and allow each directive to specify if the enablement is persistent across a Controller Level Reset.
 - ~~The Streams Directive and the Data Placement Directive are not allowed to be enabled in a namespace at the same time.~~
 - Approved in TPAR 4188, agreed to include in TP4146b:
 - The Streams Directive is not allowed to be enabled on any namespace in an Endurance Group that has Flexible Data Placement enabled.
 - All namespaces in an Endurance Group that has Flexible Data Placement enabled shall be associated with the NVM Command Set.
 - **Incompatible change**
 - Changed the support the Endurance Group Event Aggregate log page when Endurance Groups are supported to only be required if more than one Endurance Group is supported by the NVM subsystem.
 - References
 - Technical Proposal 4146

Description for Changes Document for NVM Express® NVM Command Set Specification 1.0b

New Features/Feature Enhancements/Required Changes:

- Flexible Data Placement (FDP) (optional)
 - Description of change.
 - Added fields to the NVM Command Set I/O Command Set specific Identify Namespace data structure.
 - Modified the Namespace Management command to specify creation of a namespace in non-volatile storage that has Flexible Data Placement enabled.
 - Added directives to the Write Zeroes command and Write Uncorrectable command.
 - Added the definition for the Reclaim Unit Handle Statue Descriptor in the I/O Management Receive command.
 - Added the definition of the Media Reallocation Event Type for the Flexible Data Placement Events log page.
 - **New requirement / incompatible change**
 - None
 - References:
 - Technical Proposal 4146

Description for Changes Document for NVM Express® Zoned Namespace Command Set Specification 1.1b

New Features/Feature Enhancements/Required Changes:

- Flexible Data Placement (optional)
 - Description of change.
 - Modified the Host Software Specified Fields figure to expand the range of the bytes marked as “Refer to the NVM Command Set Specification”.
 - **New requirement / incompatible change**
 - None
 - References
 - Technical Proposal 4146

Description for Changes Document for NVM Express® Management Interface Specification 1.2c

New Features/Feature Enhancements/Required Changes:

- Flexible Data Placement (optional)
 - Description of change.
 - Added new log pages to a figure listing log pages.
 - Added the Flexible Data Placement feature to a figure listing feature support.
 - **New requirement / incompatible change**
 - None
 - References
 - Technical Proposal 4146

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.
Purple: Text that moved to this location without changes.
~~Purple~~: Text that was moved to another location.
<Green Bracketed>: Notes to editor

Modify portions of the NVM Express Base Specification 2.0b as follows:

1 Introduction

...

1.5 Definitions

...

1.5.TBD_{S1} Placement Identifier

A data structure that specifies a Reclaim Group Identifier and a Placement Handle that references a Reclaim Unit. Refer to [Figure PIDRG](#) and [Figure PIDRG_NORGI](#).

1.5.TBD_{S2} Placement Handle

A namespace scoped handle that maps to an Endurance Group scoped Reclaim Unit Handle which references a Reclaim Unit in each Reclaim Group.

1.5.TBD_{S3} Reclaim Group (RG)

An entity that contains one or more Reclaim Units. Refer to section [3.2.TBD](#).

1.5.TBD_{S4} Reclaim Unit (RU)

A logical representation of non-volatile storage within a Reclaim Group that is able to be physically erased by the controller without disturbing any other Reclaim Units. Refer to section [3.2.TBD](#).

1.5.TBD_{S5} Reclaim Unit Handle (RUH)

A controller resource that references a Reclaim Unit in each Reclaim Group. Refer to section [3.2.TBD](#).

...

2 Theory of Operation

...

2.3 NVM Storage Model

...

2.3.1 Storage Entities

The NVM storage model includes the following entities:

- NVM subsystems (refer to 1.5.40)
- Domains (refer to section 3.2.4)
- Endurance Groups (refer to section 3.2.3)
- [Reclaim Groups and Reclaim Units \(refer to section 3.2.TBD\)](#)
- NVM Sets (refer to section 3.2.2)
- Namespaces (refer to section 3.2.1)

As illustrated below,

- each domain is contained in a single NVM subsystem;
- each Endurance Group is contained in a single domain and may contain either:
 - one or more NVM Sets; or
 - one or more Reclaim Groups;
- each NVM Set is contained in a single Endurance Group and each namespace is contained in a single NVM Set. Each Media Unit is contained in a single Endurance Group; and
- each Reclaim Group is contained in a single Endurance Group, each Reclaim Unit is contained in a Reclaim Group, and each namespace is contained in an Endurance Group within one or more Reclaim Units of the Reclaim Groups in that Endurance Group.

Each Endurance Group is composed of storage media, which are termed either Media Units (refer to section 8.3.2) or Reclaim Units (refer to section 3.2.TBD). Reclaim Unit Handles reference a Reclaim Unit in each Reclaim group for writing user data. For clarity, Media Units, Reclaim Unit Handles, and Reclaim Units are not shown in the examples in this section that follow.

Figure 11 shows the hierarchical relationships of these entities within a simple NVM subsystem, which has:

- one domain;
- one Endurance Group;
- one NVM Set; and
- one namespace.

Figure 11: Simple NVM Storage Hierarchy with NVM Sets




Figure FDPEX1 shows the hierarchical relationships in a simple NVM subsystem, which has:

- one domain;
- one Endurance Group;
- one Reclaim Group; and
- one namespace with user data written to the single Reclaim Group.

The placement (i.e., which Reclaim Group) of user data for a namespace is directed by each host write command to that namespace (refer to section 8.TBD).

Figure FDPEX1: Simple NVM Storage Hierarchy with One Reclaim Group

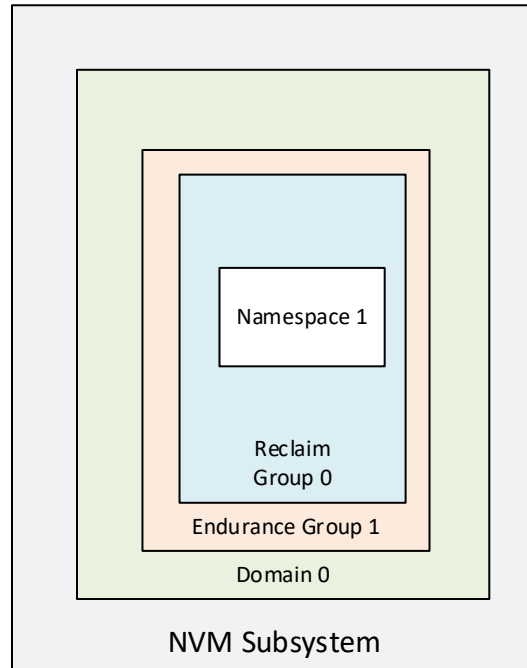


Figure FDPEX2 shows the hierarchical relationships in a simple NVM subsystem, which has:

- one domain;
- one Endurance Group;
- four Reclaim Groups; and
- one namespace with user data written to each Reclaim Group.

The placement (i.e., which Reclaim Group) of user data for a namespace is directed by each host write command to that namespace (refer to section 8.TBD).

Figure FDPEX2: Simple NVM Storage Hierarchy with Multiple Reclaim Groups

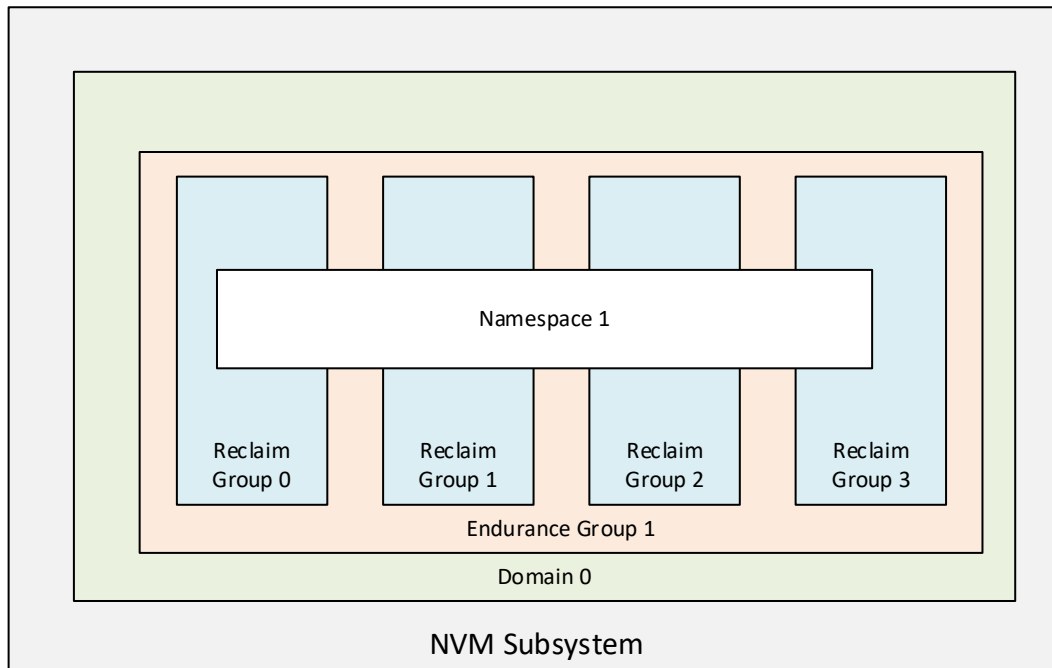
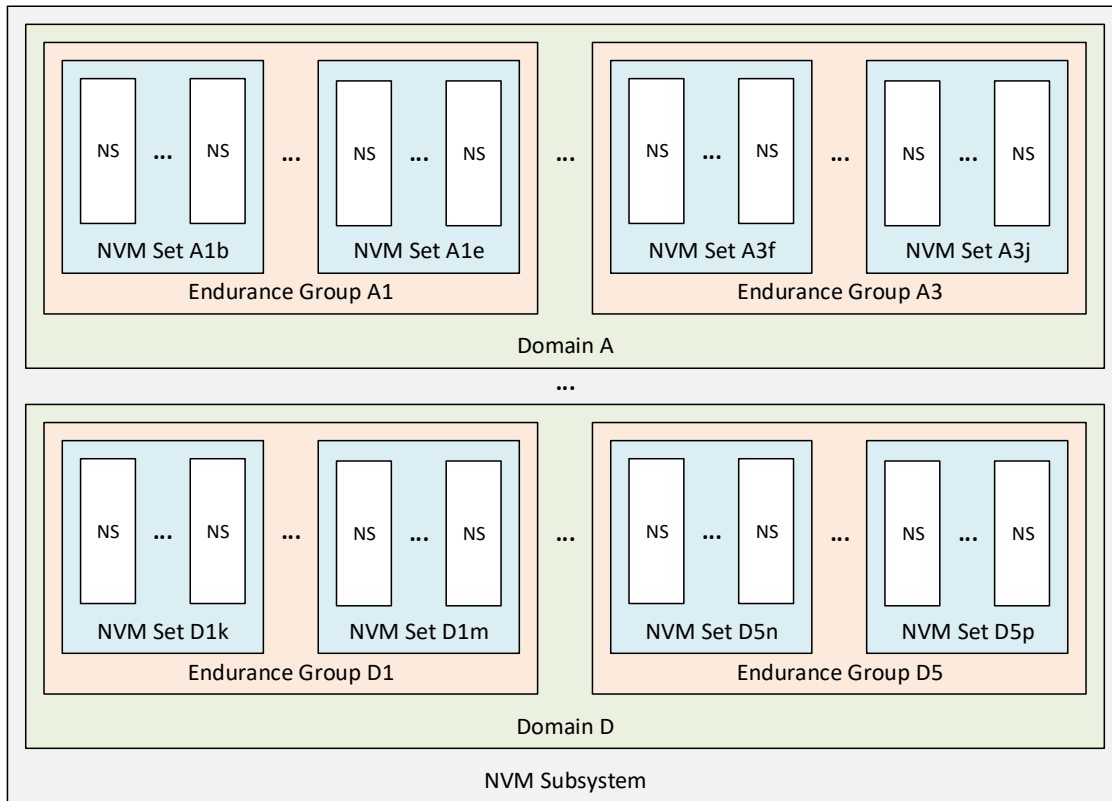


Figure 12 shows the relationships of these entities in a complex NVM subsystem, which has:

- multiple domains;
- multiple Endurance Groups per domain;
- multiple NVM Sets per Endurance Group; and
- multiple namespaces per NVM Set.

Figure 12: Complex NVM Storage Hierarchy with NVM Sets



Entity naming key (Abc):

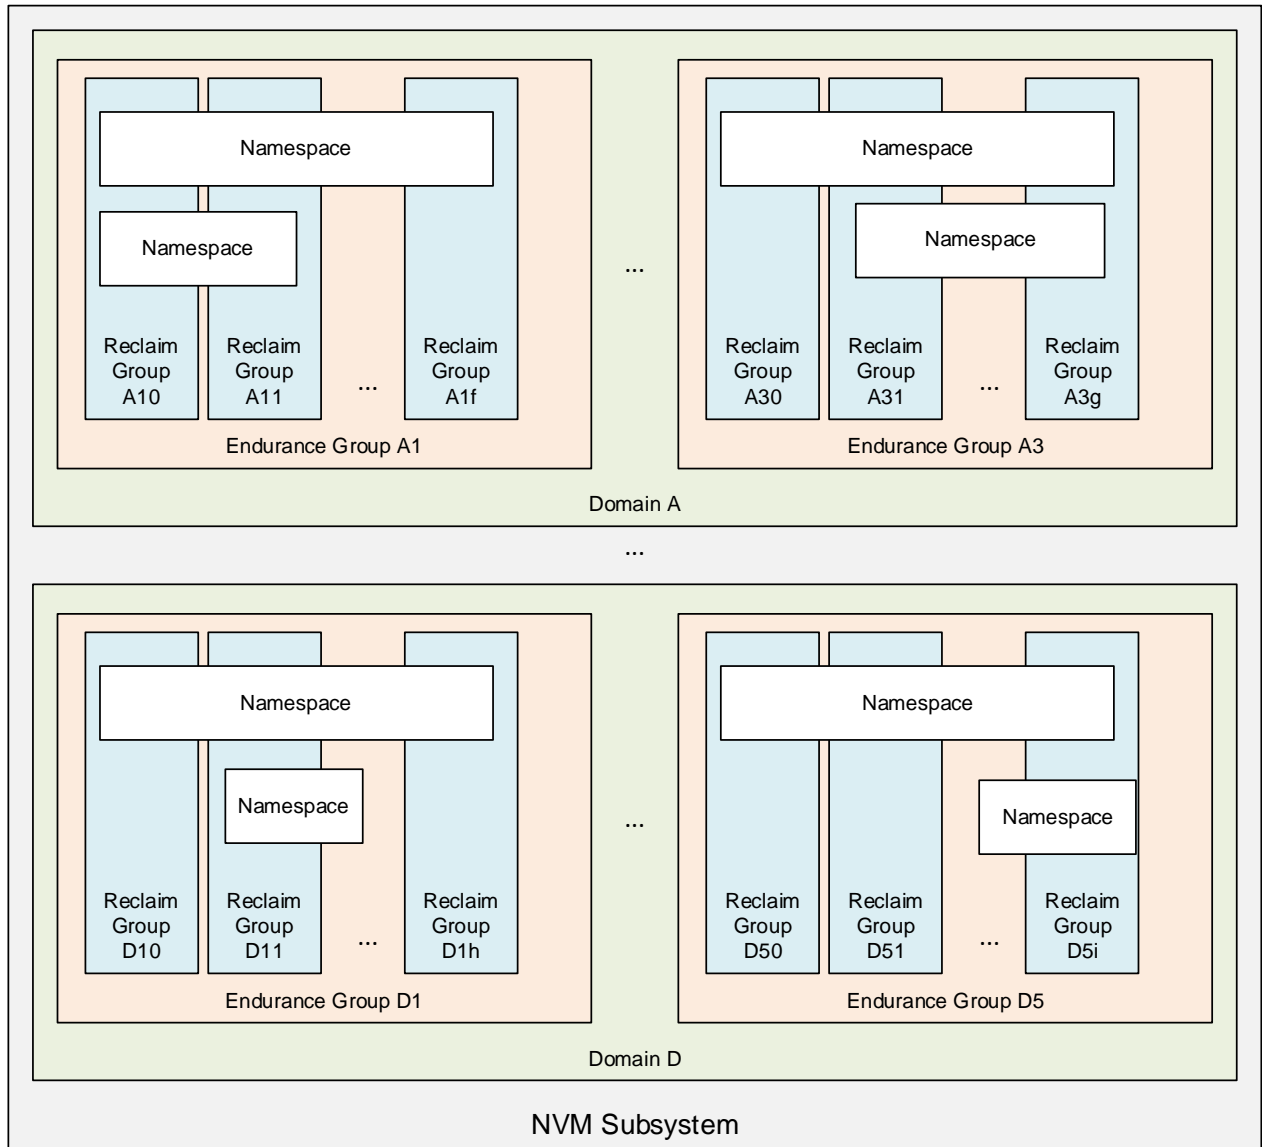
- A: Domain (capital letter)
- b: Endurance Group (digit)
- c: NVM Set (lower case letter)

Figure FDPEX3 shows the relationships in a complex NVM subsystem, which has:

- multiple domains;
- multiple Endurance Groups per domain;
- multiple Reclaim Groups per Endurance Group; and
- multiple namespaces per Endurance Group.

The placement (i.e., which Reclaim Group) of user data for a namespace is directed by each host write command to that namespace (refer to section 8.TBD).

Figure FDPEX3: Complex NVM Storage Hierarchy with Multiple Reclaim Groups



Entity naming key (Abc):

- A: Domain (capital letter)
- b: Endurance Group (digit)
- c: Reclaim Group (digit or lower case letter for maximum number)

The support **Reporting** of Endurance Groups, **Reclaim Groups within an Endurance Group**, or NVM Sets **within an Endurance Group** is optional, but the storage model supports these concepts. An NVM subsystem may be shipped by the vendor with its storage entities configured, or it may be configured or re-configured by the customer. Typical changes to the configuration are creation and deletion of namespaces.

An NVM subsystem that does not support multiple NVM Sets does not require reporting of NVM Sets. An NVM subsystem that does not support multiple Endurance Groups does not require reporting of Endurance Groups.

...

3 NVM Express Architecture

3.1 NVM Controller Architecture

...

3.1.2 Controller Types

...

3.1.2.1. I/O Controller

...

3.1.2.1.1 Command Support

...

Figure 23: I/O Controller –Common I/O Command Support

Command	Command Support Requirements ¹
...	
Reservation Release	O ²
I/O Management Receive	O ³
I/O Management Send	O ³
Notes:	
1. O/M/P definition: O = Optional, M = Mandatory, P = Prohibited 2. Mandatory if reservations are supported as indicated in the Identify Controller data structure. 3. Mandatory for controllers that support the Flexible Data Placement capability (refer to section 8.TBD). Refer to the FDPS bit in Figure 275.	

...

3.1.2.1.2 Log Page Support

...

Figure 24: I/O Controller – Log Page Support

Log Page Name	Log Page Support Requirements ¹
...	
Rotational Media Information	O
FDP Configurations	O ⁴
Reclaim Unit Handle Usage	O ⁴
FDP Statistics	O ⁴
FDP Events	O ⁴
Notes:	
1. O/M/P definition: O = Optional, M = Mandatory, P = Prohibited 2. Mandatory for controllers that support Fixed Capacity Management (refer to section 8.3.2). 3. Optional for NVM Express revision 1.4 and earlier. 4. Mandatory for controllers that support the Flexible Data Placement capability (refer to section 8.TBD). Refer to the FDPS bit in Figure 275.	

...

3.1.2.1.3 Features Support

...

Figure 25: I/O Controller – Feature Support

Feature Name	Feature Support Requirements ¹	Logged in Persistent Event Log ¹
...		
Rotational Media	O	O
Flexible Data Placement	O ⁶	O ⁶
Flexible Data Placement Events	O ⁶	O ⁶
Notes: 1. O/M/P/NR definition: O = Optional, M = Mandatory, P = Prohibited, NR = Not Recommended 2. The feature is mandatory for NVMe over PCIe. This feature is not supported for NVMe over Fabrics. 3. Mandatory if reservations are supported as indicated in the Identify Controller data structure. 4. Mandatory if reservations are supported by the namespace as indicated by a non-zero value in the Reservation Capabilities (RESCAP) field in the Identify Namespace data structure. 5. This feature is optional for NVM subsystems that do not implement a Management Endpoint. For NVM subsystems that implement any Management Endpoint refer to the NVMe Management Interface Specification. 6. Mandatory for controllers that support the Flexible Data Placement capability (refer to section 8.TBD). Refer to the FDPS bit in Figure 275.		

...

3.1.2.2 Administrative Controller

...

3.1.2.2.1 Log Page Support

...

Figure 29: Administrative Controller – Log Page Support

Log Page Name	Command Support Requirements ¹
...	
Rotational Media	P
Boot Partition	O
FDP Configurations	P
Reclaim Unit Handle Usage	P
FDP Statistics	P
FDP Events	P

Figure 29: Administrative Controller – Log Page Support

Log Page Name	Command Support Requirements ¹
I/O Command Set Specific Log Pages	P
Notes: 1. O/M/P definition: O = Optional, M = Mandatory, P = Prohibited 2. Optional if Set Features command is not supported (refer to Figure 28). 3. Optional if NVMe-MI Send command and NVMe-MI Receive command is not supported (refer to Figure 28). 4. Optional for NVMe Express revision 1.4 and earlier.	

3.1.2.2.3 Features Support

...

Figure 30: Administrative Controller – Feature Support

Feature Name	Feature Support Requirements ¹	Logged in Persistent Event Log ¹
...		
Flexible Data Placement	P	P
Flexible Data Placement Events	P	P
...		

...

3.1.2.3 Discovery Controller

...

3.1.2.3.1 Log Page Support

...

Figure 33: Discovery Controller – Log Page Support

Log Page Name	Command Support Requirements ¹
...	
Boot Partition	P
FDP Configurations	P
Reclaim Unit Handle Usage	P
FDP Statistics	P
FDP Events	P
I/O Command Set Specific Log Pages	P
Notes: 1. O/M/P definition: O = Optional, M = Mandatory, P = Prohibited 2. Optional if Set Features command is not supported (refer to Figure 32). 3. Optional if NVMe-MI Send command and NVMe-MI Receive command is not supported (refer to Figure 32). 4. Optional for versions 1.1 and earlier of the NVMe over Fabrics specification.	

...

3.1.2.3.4 Features Support

...

Figure 34: Discovery Controller – Feature Support

Feature Name	Feature Support Requirements ¹	Logged in Persistent Event Log ¹
...		
Flexible Data Placement	P	P
Flexible Data Placement Events	P	P
...		

...

3.2 NVM Subsystem Entities

...

3.2.3 Endurance Groups

Endurance may be managed within a single NVM Set (refer to section 3.2.2) or across a collection of NVM Sets. Each NVM Set is associated with an Endurance Group (refer to Figure 278). If two or more NVM Sets have the same Endurance Group Identifier, then endurance is managed by the NVM subsystem across that collection of NVM Sets. If only one NVM Set is associated with a specific Endurance Group Identifier, then endurance is managed locally to that NVM Set.

If NVM Sets are not supported, then endurance is managed by the NVM subsystem:

- within each Endurance Group if Endurance Groups are supported; or
- within the domain if Endurance Groups are not supported.

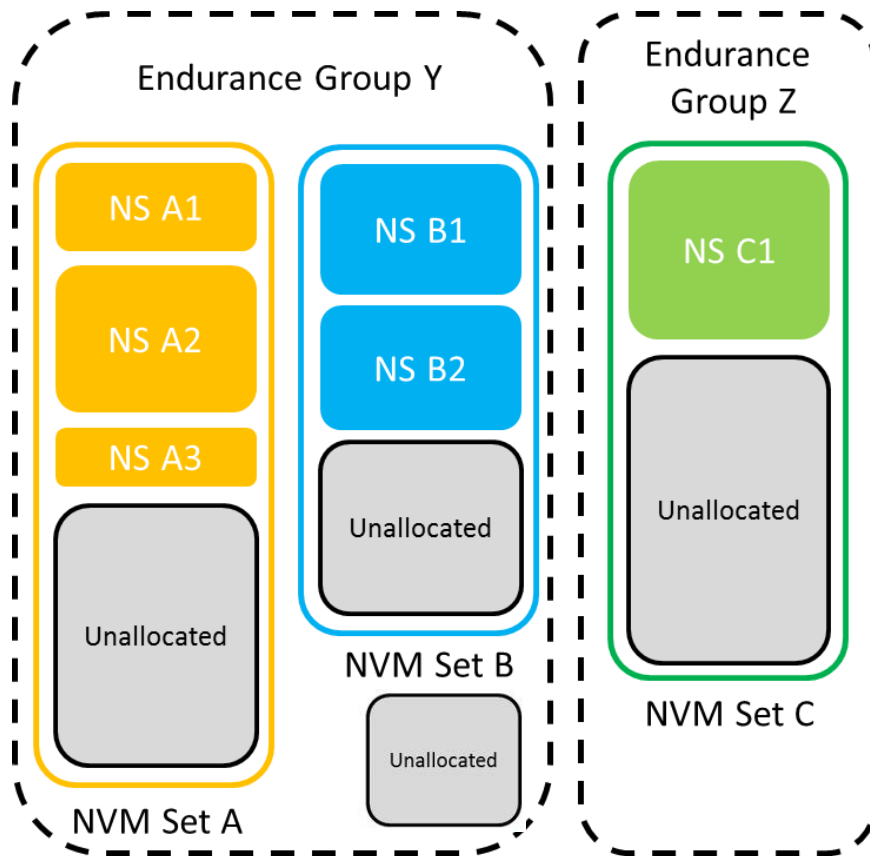
An Endurance Group shall be part of only one domain (refer to section 3.2.4).

An Endurance Group Identifier is a 16-bit value that specifies the Endurance Group with which an action is associated. An Endurance Group Identifier value of 0h is reserved and is not a valid Endurance Group Identifier. Unless otherwise specified, if the host specifies an Endurance Group Identifier cleared to 0h for a command that requires an Endurance Group Identifier, then that command shall abort with a status code of Invalid Field in Command.

The information that describes an Endurance Group is indicated in the Endurance Group Information log page (refer to section 5.16.1.10).

Figure 74 shows Endurance Groups added to the example in Figure 72. In this example, the endurance of NVM Set A and NVM Set B are managed together as part of Endurance Group Y, while the endurance of NVM Set C is managed only within NVM Set C which is the only NVM Set that is part of Endurance Group Z.

Figure 74: NVM Sets and Associated Namespaces



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;
- support the Endurance Group Information log page; and
- support the Endurance Group Event Aggregate log page **if more than one Endurance Group is supported in the NVM subsystem.**

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 shall be ignored by the controller.

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.

...

3.2.TBD Reclaim Groups, Reclaim Unit Handles, and Reclaim Units

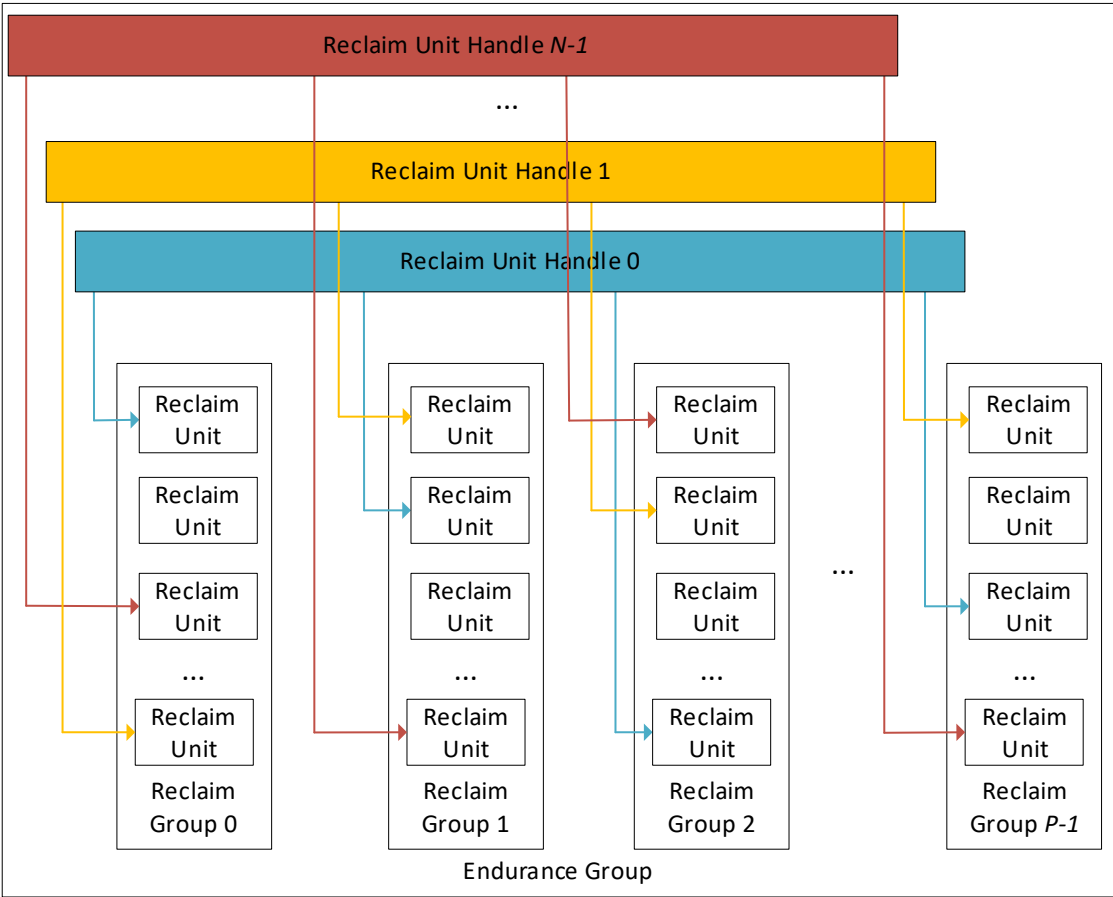
If Flexible Data Placement is enabled in an Endurance Group (refer to section 5.27.1.TBD), then the logical view of the non-volatile storage capacity in that Endurance Group is shown in Figure X and consists of:

- a set of one or more Reclaim Groups numbered from 0 to $P-1$ where P is the value of the Number of Reclaim Groups field in the FDP Configuration Descriptor (refer to Figure FDPD). A Reclaim Group consists of one or more Reclaim Units; and

- one or more Reclaim Unit Handles numbered from 0 to $N-1$ where N is the value of the Number of Reclaim Unit Handles field in the FDP Configuration Descriptor (refer to **Figure FDPD**).

A Reclaim Unit Handle consists of a reference to a Reclaim Unit in each Reclaim Group where user data from a write command is placed. A Reclaim Unit referenced by the Reclaim Unit Handle is only allowed to be referenced by at most one Reclaim Unit Handle. However, a specific Reclaim Unit is referenced by the same or different Reclaim Unit Handles as the Reclaim Unit is cycled from erased and back into use. When a Reclaim Unit is written to capacity, the controller updates that Reclaim Unit Handle to reference a different Reclaim Unit that is available for writing user data (e.g., non-volatile media that has been erased which is required prior to writing for program in place memories) and has not been written with any user data (i.e., an empty Reclaim Unit). Refer to section **8.TBD** for the details of how a host is able to issue a write command and place the user data into a Reclaim Unit.

Figure X: Flexible Data Placement Logical View of Non-Volatile Storage



...

3.3 NVM Queue Models

...

3.3.3 Queueing Data Structures

...

3.3.3.2.1 Common Completion Queue Entry

...

3.3.3.2.1 Status Field Definition

...

3.3.3.2.1.1 Generic Command Status Definition

...

Figure 94: Status Code – Generic Command Status Values

Value	Description	I/O Command Set Specific	I/O Command Set(s) ¹
...			
24h	<p>Admin Command Media Not Ready: The Admin command requires access to media and the media is not ready. The Do Not Retry bit indicates whether re-issuing the command at a later time may succeed. This status code shall only be returned:</p> <ul style="list-style-type: none"> a) for Admin commands; and b) if the controller is in Controller Ready Independent of Media mode (CC.CRIME is bit is set to '1'). <p>This status code shall not be returned with the Do Not Retry bit cleared to '0' after the amount of time indicated by the Controller Ready With Media Timeout field (CRTO.CRWMT) after the controller is enabled (i.e., CC.EN transitions from '0' to '1').</p> <p>Refer to Figure 103 for the list of Admin commands permitted to return this status code.</p>	No	
29h	FDP Disabled: The command is not allowed when Flexible Data Placement is disabled.	No	
2Ah	<p>Invalid Placement Handle List: The Placement Handle List is invalid due to:</p> <ul style="list-style-type: none"> • a Reclaim Unit Handle Identifier that is: <ul style="list-style-type: none"> ○ valid but restricted to be used by the command; or ○ invalid; or • the Placement Handle List number of entries exceeded the maximum number allowed. 	No	
2Bh to 7Fh	Reserved		
...			
<p>Key: NVM – NVM Command Set ZNS – Zoned Namespace Command Set KV – Key Value Command Set</p> <p>NOTES: 1. This column is blank unless the value is I/O Command Set specific</p>			

...

5 Admin Command Set

5.16 Get Log Page command

Figure 198: Get Log Page – Command Dword 11

Bits	Description												
31:16	<p>Log Specific Identifier: This field specifies an identifier that is required for a particular log page. If not defined for the log page specified by the Log Page Identifier field, this field is reserved.</p> <p>The log pages that require a log specific identifier are indicated in the table below.</p> <table> <tr> <th>Log Page</th><th>Identifier</th></tr> <tr> <td>Endurance Group Information</td><td>Endurance Group Identifier (refer to section 3.2.3)</td></tr> <tr> <td>Rotational Media Information</td><td>NVM Set Identifier (refer to section 3.2.2)</td></tr> <tr> <td>Predictable Latency Per NVM Set</td><td>Domain Identifier (refer to section 3.2.4)⁴</td></tr> <tr> <td>Media Unit Status</td><td>Domain Identifier (refer to section 3.2.4)⁴</td></tr> <tr> <td>Supported Capacity Configuration List</td><td>Domain Identifier (refer to section 3.2.4)⁴</td></tr> </table>	Log Page	Identifier	Endurance Group Information	Endurance Group Identifier (refer to section 3.2.3)	Rotational Media Information	NVM Set Identifier (refer to section 3.2.2)	Predictable Latency Per NVM Set	Domain Identifier (refer to section 3.2.4)⁴	Media Unit Status	Domain Identifier (refer to section 3.2.4)⁴	Supported Capacity Configuration List	Domain Identifier (refer to section 3.2.4)⁴
Log Page	Identifier												
Endurance Group Information	Endurance Group Identifier (refer to section 3.2.3)												
Rotational Media Information	NVM Set Identifier (refer to section 3.2.2)												
Predictable Latency Per NVM Set	Domain Identifier (refer to section 3.2.4)⁴												
Media Unit Status	Domain Identifier (refer to section 3.2.4)⁴												
Supported Capacity Configuration List	Domain Identifier (refer to section 3.2.4)⁴												
15:00	<p>Number of Dwords (NUMDU): This field specifies the most significant 16 bits of the number of dwords to return unless otherwise specified.</p>												
<p>NOTES:</p> <p>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 5.17.2.17), then the controller shall abort the command with Invalid Field in Command.</p>													

5.16.1 Log Specific Information

Figure 202: Get Log Page – Log Page Identifiers

Log Identifier	Scope	Log Page Name	Reference Section
...
16h	Endurance Group	Rotational Media Information	5.16.1.22
20h	Endurance Group	FDP Configurations	5.16.1.TBD _{Log}
21h	Endurance Group	Reclaim Unit Handle Usage	5.16.1.TBD _{Log1}
22h	Endurance Group	FDP Statistics	5.16.1.TBD _{Log2}
23h	Endurance Group	FDP Events	5.16.1.TBD _{Log3}
47h 24h to 6Fh	Reserved

5.16.1.10 Endurance Group Information (Log Identifier 09h)

This log page is used to provide endurance information based on the Endurance Group (refer to section 3.2.3). An Endurance Group contains capacity that may be allocated to zero or more NVM Sets. Capacity that has not been allocated to an NVM Set is unallocated Endurance Group capacity. The information

provided is over the life of the Endurance Group. The Endurance Group Identifier is specified in the Log Specific Identifier field in Command Dword 11 of the Get Log Page command as defined in [Figure FDPLog09](#). The log page is 512 bytes in size.

Figure FDPLog09: Endurance Group Identifier – Log Specific Identifier

Bits	Description
15:00	Endurance Group Identifier (ENDGID): This field specifies the identifier for the Endurance Group (refer to section 3.2.3) used for this log page.

...

5.16.1.11 Predictable Latency Per NVM Set (Log Identifier 0Ah)

This log page may be used to determine the current window for the specified NVM Set when Predictable Latency Mode is enabled and any events that have occurred for the specified NVM Set. There is one log page for each NVM Set when Predictable Latency Mode is supported. ~~Command Dword 11 (refer to Figure 198) specifies the~~ The NVM Set for which the log page is to be returned is specified in the Log Specific Identifier field in Command Dword 11 of the Get Log Page command as defined in [Figure FDPLog0A](#). The log page is 512 bytes in size.

Figure FDPLog0A: NVM Set Identifier – Log Specific Identifier

Bits	Description
15:00	NVM Set Identifier (NVMSETID): This field specifies the identifier for the NVM Set (refer to section 3.2.2) used for this log page.

...

5.16.1.16 Media Unit Status (Log Identifier 10h)

This log page is used to describe the configuration and wear of Media Units (refer to section 8.3). 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 249) 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 198~~ [Figure FDPLog10](#)), if accessible. If the information is not accessible, then the log page is not available (refer to section 8.1.4). 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 249) 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 8.3.

Figure FDPLog10: Domain Identifier – Log Specific Identifier

Bits	Description
15:00	<p>Domain Identifier (DID): This field specifies the identifier for the domain (refer to section 3.2.4) used for this log page.</p> <p>If the NVM subsystem does not support multiple domains, then this field is reserved.</p> <p>If this field specifies a non-zero Domain Identifier that is not reported in the Domain List (refer to section 5.17.2.17), then the controller shall abort the command with a status code of Invalid Field in Command.</p>

...

5.16.1.17 Supported Capacity Configuration List (Log Identifier 11h)

This log page is used to provide a list of Supported Capacity Configuration Descriptors (refer to Figure 250). Each entry in the list defines a different configuration of Endurance Groups supported by the specified domain specified in the Log Specific Identifier field in Command Dword 11 of the Get Log Page command as defined in Figure FDPLog10.

...

5.16.1.22 Rotational Media Information Log (Log Identifier 16h)

This log page provides rotational media information (refer to section 8.20) for Endurance Groups that store data on rotational media. The information provided is retained across power cycles and resets.

The Endurance Group Identifier is specified in the Log Specific Identifier field in Command Dword 11 of the Get Log Page command as defined in Figure FDPLog09.

...

5.16.1.TBD_{Log} Flexible Data Placement (FDP) Configurations (Log Page Identifier 20h)

The FDP Configurations log page (refer to Figure FDPConfig) identifies a list of static FDP configurations that are allowed to be applied to the specified Endurance Group. The Endurance Group Identifier in the Log Specific Identifier field as defined in Figure FDPLog09 specifies the Endurance Group. The creation of an Endurance Group or the enablement of an FDP configuration in an Endurance Group within the NVM subsystem may affect which FDP configurations are currently allowed (refer to the FDP Configuration Valid bit in Figure FDPD). Refer to section 8.TBD for the usage of this log page by the host.

Figure FDPConfig: FDP Configurations Log Page

Bytes	Description
Header	
01:00	Number of FDP Configurations (N): This field indicates the number of FDP Configuration Descriptors contained in this log page. This is a 0's based number.
02	Version: This field indicates the version of the log page that includes the header and the FDP Configuration Descriptor List. This field shall be cleared to 0h.
03	Reserved
07:04	Size: This field identifies the size of this log page in bytes.
15:08	Reserved
FDP Configuration Descriptor List	
Variable Sized: 16	FDP Configuration Descriptor 0 (FDPD 0): The FDP Configuration Descriptor (refer to Figure FDPD) of the first configuration.
Variable Sized	FDP Configuration Descriptor 1 (FDPD 1): The FDP Configuration Descriptor (refer to Figure FDPD) of the second configuration, if any.

Figure FDPConfig: FDP Configurations Log Page

Bytes	Description
...	...
Variable Sized	FDP Configuration Descriptor N-1 (FDPD N-1): The FDP Configuration Descriptor (refer to Figure FDPD) of the last configuration, if any.

[Figure FDPD](#) defines the format of an FDP Configuration Descriptor. The Reclaim Unit Handles in the Reclaim Unit Handle list shall be listed in ascending order of Reclaim Unit Handle Identifier.

Figure FDPD: FDP Configuration Descriptor

Bytes	Description										
01:00	Descriptor Size: This field indicates the size in bytes of this FDP Configuration Descriptor. FDP Attributes: This field identifies attributes of this FDP configuration.										
02	<table><tr><th>Bits</th><th>Description</th></tr><tr><td>7</td><td>FDP Configuration Valid: This bit is set to '1' to indicate that this FDP configuration is valid. This bit is cleared to '0' to indicate that this entry is not currently available and the host should ignore this FDP configuration.</td></tr><tr><td>6:5</td><td>Reserved</td></tr><tr><td>4</td><td>FDP Volatile Write Cache (FDPVWC): If this bit is set to '1', then a volatile write cache is present for this FDP configuration. If this bit is cleared to '0', then a volatile write cache is not present for this FDP configuration. Refer to section 5.27.1.4. If the controller reports one or more FDP configurations with this bit set to '1', then the Volatile Write Cache Present (VWCP) bit in the Identify Controller data structure (refer to Figure 275) shall be set to '1'. This results in the VWCP bit being set to '1' even though there is no volatile write cache in the controller to support hosts not aware of the Flexible Data Placement capability (refer to section 8.TBD). Refer to section 7.1 to determine if a volatile write cache is present if a namespace exists in an Endurance Group that has Flexible Data Placement enabled.</td></tr><tr><td>3:0</td><td>Reclaim Group Identifier Format (RGIF): This field identifies the number of most-significant bits in the Placement Identifier that contain the Reclaim Group Identifier (refer to Figure PIDRG and Figure PIDRG_NORGI). If the NRG field in this data structure is set to 1h (i.e., there is a single Reclaim Group), then this field may be cleared to 0h. If the NRG field is greater than 1h, then this field shall be a non-zero value.</td></tr></table>	Bits	Description	7	FDP Configuration Valid: This bit is set to '1' to indicate that this FDP configuration is valid. This bit is cleared to '0' to indicate that this entry is not currently available and the host should ignore this FDP configuration.	6:5	Reserved	4	FDP Volatile Write Cache (FDPVWC): If this bit is set to '1', then a volatile write cache is present for this FDP configuration. If this bit is cleared to '0', then a volatile write cache is not present for this FDP configuration. Refer to section 5.27.1.4. If the controller reports one or more FDP configurations with this bit set to '1', then the Volatile Write Cache Present (VWCP) bit in the Identify Controller data structure (refer to Figure 275) shall be set to '1'. This results in the VWCP bit being set to '1' even though there is no volatile write cache in the controller to support hosts not aware of the Flexible Data Placement capability (refer to section 8.TBD). Refer to section 7.1 to determine if a volatile write cache is present if a namespace exists in an Endurance Group that has Flexible Data Placement enabled.	3:0	Reclaim Group Identifier Format (RGIF): This field identifies the number of most-significant bits in the Placement Identifier that contain the Reclaim Group Identifier (refer to Figure PIDRG and Figure PIDRG_NORGI). If the NRG field in this data structure is set to 1h (i.e., there is a single Reclaim Group), then this field may be cleared to 0h. If the NRG field is greater than 1h, then this field shall be a non-zero value.
	Bits	Description									
	7	FDP Configuration Valid: This bit is set to '1' to indicate that this FDP configuration is valid. This bit is cleared to '0' to indicate that this entry is not currently available and the host should ignore this FDP configuration.									
	6:5	Reserved									
	4	FDP Volatile Write Cache (FDPVWC): If this bit is set to '1', then a volatile write cache is present for this FDP configuration. If this bit is cleared to '0', then a volatile write cache is not present for this FDP configuration. Refer to section 5.27.1.4. If the controller reports one or more FDP configurations with this bit set to '1', then the Volatile Write Cache Present (VWCP) bit in the Identify Controller data structure (refer to Figure 275) shall be set to '1'. This results in the VWCP bit being set to '1' even though there is no volatile write cache in the controller to support hosts not aware of the Flexible Data Placement capability (refer to section 8.TBD). Refer to section 7.1 to determine if a volatile write cache is present if a namespace exists in an Endurance Group that has Flexible Data Placement enabled.									
3:0	Reclaim Group Identifier Format (RGIF): This field identifies the number of most-significant bits in the Placement Identifier that contain the Reclaim Group Identifier (refer to Figure PIDRG and Figure PIDRG_NORGI). If the NRG field in this data structure is set to 1h (i.e., there is a single Reclaim Group), then this field may be cleared to 0h. If the NRG field is greater than 1h, then this field shall be a non-zero value.										
03	Vendor Specific Size (VSS): This field identifies the size in bytes of the Vendor Specific field.										
07:04	Number of Reclaim Groups (NRG): This field indicates the number of Reclaim Groups in this FDP configuration. This field shall be a non-zero value.										
09:08	Number of Reclaim Unit Handles (NRUH): This field identifies the number of Reclaim Unit Handle Descriptors in the Reclaim Unit Handle List. This field shall be a non-zero value.										
11:10	Max Placement Identifiers (MAXPIDS): This field indicates the maximum value allowed in the Number of Placement Identifiers (NPID) field (refer to Figure PIU) in an I/O Management Send command for the Reclaim Unit Handle Update Operation (refer to section 7.TBDFDPS.1.1) for this FDP configuration. This is a 0's based number. The value of this field shall be less than the product of the NRG field and the NRUH field.										
15:12	Number of Namespaces Supported: This field indicates the number of namespaces allowed to be created in the Endurance Group if Flexible Data Placement is enabled utilizing this FDP configuration. If the MNAN field is set to a value greater than 0h (refer to Figure 275), then this field shall be less than or equal to the value in the MNAN field. If the MNAN field is cleared to 0h, then this field shall be less than or equal to the value in the NN field (refer to Figure 275).										

Figure FDPD: FDP Configuration Descriptor

Bytes	Description
23:16	Reclaim Unit Nominal Size (RUNS): This field indicates the nominal size, in bytes, of each Reclaim Unit of non-volatile storage within each Reclaim Group.
27:24	Estimated Reclaim Unit Time Limit (ERUTL): This field indicates the estimated maximum time in seconds that a Reclaim Unit is allowed to be referenced by a Reclaim Unit Handle (refer to section 3.2.TBD and section 8.TBD) before the controller is allowed to modify the Reclaim Unit Handle to reference a different Reclaim Unit. If this field is cleared to 0h, then the Estimated Reclaim Unit Time Limit is not reported.
63:28	Reserved
Reclaim Unit Handle Descriptor List	
67:64	Reclaim Unit Handle Descriptor 0 (RUHD 0): This field contains the Reclaim Unit Handle Descriptor (refer to Figure FDP RUHD) for Reclaim Unit Handle 0 (refer to Figure X).
71:68	Reclaim Unit Handle Descriptor 1 (RUHD 1): This field contains the Reclaim Unit Handle Descriptor (refer to Figure FDP RUHD) for Reclaim Unit Handle 1 (refer to Figure X), if any.
...	
$((NRUH-1)*4)+67$: $((NRUH-1)*4)+64$	Reclaim Unit Handle Descriptor NRUH-1 (RUHD NRUH-1): This field contains the Reclaim Unit Handle Descriptor (refer to Figure FDP RUHD) for Reclaim Unit Handle NRUH-1 (refer to Reclaim Unit Handle $N-1$ in Figure X), if any.
$(NRUH*4)+63$:VSS: $(NRUH*4)+64$	Vendor Specific. This field is valid if the VSS field is greater than 0h.
$(NRUH*4)+64$:VSS	Padding to align the size of this descriptor to the next 8-byte boundary, if necessary. This field shall be cleared to 0h.

Figure FDP RUHD: Reclaim Unit Handle Descriptor

Bytes	Description											
00	Reclaim Unit Handle Type (RUHT): This field indicates the Reclaim Unit Handle type of the Reclaim Unit Handle.											
	Value	Description	0h	Reserved	1h	Initially Isolated: The user data from a write command that utilizes this type of Reclaim Unit Handle is originally isolated in the referenced Reclaim Unit from other user data in write commands that utilize a different Reclaim Unit Handle in the same Reclaim Group. If the controller moves the user data due to vendor specific operations (i.e., garbage collection), then the controller is allowed to move that user data to a Reclaim Unit in the same Reclaim Group that contains other user data moved by the controller that was written by the host utilizing any Reclaim Unit Handle of the same type. Refer to section 8.TBD.	2h	Persistently Isolated: The user data from a write command that utilizes this type of Reclaim Unit Handle is originally isolated in the referenced Reclaim Unit from other user data in write commands that utilize a different Reclaim Unit Handle. If the controller moves the user data due to vendor specific operations (i.e., garbage collection), then the controller shall move that user data to a Reclaim Unit in the same Reclaim Group that only contains user data that was written by the host utilizing the same Reclaim Unit Handle. Refer to section 8.TBD.	3h to BFh	Reserved	C0h to FFh	Vendor Specific
	Value	Description										
	0h	Reserved										
	1h	Initially Isolated: The user data from a write command that utilizes this type of Reclaim Unit Handle is originally isolated in the referenced Reclaim Unit from other user data in write commands that utilize a different Reclaim Unit Handle in the same Reclaim Group. If the controller moves the user data due to vendor specific operations (i.e., garbage collection), then the controller is allowed to move that user data to a Reclaim Unit in the same Reclaim Group that contains other user data moved by the controller that was written by the host utilizing any Reclaim Unit Handle of the same type. Refer to section 8.TBD.										
	2h	Persistently Isolated: The user data from a write command that utilizes this type of Reclaim Unit Handle is originally isolated in the referenced Reclaim Unit from other user data in write commands that utilize a different Reclaim Unit Handle. If the controller moves the user data due to vendor specific operations (i.e., garbage collection), then the controller shall move that user data to a Reclaim Unit in the same Reclaim Group that only contains user data that was written by the host utilizing the same Reclaim Unit Handle. Refer to section 8.TBD.										
	3h to BFh	Reserved										
C0h to FFh	Vendor Specific											
03:01	Reserved											

The format of the Placement Identifier used in the Data Placement Directive (refer to section 8.7.TBD) for a specific FDP configuration is dependent on the number of Reclaim Groups supported by the FDP

configuration in use and the number of bits allocated for the Reclaim Group Identifier field. If the NRG field (refer to [Figure FDPD](#)) is set to 1h indicating that there is only one Reclaim Group and zero bits are allocated to the Reclaim Group Identifier (i.e., the RGIF field (refer to [Figure FDPD](#)) is cleared to 0h), then the format of the Placement Identifier is defined in [Figure PIDRG_NORGI](#). If one or more bits are allocated to the Reclaim Group Identifier (i.e., the RGIF field is non-zero), then the format of the Placement Identifier is defined in [Figure PIDRG](#).

Figure PIDRG_NORGI: Placement Identifier Format without Reclaim Group Identifier

Bits	Description
15:00	Placement Handle (PHNDL): This field specifies a Placement Handle associated with the namespace that is used by the controller to determine the Reclaim Unit Handle.

Figure PIDRG: Placement Identifier Format with a non-zero RGIF

Bits	Description
15:(16-RGIF)	Reclaim Group Identifier (RGID): This field specifies the Reclaim Group used by the write command. If the NRG field (refer to Figure FDPD) is set to 1h (i.e., there is only one Reclaim Group), then this field shall be ignored by the controller.
(15-RGIF):00	Placement Handle (PHNDL): This field specifies a Placement Handle associated with the namespace that is used by the controller to determine the Reclaim Unit Handle.

5.16.1.TBD_{Log1} Reclaim Unit Handle Usage (Log Page Identifier 21h)

The Reclaim Unit Handle Usage log page (refer to [Figure FigLogRUHU](#)) is used to provide information about the Reclaim Unit Handles associated with the Placement Handles of the namespaces in the specified Endurance Group. The Endurance Group is specified by the Endurance Group Identifier field in the Log Specific Identifier field as defined in [Figure FDPLog09](#).

If Flexible Data Placement is enabled in the Endurance Group specified by the Endurance Group Identifier field and a Get Log Page command specifies the Reclaim Unit Handle Usage log page, then the NSID field in that command is reserved.

If Flexible Data Placement is disabled in the specified Endurance Group and the Get Log Page command specifying this log page, then the controller shall abort the command with a status code of FDP Disabled.

Figure FigLogRUHU: Reclaim Unit Handle Usage Log Page

Bytes	Description
01:00	Number of Reclaim Unit Handles (NRUH): This field identifies the number of Reclaim Unit Handle Usage Descriptors in the Reclaim Unit Handle Usage Descriptor List. This field shall be a non-zero value. This field shall be set to the value of the NRUH field in the current FDP configuration (refer to Figure FDPD) for the specified Endurance Group.
07:02	Reserved
Reclaim Unit Handle Usage Descriptor List	
15:08	Reclaim Unit Handle Usage Descriptor 0 (RUHUD 0): Contains the Reclaim Unit Handle Usage Descriptor (refer to Figure FigLogRUHUD) for Reclaim Unit Handle 0.

Figure FigLogRUHU: Reclaim Unit Handle Usage Log Page

Bytes	Description
23:16	Reclaim Unit Handle Usage Descriptor 1 (RUHUD 1): Contains the Reclaim Unit Handle Usage Descriptor (refer to Figure FigLogRUHUD) for Reclaim Unit Handle 1.
...	
(NRUH-1)*8+15: (NRUH-1)*8+8	Reclaim Unit Handle Usage Descriptor NRUH-1 (RUHUD NRUH-1): Contains the Reclaim Unit Handle Usage Descriptor (refer to Figure FigLogRUHUD) for Reclaim Unit Handle NRUH-1.

The Reclaim Unit Handle Usage Descriptor is defined in [Figure FigLogRUHUD](#).

Figure FigLogRUHUD: Reclaim Unit Handle Usage Descriptor

Bytes	Description										
00	Reclaim Unit Handle Attributes: This field identifies attributes associated with the Reclaim Unit Handle Usage.										
	<table><tr><th>Value</th><th>Description</th></tr><tr><td>0h</td><td>Not used by a namespace.</td></tr><tr><td>1h</td><td>Host Specified: The Reclaim Unit Handle Identifier is used by a namespace created with a Namespace Management command explicitly requesting the usage of this Reclaim Unit Handle (i.e., specifying a Placement Handle List).</td></tr><tr><td>2h</td><td>Controller Specified: The Reclaim Unit Handle Identifier is used by a namespace created with a Namespace Management command requesting the controller to select the one and only Reclaim Unit Handle (i.e., not specifying a Placement Handle List).</td></tr><tr><td>3h to FFh</td><td>Reserved</td></tr></table>	Value	Description	0h	Not used by a namespace.	1h	Host Specified: The Reclaim Unit Handle Identifier is used by a namespace created with a Namespace Management command explicitly requesting the usage of this Reclaim Unit Handle (i.e., specifying a Placement Handle List).	2h	Controller Specified: The Reclaim Unit Handle Identifier is used by a namespace created with a Namespace Management command requesting the controller to select the one and only Reclaim Unit Handle (i.e., not specifying a Placement Handle List).	3h to FFh	Reserved
	Value	Description									
	0h	Not used by a namespace.									
	1h	Host Specified: The Reclaim Unit Handle Identifier is used by a namespace created with a Namespace Management command explicitly requesting the usage of this Reclaim Unit Handle (i.e., specifying a Placement Handle List).									
2h	Controller Specified: The Reclaim Unit Handle Identifier is used by a namespace created with a Namespace Management command requesting the controller to select the one and only Reclaim Unit Handle (i.e., not specifying a Placement Handle List).										
3h to FFh	Reserved										
07:01	Reserved										

In the list of Reclaim Unit Handle Usage Descriptors, no more than one entry shall have the Reclaim Unit Handle Attributes field set to Controller Specified (i.e., 2h).

5.16.1.TBD_{Log2} Flexible Data Placement (FDP) Statistics (Log Page Identifier 22h)

The FDP Statistics log page (refer to [Figure FigLogDS](#)) is used to provide information about the FDP configuration over the life of the FDP configuration in an Endurance Group. The Endurance Group is specified by the Endurance Group Identifier field in the Log Specific Identifier field as defined in [Figure FDPLog09](#).

If Flexible Data Placement is enabled in the specified Endurance Group and the Get Log Page command specifies this log page, then the NSID field in that command is reserved.

If Flexible Data Placement is disabled in the specified Endurance Group and the Get Log Page command specifies this log page, then the controller shall abort the command with a status code of FDP Disabled.

Figure FigLogDS: FDP Statistics Log Page

Bytes	Description
15:00	<p>Host Bytes with Metadata Written (HBMW): Contains the total number of bytes of user data the host has written to the specified Endurance Group as part of processing I/O commands as defined in the appropriate I/O Command Set specification. This value does not include controller writes due to internal operations such as garbage collection.</p> <p>This field shall not wrap once the value FFFFFFFF_FFFFFFFF_FFFFFFFF_FFFFFFFFh is reached.</p> <p>If the Flexible Data Placement Feature value is modified by a Set Features command, then this field shall be cleared to 0h.</p>
31:16	<p>Media Bytes with Metadata Written (MBMW): Contains the total number of bytes of user data that have been written to the specified Endurance Group including both host and controller writes (e.g., garbage collection and background media scan operations) as part of processing as defined in the appropriate I/O Command Set specification.</p> <p>This field shall not wrap once the value FFFFFFFF_FFFFFFFF_FFFFFFFF_FFFFFFFFh is reached.</p> <p>If the Flexible Data Placement Feature value is modified by a Set Features command, then this field shall be cleared to 0h.</p>
47:32	<p>Media Bytes Erased (MBE): Contains the total number of bytes of data that have been erased in the specified Endurance Group.</p> <p>This field shall not wrap once the value FFFFFFFF_FFFFFFFF_FFFFFFFF_FFFFFFFFh is reached.</p> <p>If the Flexible Data Placement Feature value is modified by a Set Features command, then this field shall be cleared to 0h.</p>
63:48	Reserved

The values reported in this log page are not cleared to 0h by a firmware update.

The values reported in this log page include operations on all namespaces that existed in the specified Endurance Group since the Flexible Data Placement was last enabled in that Endurance Group (refer to section 5.27.1.TBD).

5.16.1.TBD_{Log3} Flexible Data Placement (FDP) Events (Log Page Identifier 23h)

The FDP Events log page is used to provide information about events affecting Reclaim Units and media usage in an Endurance Group that has Flexible Data Placement enabled. The log page is 4 KiB bytes in size. The Endurance Group is specified by the Endurance Group Identifier in the Log Specific Identifier field as defined in Figure FDPLog09.

FDP events are associated with Reclaim Unit Handles. An FDP event shall only be reported if the FDP event occurs and that FDP event is enabled for that Reclaim Unit Handle (refer to section 5.27.1.TBD1).

If:

- an FDP event occurs;
- that event is enabled; and
- the number of events recorded is the maximum supported,

then the oldest event shall be discarded.

The FDP Event Type bit (refer to Figure FigLogCDW10) specifies whether the controller reports host events or controller events in the log page. The controller shall not report both host events and controller events in the log page.

If Flexible Data Placement is enabled in the specified Endurance Group and a Get Log Page command specifies this log page, then the NSID field in that command is reserved.

If Flexible Data Placement is disabled in the specified Endurance Group, then a Get Log Page command that specifies this log page shall be aborted with a status code of FDP Disabled.

Figure FigLogCDW10: Command Dword 10 – Log Specific Field

Bits	Description
14:09	Reserved
08	FDP Event Type (FDPET): This bit specifies the type of events to be reported in the log page. If this bit is set to '1', then the controller shall report host events. If this bit is cleared to '0', then the controller shall report controller events.

Figure FigLogFDPE: FDP Events Log Page

Bytes	Description
03:00	Number of FDP Events (N): This field indicates the number of FDP events (refer to Figure FigLogDED) contained in this log. If Flexible Data Placement transitions from disabled to enabled in an Endurance Group, then this field shall be cleared to 0h for that Endurance Group.
63:04	Reserved
FDP Event List	
127:64	FDP Event 1: The oldest FDP event, if any.
...	...
(N*64)+63:(N*64)	FDP Event N: The most recent FDP event, if any.

The FDP events shall be listed in ascending order of occurrence. Because the Timestamp feature may or may not have been set by the host after each Controller Level Reset and because the Event Timestamp field is the time at which the event was created, the FDP events may or may not be in order of ascending value of the Event Timestamp field.

The number of FDP events reported in the log page is equal to the value of the Number of FDP Events field. The values of the bytes in this log page which follow the FDP Event List is implementation specific.

The format of the FDP event is described in [Figure FigLogDED](#).

Figure FigLogDED: FDP Event

Bytes	Description																																																																																																	
00	Event Type: This field indicates the type of event:																																																																																																	
	<table><tr><th>Value</th><th>O/M¹</th><th>ETS²</th><th>Description</th><th>Reference</th></tr><tr><td colspan="5">Host Events</td></tr><tr><td>0h</td><td>O</td><td>No</td><td>Reclaim Unit Not Fully Written To Capacity</td><td>5.16.1.TBD_{Log3.1.1}</td></tr><tr><td>1h</td><td>O</td><td>No</td><td>Reclaim Unit Time Limit Exceeded</td><td>5.16.1.TBD_{Log3.1.2}</td></tr><tr><td>2h</td><td>O</td><td>No</td><td>Controller Level Reset Modified Reclaim Unit Handles</td><td>5.16.1.TBD_{Log3.1.3}</td></tr><tr><td>3h</td><td>O</td><td>No</td><td>Invalid Placement Identifier</td><td>5.16.1.TBD_{Log3.1.4}</td></tr><tr><td>4h to 6Fh</td><td></td><td colspan="3">Reserved</td></tr><tr><td>70h to 7Fh</td><td>O</td><td></td><td colspan="2">Vendor specific</td></tr><tr><td colspan="5">Controller Events</td></tr><tr><td>80h</td><td>O</td><td>Yes</td><td>Media Reallocated</td><td>5.16.1.TBD_{Log3.2.1}</td></tr><tr><td>81h</td><td>O</td><td>No</td><td>Implicitly Modified Reclaim Unit Handle</td><td>5.16.1.TBD_{Log3.2.2}</td></tr><tr><td>82h to EFh</td><td></td><td colspan="3">Reserved</td></tr><tr><td>F0h to FFh</td><td>O</td><td></td><td colspan="2">Vendor specific</td></tr><tr><td colspan="5">Note: 1. O/M definition: O = Optional, M = Mandatory. 2. Event Type Specific field: Yes = the field is utilized. No = the field is not utilized.</td></tr><tr><td rowspan="5">01</td><td>FDP Event Flags (FDPEF):</td></tr><tr><td><table><tr><th>Bits</th><th>Description</th></tr><tr><td>07:03</td><td>Reserved</td></tr><tr><td>02</td><td>Location Valid (LV): If this bit is set to '1', then the Reclaim Group Identifier field and the Reclaim Unit Handle Identifier field contain a reported value. If this bit is cleared to '0', then the Reclaim Group Identifier field and the Reclaim Unit Handle Identifier field do not contain a reported values.</td></tr><tr><td>01</td><td>NSID Valid (NSIDV): If this bit is set to '1', the NSID field contains a reported value. If this bit is cleared to '0', then the NSID field does not contain a reported value.</td></tr><tr><td>00</td><td>Placement Identifier Valid (PIV): If this bit is set to '1', then the Placement Identifier field contains a reported value. If this bit is cleared to '0', then the Placement Identifier field does not contain a reported value.</td></tr></table></td></tr><tr><td>03:02</td><td>Placement Identifier (PID): This field indicates the Placement Identifier associated with the FDP event. This field is reserved if the PIV bit is cleared to '0'. If the FDP event is an Invalid Placement Identifier event (i.e., Event Type 03h), then this field is Placement Identifier from the write command submitted by the host.</td></tr><tr><td>11:04</td><td>Event Timestamp: This field indicates a timestamp indicating the current value of the Timestamp feature (refer to section 5.27.1.11) when this FDP event occurred. The format of this field is as defined in Figure 340).</td></tr><tr><td>15:12</td><td>Namespace Identifier (NSID): This field indicates the identifier of the namespace associated with the FDP event. If the NSIDV bit is cleared to '0', then this field shall be cleared to 0h and the host should ignore this field.</td></tr><tr><td>31:16</td><td>Event Type Specific: Each Event Type indicates the use of this field and the format of this field.</td></tr><tr><td>33:32</td><td>Reclaim Group Identifier: This field indicates the Reclaim Group associated with the event. If the LV bit is cleared to '0', then this field shall be cleared to 0h and the host should ignore this field. If the FDP event is an Invalid Placement Identifier event (i.e., Event Type 03h), then this field indicates the Reclaim Group selected by the controller.</td></tr><tr><td>35:34</td><td>Reclaim Unit Handle Identifier: This field indicates the Reclaim Unit Handle associated with the FDP event. If the LV bit is cleared to '0', then this field shall be cleared to 0h and the host should ignore this field. If the FDP event is an Invalid Placement Identifier event (i.e., Event Type 03h), then this field indicates the Reclaim Unit Handle selected by the controller.</td></tr><tr><td>39:36</td><td>Reserved</td></tr></table>	Value	O/M ¹	ETS ²	Description	Reference	Host Events					0h	O	No	Reclaim Unit Not Fully Written To Capacity	5.16.1.TBD _{Log3.1.1}	1h	O	No	Reclaim Unit Time Limit Exceeded	5.16.1.TBD _{Log3.1.2}	2h	O	No	Controller Level Reset Modified Reclaim Unit Handles	5.16.1.TBD _{Log3.1.3}	3h	O	No	Invalid Placement Identifier	5.16.1.TBD _{Log3.1.4}	4h to 6Fh		Reserved			70h to 7Fh	O		Vendor specific		Controller Events					80h	O	Yes	Media Reallocated	5.16.1.TBD _{Log3.2.1}	81h	O	No	Implicitly Modified Reclaim Unit Handle	5.16.1.TBD _{Log3.2.2}	82h to EFh		Reserved			F0h to FFh	O		Vendor specific		Note: 1. O/M definition: O = Optional, M = Mandatory. 2. Event Type Specific field: Yes = the field is utilized. No = the field is not utilized.					01	FDP Event Flags (FDPEF):	<table><tr><th>Bits</th><th>Description</th></tr><tr><td>07:03</td><td>Reserved</td></tr><tr><td>02</td><td>Location Valid (LV): If this bit is set to '1', then the Reclaim Group Identifier field and the Reclaim Unit Handle Identifier field contain a reported value. If this bit is cleared to '0', then the Reclaim Group Identifier field and the Reclaim Unit Handle Identifier field do not contain a reported values.</td></tr><tr><td>01</td><td>NSID Valid (NSIDV): If this bit is set to '1', the NSID field contains a reported value. If this bit is cleared to '0', then the NSID field does not contain a reported value.</td></tr><tr><td>00</td><td>Placement Identifier Valid (PIV): If this bit is set to '1', then the Placement Identifier field contains a reported value. If this bit is cleared to '0', then the Placement Identifier field does not contain a reported value.</td></tr></table>	Bits	Description	07:03	Reserved	02	Location Valid (LV): If this bit is set to '1', then the Reclaim Group Identifier field and the Reclaim Unit Handle Identifier field contain a reported value. If this bit is cleared to '0', then the Reclaim Group Identifier field and the Reclaim Unit Handle Identifier field do not contain a reported values.	01	NSID Valid (NSIDV): If this bit is set to '1', the NSID field contains a reported value. If this bit is cleared to '0', then the NSID field does not contain a reported value.	00	Placement Identifier Valid (PIV): If this bit is set to '1', then the Placement Identifier field contains a reported value. If this bit is cleared to '0', then the Placement Identifier field does not contain a reported value.	03:02	Placement Identifier (PID): This field indicates the Placement Identifier associated with the FDP event. This field is reserved if the PIV bit is cleared to '0'. If the FDP event is an Invalid Placement Identifier event (i.e., Event Type 03h), then this field is Placement Identifier from the write command submitted by the host.	11:04	Event Timestamp: This field indicates a timestamp indicating the current value of the Timestamp feature (refer to section 5.27.1.11) when this FDP event occurred. The format of this field is as defined in Figure 340).	15:12	Namespace Identifier (NSID): This field indicates the identifier of the namespace associated with the FDP event. If the NSIDV bit is cleared to '0', then this field shall be cleared to 0h and the host should ignore this field.	31:16	Event Type Specific: Each Event Type indicates the use of this field and the format of this field.	33:32	Reclaim Group Identifier: This field indicates the Reclaim Group associated with the event. If the LV bit is cleared to '0', then this field shall be cleared to 0h and the host should ignore this field. If the FDP event is an Invalid Placement Identifier event (i.e., Event Type 03h), then this field indicates the Reclaim Group selected by the controller.	35:34	Reclaim Unit Handle Identifier: This field indicates the Reclaim Unit Handle associated with the FDP event. If the LV bit is cleared to '0', then this field shall be cleared to 0h and the host should ignore this field. If the FDP event is an Invalid Placement Identifier event (i.e., Event Type 03h), then this field indicates the Reclaim Unit Handle selected by the controller.	39:36	Reserved
	Value	O/M ¹	ETS ²	Description	Reference																																																																																													
	Host Events																																																																																																	
	0h	O	No	Reclaim Unit Not Fully Written To Capacity	5.16.1.TBD _{Log3.1.1}																																																																																													
	1h	O	No	Reclaim Unit Time Limit Exceeded	5.16.1.TBD _{Log3.1.2}																																																																																													
	2h	O	No	Controller Level Reset Modified Reclaim Unit Handles	5.16.1.TBD _{Log3.1.3}																																																																																													
	3h	O	No	Invalid Placement Identifier	5.16.1.TBD _{Log3.1.4}																																																																																													
	4h to 6Fh		Reserved																																																																																															
	70h to 7Fh	O		Vendor specific																																																																																														
	Controller Events																																																																																																	
	80h	O	Yes	Media Reallocated	5.16.1.TBD _{Log3.2.1}																																																																																													
	81h	O	No	Implicitly Modified Reclaim Unit Handle	5.16.1.TBD _{Log3.2.2}																																																																																													
	82h to EFh		Reserved																																																																																															
	F0h to FFh	O		Vendor specific																																																																																														
Note: 1. O/M definition: O = Optional, M = Mandatory. 2. Event Type Specific field: Yes = the field is utilized. No = the field is not utilized.																																																																																																		
01	FDP Event Flags (FDPEF):																																																																																																	
	<table><tr><th>Bits</th><th>Description</th></tr><tr><td>07:03</td><td>Reserved</td></tr><tr><td>02</td><td>Location Valid (LV): If this bit is set to '1', then the Reclaim Group Identifier field and the Reclaim Unit Handle Identifier field contain a reported value. If this bit is cleared to '0', then the Reclaim Group Identifier field and the Reclaim Unit Handle Identifier field do not contain a reported values.</td></tr><tr><td>01</td><td>NSID Valid (NSIDV): If this bit is set to '1', the NSID field contains a reported value. If this bit is cleared to '0', then the NSID field does not contain a reported value.</td></tr><tr><td>00</td><td>Placement Identifier Valid (PIV): If this bit is set to '1', then the Placement Identifier field contains a reported value. If this bit is cleared to '0', then the Placement Identifier field does not contain a reported value.</td></tr></table>	Bits	Description	07:03	Reserved	02	Location Valid (LV): If this bit is set to '1', then the Reclaim Group Identifier field and the Reclaim Unit Handle Identifier field contain a reported value. If this bit is cleared to '0', then the Reclaim Group Identifier field and the Reclaim Unit Handle Identifier field do not contain a reported values.	01	NSID Valid (NSIDV): If this bit is set to '1', the NSID field contains a reported value. If this bit is cleared to '0', then the NSID field does not contain a reported value.	00	Placement Identifier Valid (PIV): If this bit is set to '1', then the Placement Identifier field contains a reported value. If this bit is cleared to '0', then the Placement Identifier field does not contain a reported value.																																																																																							
	Bits	Description																																																																																																
	07:03	Reserved																																																																																																
	02	Location Valid (LV): If this bit is set to '1', then the Reclaim Group Identifier field and the Reclaim Unit Handle Identifier field contain a reported value. If this bit is cleared to '0', then the Reclaim Group Identifier field and the Reclaim Unit Handle Identifier field do not contain a reported values.																																																																																																
01	NSID Valid (NSIDV): If this bit is set to '1', the NSID field contains a reported value. If this bit is cleared to '0', then the NSID field does not contain a reported value.																																																																																																	
00	Placement Identifier Valid (PIV): If this bit is set to '1', then the Placement Identifier field contains a reported value. If this bit is cleared to '0', then the Placement Identifier field does not contain a reported value.																																																																																																	
03:02	Placement Identifier (PID): This field indicates the Placement Identifier associated with the FDP event. This field is reserved if the PIV bit is cleared to '0'. If the FDP event is an Invalid Placement Identifier event (i.e., Event Type 03h), then this field is Placement Identifier from the write command submitted by the host.																																																																																																	
11:04	Event Timestamp: This field indicates a timestamp indicating the current value of the Timestamp feature (refer to section 5.27.1.11) when this FDP event occurred. The format of this field is as defined in Figure 340).																																																																																																	
15:12	Namespace Identifier (NSID): This field indicates the identifier of the namespace associated with the FDP event. If the NSIDV bit is cleared to '0', then this field shall be cleared to 0h and the host should ignore this field.																																																																																																	
31:16	Event Type Specific: Each Event Type indicates the use of this field and the format of this field.																																																																																																	
33:32	Reclaim Group Identifier: This field indicates the Reclaim Group associated with the event. If the LV bit is cleared to '0', then this field shall be cleared to 0h and the host should ignore this field. If the FDP event is an Invalid Placement Identifier event (i.e., Event Type 03h), then this field indicates the Reclaim Group selected by the controller.																																																																																																	
35:34	Reclaim Unit Handle Identifier: This field indicates the Reclaim Unit Handle associated with the FDP event. If the LV bit is cleared to '0', then this field shall be cleared to 0h and the host should ignore this field. If the FDP event is an Invalid Placement Identifier event (i.e., Event Type 03h), then this field indicates the Reclaim Unit Handle selected by the controller.																																																																																																	
39:36	Reserved																																																																																																	

Figure FigLogDED: FDP Event

Bytes	Description
Vendor Specific Information	
63:40	Vendor specific. This field may be used by any Event Type (i.e., this field is not only for Vendor specific Event Types).

If the controller is not able to report the namespace identifier associated with the event, then the NSIDV bit shall be cleared to '0'.

If the controller is not able to report the specific Reclaim Unit Handle, then:

- the PIV bit shall be cleared to '0'; and
- the LV bit shall be cleared to '0'.

5.16.1.TBD_{Log3}.1 Host Events

5.16.1.TBD_{Log3}.1.1 Reclaim Unit Not Fully Written To Capacity (Event Type 0h)

This FDP event indicates that the Reclaim Unit Handle reported in the Placement Identifier (refer to [Figure PIDRG](#) and [Figure PIDRG_NORGI](#)) was modified to reference a different Reclaim Unit due to an I/O Management Send command issued by the host prior to the previously referenced Reclaim Unit being written to capacity.

5.16.1.TBD_{Log3}.1.2 Reclaim Unit Time Limit Exceeded (Event Type 1h)

This FDP event indicates that the capacity of a Reclaim Unit was not completely written within the time period defined in the Estimated Reclaim Unit Time Limit field (refer to [Figure FDPD](#)) due to the controller modifying the Reclaim Unit Handle to reference a different Reclaim Unit.

5.16.1.Log_{TBD3}.1.3 Controller Level Reset Modified Reclaim Unit Handles (Event Type 2h)

This FDP event indicates that one or more Reclaim Unit Handles were modified to reference a different Reclaim Unit due to a Controller Level Reset.

In the FDP event, the NSID field is reserved and the Placement Identifier field is reserved.

5.16.1.Log_{TBD3}.1.4 Invalid Placement Identifier (Event Type 3h)

This FDP event indicates that a write command specified an invalid Placement Identifier (refer to [Figure PIDRG](#) and [Figure PIDRG_NORGI](#)) due to an invalid Placement Handle or an invalid Reclaim Group Identifier.

5.16.1.TBD_{Log3}.2 Controller Events

5.16.1.TBD_{Log3}.2.1 Media Reallocated (Event Type 80h)

This FDP event indicates that user data contained in a Reclaim Unit was written by the host specifying a Reclaim Unit Handle with a Reclaim Unit Handle type of Initially Isolated and was moved by the controller to a different Reclaim Unit (e.g., controller performed garbage collection).

The Event Type Specific field is I/O Command Set specific. Refer to the applicable I/O Command Set specification to determine if this log page is supported and the definition of the Event Type Specific field for this FDP event.

5.16.1.TBD_{Log3}.2.2 Implicitly Modified Reclaim Unit Handle (Event Type 81h)

This FDP event indicates that the controller modified a Reclaim Unit Handle to reference a different Reclaim Unit where that modification was not due to any host action (e.g., a write command that has a size greater than the remaining capacity of the initial Reclaim Unit being written).

...

5.17 Identify command

...

5.17.2 Identify Data Structures**5.17.2.1 Identify Controller Data Structure (CNS 01h)**

...

Figure 275: Identify – Identify Controller Data Structure, I/O Command Set Independent

Bytes	I/O ¹	Admin ¹	Disc ¹	Description
...				

Figure 275: Identify – Identify Controller Data Structure, I/O Command Set Independent

Bytes	I/O ¹	Admin ¹	Disc ¹	Description																				
99:96	M	M	R	Controller Attributes (CTRATT): This field indicates attributes of the controller.																				
				<table><tr><th>Bits</th><th>Description</th></tr><tr><td colspan="2">...</td></tr><tr><td>19</td><td>Flexible Data Placement Support (FDPS): If this bit is set to '1', then the controller supports the Flexible Data Placement capability (refer to section 8.TBD). If this bit is cleared to '0', then the controller does not support the Flexible Data Placement capability. If Fixed Capacity Management is supported (i.e., the FCM bit is set to '1'), then this bit shall be cleared to '0'.</td></tr><tr><td>15</td><td>Extended LBA Formats Supported (ELBAS): If set to '1' indicates that the controller supports the I/O command set specific extended protection information formats (refer to the Protection Information Formats section of the applicable I/O command set specification). If cleared to '0' indicates that the controller does not support the I/O command set specific extended protection information formats (refer to the Protection Information Formats section of the NVM Command Set Specification). Refer to the LBA Format Extension Enable (LBAFEE) field in the Host Behavior Support feature (refer to section 5.27.1.18) for details for host software to enable the controller to operate on namespaces using the protection information formats. NOTE: This bit field applies to all I/O Command Sets. The original name has been retained for historical continuity.</td></tr><tr><td colspan="2">...</td></tr><tr><td>12</td><td>Variable Capacity Management (VCM): If set to '1', then the controller supports Variable Capacity Management (refer to section 8.3.3). If cleared to '0', then the controller does not support Variable Capacity Management.</td></tr><tr><td>11</td><td>Fixed Capacity Management (FCM): If set to '1', then the controller supports Fixed Capacity Management (refer to section 8.3.2). If cleared to '0', then the controller does not support Fixed Capacity Management. If the Flexible Data Placement capability is supported (i.e., the FDPS bit is set to '1'), then this bit shall be cleared to '0'.</td></tr><tr><td colspan="2">...</td></tr><tr><td>2</td><td>NVM Sets: If set to '1', then the controller supports NVM Sets (refer to section 3.2.2). If cleared to '0', then the controller does not support NVM Sets.</td></tr><tr><td colspan="2">...</td></tr></table>	Bits	Description	...		19	Flexible Data Placement Support (FDPS): If this bit is set to '1', then the controller supports the Flexible Data Placement capability (refer to section 8.TBD). If this bit is cleared to '0', then the controller does not support the Flexible Data Placement capability. If Fixed Capacity Management is supported (i.e., the FCM bit is set to '1'), then this bit shall be cleared to '0'.	15	Extended LBA Formats Supported (ELBAS): If set to '1' indicates that the controller supports the I/O command set specific extended protection information formats (refer to the Protection Information Formats section of the applicable I/O command set specification). If cleared to '0' indicates that the controller does not support the I/O command set specific extended protection information formats (refer to the Protection Information Formats section of the NVM Command Set Specification). Refer to the LBA Format Extension Enable (LBAFEE) field in the Host Behavior Support feature (refer to section 5.27.1.18) for details for host software to enable the controller to operate on namespaces using the protection information formats. NOTE: This bit field applies to all I/O Command Sets. The original name has been retained for historical continuity.	...		12	Variable Capacity Management (VCM): If set to '1', then the controller supports Variable Capacity Management (refer to section 8.3.3). If cleared to '0', then the controller does not support Variable Capacity Management.	11	Fixed Capacity Management (FCM): If set to '1', then the controller supports Fixed Capacity Management (refer to section 8.3.2). If cleared to '0', then the controller does not support Fixed Capacity Management. If the Flexible Data Placement capability is supported (i.e., the FDPS bit is set to '1'), then this bit shall be cleared to '0'.	...		2	NVM Sets: If set to '1', then the controller supports NVM Sets (refer to section 3.2.2). If cleared to '0', then the controller does not support NVM Sets.	...	
				Bits	Description																			
				...																				
				19	Flexible Data Placement Support (FDPS): If this bit is set to '1', then the controller supports the Flexible Data Placement capability (refer to section 8.TBD). If this bit is cleared to '0', then the controller does not support the Flexible Data Placement capability. If Fixed Capacity Management is supported (i.e., the FCM bit is set to '1'), then this bit shall be cleared to '0'.																			
				15	Extended LBA Formats Supported (ELBAS): If set to '1' indicates that the controller supports the I/O command set specific extended protection information formats (refer to the Protection Information Formats section of the applicable I/O command set specification). If cleared to '0' indicates that the controller does not support the I/O command set specific extended protection information formats (refer to the Protection Information Formats section of the NVM Command Set Specification). Refer to the LBA Format Extension Enable (LBAFEE) field in the Host Behavior Support feature (refer to section 5.27.1.18) for details for host software to enable the controller to operate on namespaces using the protection information formats. NOTE: This bit field applies to all I/O Command Sets. The original name has been retained for historical continuity.																			
				...																				
				12	Variable Capacity Management (VCM): If set to '1', then the controller supports Variable Capacity Management (refer to section 8.3.3). If cleared to '0', then the controller does not support Variable Capacity Management.																			
				11	Fixed Capacity Management (FCM): If set to '1', then the controller supports Fixed Capacity Management (refer to section 8.3.2). If cleared to '0', then the controller does not support Fixed Capacity Management. If the Flexible Data Placement capability is supported (i.e., the FDPS bit is set to '1'), then this bit shall be cleared to '0'.																			
				...																				
2	NVM Sets: If set to '1', then the controller supports NVM Sets (refer to section 3.2.2). If cleared to '0', then the controller does not support NVM Sets.																							
...																								

525	M	M	R	<p>Volatile Write Cache (VWC): This field indicates attributes related to the presence of a volatile write cache in the controller.</p> <p>Bits 7:3 are reserved.</p> <p>Bits 2:1 indicate Flush command behavior (refer to the Flush Command section of the NVM Command Set Specification7.1) if the NSID value is set to FFFFFFFFh as follows:</p> <table><tr><th>Value</th><th>Definition</th></tr><tr><td>00b</td><td>Support for the NSID field set to FFFFFFFFh is not indicated. Only controllers compliant with NVM Express Base Specification revision 1.3 and earlier shall be allowed to return this value.</td></tr><tr><td>01b</td><td>Reserved.</td></tr><tr><td>10b</td><td>The Flush command does not support the NSID field set to FFFFFFFFh. The controller shall fail a Flush command with the NSID set to FFFFFFFFh with a status code of Invalid Namespace or Format.</td></tr><tr><td>11b</td><td>The Flush command supports the NSID field set to FFFFFFFFh.</td></tr></table> <p>Bit 0 if set to '1' indicates that a volatile write cache is present. If cleared to '0', a volatile write cache is not present.</p> <p>If a volatile write cache is present, then the host controls whether the volatile write cache is enabled with a Set Features command specifying the Volatile Write Cache feature identifier (refer to section 5.27.1.4). The Flush command (refer to the Flush Command section of the NVM Command Set Specification) is used to request that the contents of a volatile write cache be made non-volatile.</p> <table><tr><th>Bits</th><th>Description</th></tr><tr><td>07:03</td><td>Reserved</td></tr><tr><td rowspan="5">02:01</td><td>Flush Behavior: This field indicates Flush command behavior (refer to the Flush Command section of the NVM Command Set Specification7.1) if the NSID value is set to FFFFFFFFh as follows:</td></tr><tr><td><table><tr><th>Value</th><th>Definition</th></tr><tr><td>00b</td><td>Support for the NSID field set to FFFFFFFFh is not indicated. Only controllers compliant with NVM Express Base Specification revision 1.3 and earlier shall be allowed to return this value.</td></tr><tr><td>01b</td><td>Reserved.</td></tr><tr><td>10b</td><td>The Flush command does not support the NSID field set to FFFFFFFFh. The controller shall fail a Flush command with the NSID set to FFFFFFFFh with a status code of Invalid Namespace or Format.</td></tr><tr><td>11b</td><td>The Flush command supports the NSID field set to FFFFFFFFh.</td></tr></table></td></tr><tr><td rowspan="4">00</td><td>Volatile Write Cache Present: This bit is set to '1' to indicate that a volatile write cache is present in the controller. If This bit is cleared to '0', to indicate that volatile write cache is not present in the controller.</td></tr><tr><td>If a volatile write cache is present, then the host controls whether the volatile write cache is enabled with a Set Features command specifying the Volatile Write Cache feature identifier (refer to section 5.27.1.4). The Flush command (refer to the Flush Command section of the NVM Command Set Specification) is used to request that the contents of a volatile write cache be made non-volatile.</td></tr><tr><td></td></tr><tr><td></td></tr></table>	Value	Definition	00b	Support for the NSID field set to FFFFFFFFh is not indicated. Only controllers compliant with NVM Express Base Specification revision 1.3 and earlier shall be allowed to return this value.	01b	Reserved.	10b	The Flush command does not support the NSID field set to FFFFFFFFh. The controller shall fail a Flush command with the NSID set to FFFFFFFFh with a status code of Invalid Namespace or Format.	11b	The Flush command supports the NSID field set to FFFFFFFFh.	Bits	Description	07:03	Reserved	02:01	Flush Behavior: This field indicates Flush command behavior (refer to the Flush Command section of the NVM Command Set Specification7.1) if the NSID value is set to FFFFFFFFh as follows:	<table><tr><th>Value</th><th>Definition</th></tr><tr><td>00b</td><td>Support for the NSID field set to FFFFFFFFh is not indicated. Only controllers compliant with NVM Express Base Specification revision 1.3 and earlier shall be allowed to return this value.</td></tr><tr><td>01b</td><td>Reserved.</td></tr><tr><td>10b</td><td>The Flush command does not support the NSID field set to FFFFFFFFh. The controller shall fail a Flush command with the NSID set to FFFFFFFFh with a status code of Invalid Namespace or Format.</td></tr><tr><td>11b</td><td>The Flush command supports the NSID field set to FFFFFFFFh.</td></tr></table>	Value	Definition	00b	Support for the NSID field set to FFFFFFFFh is not indicated. Only controllers compliant with NVM Express Base Specification revision 1.3 and earlier shall be allowed to return this value.	01b	Reserved.	10b	The Flush command does not support the NSID field set to FFFFFFFFh. The controller shall fail a Flush command with the NSID set to FFFFFFFFh with a status code of Invalid Namespace or Format.	11b	The Flush command supports the NSID field set to FFFFFFFFh.	00	Volatile Write Cache Present: This bit is set to '1' to indicate that a volatile write cache is present in the controller. If This bit is cleared to '0', to indicate that volatile write cache is not present in the controller.	If a volatile write cache is present, then the host controls whether the volatile write cache is enabled with a Set Features command specifying the Volatile Write Cache feature identifier (refer to section 5.27.1.4). The Flush command (refer to the Flush Command section of the NVM Command Set Specification) is used to request that the contents of a volatile write cache be made non-volatile.		
Value	Definition																																			
00b	Support for the NSID field set to FFFFFFFFh is not indicated. Only controllers compliant with NVM Express Base Specification revision 1.3 and earlier shall be allowed to return this value.																																			
01b	Reserved.																																			
10b	The Flush command does not support the NSID field set to FFFFFFFFh. The controller shall fail a Flush command with the NSID set to FFFFFFFFh with a status code of Invalid Namespace or Format.																																			
11b	The Flush command supports the NSID field set to FFFFFFFFh.																																			
Bits	Description																																			
07:03	Reserved																																			
02:01	Flush Behavior: This field indicates Flush command behavior (refer to the Flush Command section of the NVM Command Set Specification7.1) if the NSID value is set to FFFFFFFFh as follows:																																			
	<table><tr><th>Value</th><th>Definition</th></tr><tr><td>00b</td><td>Support for the NSID field set to FFFFFFFFh is not indicated. Only controllers compliant with NVM Express Base Specification revision 1.3 and earlier shall be allowed to return this value.</td></tr><tr><td>01b</td><td>Reserved.</td></tr><tr><td>10b</td><td>The Flush command does not support the NSID field set to FFFFFFFFh. The controller shall fail a Flush command with the NSID set to FFFFFFFFh with a status code of Invalid Namespace or Format.</td></tr><tr><td>11b</td><td>The Flush command supports the NSID field set to FFFFFFFFh.</td></tr></table>	Value	Definition	00b	Support for the NSID field set to FFFFFFFFh is not indicated. Only controllers compliant with NVM Express Base Specification revision 1.3 and earlier shall be allowed to return this value.	01b	Reserved.	10b	The Flush command does not support the NSID field set to FFFFFFFFh. The controller shall fail a Flush command with the NSID set to FFFFFFFFh with a status code of Invalid Namespace or Format.	11b	The Flush command supports the NSID field set to FFFFFFFFh.																									
	Value	Definition																																		
	00b	Support for the NSID field set to FFFFFFFFh is not indicated. Only controllers compliant with NVM Express Base Specification revision 1.3 and earlier shall be allowed to return this value.																																		
	01b	Reserved.																																		
10b	The Flush command does not support the NSID field set to FFFFFFFFh. The controller shall fail a Flush command with the NSID set to FFFFFFFFh with a status code of Invalid Namespace or Format.																																			
11b	The Flush command supports the NSID field set to FFFFFFFFh.																																			
00	Volatile Write Cache Present: This bit is set to '1' to indicate that a volatile write cache is present in the controller. If This bit is cleared to '0', to indicate that volatile write cache is not present in the controller.																																			
	If a volatile write cache is present, then the host controls whether the volatile write cache is enabled with a Set Features command specifying the Volatile Write Cache feature identifier (refer to section 5.27.1.4). The Flush command (refer to the Flush Command section of the NVM Command Set Specification) is used to request that the contents of a volatile write cache be made non-volatile.																																			

5.17.2.8 I/O Command Set Independent Identify Namespace data structure (CNS 08h)

< **Note to Editor:** Added the Scope column from TP4095a>

Figure 280: Identify – I/O Command Set Independent Identify Namespace Data Structure

Bytes	O/M ¹	Description	Reported ²				
...							
00	M	Common Namespace Features (NSFEAT): This field defines features of the namespace. Bits 7:5 are reserved. Bit 4 Rotational Media (RMEDIA) if set to '1' indicates that the namespace stores data on rotational media (refer to section 8.20). If cleared to '0', indicates that the namespace does not store data on rotational media. Bit 3 (UIDREUSE) if set to '1' indicates that the value in the NGUID field for this namespace, if non-zero, is never reused by the controller and that the value in the EUI64 field for this namespace, if non-zero, is never reused by the controller. If cleared to '0', then the NGUID value may be reused and the EUI64 value may be reused by the controller for a new namespace created after this namespace is deleted. This bit shall be cleared to '0' if both NGUID and EUI64 fields are cleared to 0h. Refer to section 4.5.1. Bit 2:0 are reserved.	No				
		<table><tr><th>Bits</th><th>Description</th></tr><tr><td>7:56</td><td>Reserved</td></tr></table>		Bits	Description	7:5 6	Reserved
		Bits		Description			
		7:5 6		Reserved			
		5		Volatile Write Cache Not Present (VWCNP): This bit indicates that a volatile write cache is not present for this namespace. If this bit is set to '1', then there is no volatile write cache present for this namespace and the host should ignore the Volatile Write Cache Present bit in the Identify Controller data structure (refer to Figure 275) for this namespace. If this bit is cleared to '0', then the presence of a volatile write cache for this namespace is defined in the Volatile Write Cache bit in the Identify Controller data structure (refer to Figure 275).			
		4		Rotational Media (RMEDIA): If this bit is set to '1', then the namespace stores data on rotational media (refer to section 8.20). If this bit is cleared to '0', then the namespace does not store data on rotational media.			
		3		UID Reuse (UIDREUSE): If this bit is set to '1', then the value in the NGUID field for this namespace, if non-zero, is never reused by the controller and that the value in the EUI64 field for this namespace, if non-zero, is never reused by the controller. If this bit is cleared to '0', then the NGUID value may be reused and the EUI64 value may be reused by the controller for a new namespace created after this namespace is deleted. This bit shall be cleared to '0' if both NGUID and EUI64 fields are cleared to 0h. Refer to section 4.5.1.			
		2:0		Reserved			

5.27 Set Features command

...

5.27.1 Feature Specific Information

...

<The scope column in Figure 316 is from TP4074a>

Figure 316: Set Features – Feature Identifiers

Feature Identifier	Current Setting Persists Across Power Cycle and Reset ²	Uses Memory Buffer for Attributes	Feature Name	Scope
...				
1Ah	Yes	No	Spinup Control	NVM subsystem
1Dh	Yes	No	Flexible Data Placement	Endurance Group
1Eh	Yes	Yes	FDP Events	Reclaim Unit Handle
...				
NOTES:				
<ol style="list-style-type: none"> 1. The behavior of a controller in response to an inactive namespace ID to a vendor specific Feature Identifier is vendor specific. 2. This column is only valid if the feature is not saveable (refer to section 4.2). If the feature is saveable, then this column is not used. 3. The controller does not save settings for the Host Memory Buffer feature across power states and reset events, however, host software may restore the previous values. Refer to section 8.9. 4. The feature does not use a memory buffer for Set Features commands and does use a memory buffer for Get Features commands. Refer to section 8.9. 5. Selection of a UUID may be supported. Refer to section 8.25. 				

...

5.27.1.4 Volatile Write Cache (Feature Identifier 06h), (Optional)

This Feature controls the volatile write cache, if present, on the controller. If a volatile write cache is present (refer to the VWC field in Figure 275), then this feature shall be supported. The attributes are specified in Command Dword 11.

Note: If the controller is able to guarantee that data present in a write cache is written to non-volatile media on loss of power, then that write cache is considered non-volatile and this feature does not apply to that write cache.

If a Get Features command is submitted for this Feature, the attributes specified in Figure 321 are returned in Dword 0 of the completion queue entry for that command.

If a volatile write cache is not present, then a Set Features command specifying the Volatile Write Cache feature identifier shall abort with a status code of Invalid Field in Command, and a Get Features command specifying the Volatile Write Cache feature identifier should abort with a status code of Invalid Field in Command.

If a volatile write cache is present and the volatile write cache is disabled (i.e., the WCE bit is cleared to '0'), then the user data written by any command to a namespace shall be persistent. Refer to the Volatile Write Cache Presence field contained in the Common Namespace Features field in [Figure 280](#).

Figure 321: Volatile Write Cache – Command Dword 11

Bits	Description
31:01	Reserved
00	Volatile Write Cache Enable (WCE): If set to '1', then the volatile write cache is enabled. If cleared to '0', then the volatile write cache is disabled.

...

5.27.1.TBD Flexible Data Placement (Feature Identifier 1Dh)

This Feature controls operation of the Flexible Data Placement capability (refer to section [8.TBD](#)) in the specified Endurance Group (refer to [Figure FigFeatSI](#)).

The attribute is specified in Command Dword 12. Effects of enabling and disabling the Flexible Data Placement are described in section [8.TBD](#).

If a Get Features command specifying this Feature is successfully completed, then the attributes described in [Figure FigFeat](#) are returned in Dword 0 of the completion queue entry for that command.

Figure FigFeatSI: Flexible Data Placement – Command Dword 11

Bits	Description
31:16	Reserved
15:00	Endurance Group Identifier (ENDGID): If Endurance Groups are supported, then this field specifies the identifier for the Endurance Group (refer to section 3.2.3) used for this Feature.

Figure FigFeat: Flexible Data Placement – Command Dword 12

Bits	Description
31:16	Reserved
15:08	Flexible Data Placement Configuration Index: This field specifies the index into the FDP Configuration Descriptor List (refer to Figure FDPConfig) identifying a valid FDP configuration to be applied to the Endurance Group when Flexible Data Placement is enabled (refer to section 8.TBD).
07:01	Reserved
00	Flexible Data Placement Enable (FDPE): This bit specifies if the Flexible Data Placement (i.e., capability) is enabled or disabled. If this bit is set to '1', then Flexible Data Placement is enabled. If this bit is cleared to '0', then Flexible Data Placement is disabled.

This Feature shall be saveable (refer to section [4.2](#)). The default value of this Feature shall be 0h.

The value of this Feature is only allowed to change if the SV bit is set to '1' in the Set Features command. Therefore, if bit 4 is set to '1' in the ONCS field in the Identify Controller data structure (refer to [Figure 275](#)) and the SV bit is cleared to '0', then the controller shall abort the command with a status code of Invalid Field in Command.

If the value of this Feature changes, then the controller is allowed to modify any fields associated with specifying any User Data Format information accessible using the Identify command (i.e., any field associated with the Format Index (refer to section 1.6.2) that includes CNS values 00h, 05h, 08h, <the following are from TP4095a> 09h, 0Ah). The host should issue Identify commands in this condition to access the values that may have been modified.

...

5.27.1.TBD1 Flexible Data Placement Events (Feature Identifier 1Eh)

This Feature controls if a controller generates Flexible Data Placement (FDP) events associated with a specific Reclaim Unit Handle. The Reclaim Unit Handle is the Reclaim Unit Handle associated with the specified Placement Handle for the specified namespace-specified. For a Set Features command, the data buffer contains the list of FDP Event Types that are to be enabled or disabled as defined in Figure FeatDS.

The attribute is specified in Command Dword 12. Effects of enabling and disabling FDP events on the Reclaim Unit Handles are described in section 8.TBD.

If Flexible Data Placement is disabled, then the controller shall abort any Set Feature command or Get Feature command specifying this Feature with a status code of FDP Disabled.

If the NSID field is set to FFFFFFFFh, then the controller shall abort the command with a status code of Invalid Field in Command.

If a Get Features command specifying this Feature is successfully completed, then the attributes for the supported list of FDP Event Types described in Figure FeatGF is returned in the data buffer and Dword 0 of the completion queue entry contains the value of NOET (i.e., the number of supported FDP Event Types that are contained in the data buffer). Completion queue entry Dword 0 is defined in Figure FeatCQ. The supported list of FDP Event Types shall be listed in ascending order of FDP Event Type.

Figure FeatCQ: FDP Events – Completion Queue Entry Dword 0

Bits	Description
31:08	Reserved
07:00	Number of FDP Event Types (NOET): This field specifies the number of FDP Event Types that are contained in the data buffer (refer to Figure FeatDS).

Figure FeatFDPE: FDP Events – Command Dword 11

Bits	Description
31:24	Reserved
23:16	Number of FDP Event Types (NOET): This field specifies the number of FDP Event Types that are contained in the data buffer (refer to Figure FeatDS). This field is ignored by the controller on a Get Features command.
15:00	Placement Handle (PHNDL): This field specifies the Placement Handle associated with the Reclaim Unit Handle affected by this command. If the specified Placement Handle is not valid for the namespace, then the controller shall abort a Set Feature command or a Get Feature command with a status code of Invalid Field in Command.

Figure FeatEn: FDP Events – Command Dword 12

Bits	Description
31:01	Reserved
00	FDP Event Enable: If this bit is set to '1' then FDP events are enabled on the Reclaim Unit Handle associated with the Placement Handle. If this bit is cleared to '0', then FDP events are disabled on the Reclaim Unit Handle associated with the Placement Handle. This field is ignored by the controller on a Get Features command.

This Feature shall be saveable (refer to section 4.2).

If:

- a Reclaim Unit Handle is shared by more than one namespace that exists in the same Endurance Group; and
- a Set Features command specifies this Feature and one of those namespaces,

then a modification to this Feature occurs to that shared Reclaim Unit Handle and the Placement Handle that is associated with the shared Reclaim Unit Handle.

Figure FeatDS: FDP Events – Set Feature Data Structure

Bytes	Description
00	FDP Event Type 0: This field contains the first FDP Event Type (refer to Figure FigLogDED).
01	FDP Event Type 1: This field contains the second FDP Event Type.
...	
NOET-1	FDP Event Type NOET-1: This field contains the last FDP Event Type.

Figure FeatGF: FDP Events – Get Feature Data Structure

Bytes	Description
01:00	Supported FDP Event Type 0: This field contains the Supported FDP Event Descriptor for the first Supported FDP Event Type.
03:02	Supported FDP Event Type 1: This field contains the Supported FDP Event Descriptor for the second Supported FDP Event Type.
...	
(NOET-1)*2+1: (NOET-1)*2	Supported FDP Event Type NOET-1: This field contains the Supported FDP Event Descriptor for the last Supported FDP Event Type.

Figure SFDPED: Supported FDP Event Descriptor

Bytes	Description	
00	FDP Event Type: This field contains the FDP Event Type (refer to Figure FigLogDED).	
01	FDP Event Type Attributes: This field contains attributes for the FDP Event Type.	
	Bits	Description
	07:01	Reserved
	00	FDP Event Enable: If this bit is set to '1' then this FDP Event Type is enabled. If this bit is cleared to '0', then this FDP Event Type is disabled.

7 I/O Commands

Figure 390: Opcodes for I/O Commands

Opcode by Field			Combined Opcode ¹	Command ²	Reference
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer ³			
0b	000 00b	00b	00h	Flush ⁴	7.1
0b	000 11b	01b	0Dh	Reservation Register	7.3
0b	000 11b	10b	0Eh	Reservation Report	7.5
0b	001 00b	01b	11h	Reservation Acquire	7.2
0b	001 01b	01b	15h	Reservation Release	7.4
0b	001 11b	01b	1Dh	I/O Management Send	7.TBD _{IOMS}
0b	001 00b	10b	12h	I/O Management Receive	7.TBD _{IOMR}
Vendor Specific					
1b	n/a	NOTE 3	80h to FFh	Vendor specific	

NOTES:

1. Opcodes not listed are I/O Command Set specific or reserved.
2. All I/O commands use the Namespace Identifier (NSID) field. The value FFFFFFFFh is not supported in this field unless footnote 4 in this figure indicates that a specific command does support that value.
3. Indicates the data transfer direction of the command. All options to the command shall transfer data as specified or transfer no data. All commands, including vendor specific commands, shall follow this convention: 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional.
4. This command may support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh.

7.1 Flush command

The Flush command is used to request that the contents of volatile write cache be made non-volatile.

If a volatile write cache is enabled (refer to section 5.27.1.4), then the Flush command shall commit data and metadata associated with the specified namespace(s) to non-volatile media. The flush applies to all commands for the specified namespace(s) completed by the controller prior to the submission of the Flush command. The controller may also flush additional data and/or metadata from any namespace.

If bits 2:1 are set to 11b in the VWC field (refer to Figure 275) and the specified NSID is FFFFFFFFh, then the Flush command applies to all namespaces attached to the controller processing the Flush command. If bits 2:1 are set to 10b in the VWC field and the specified NSID is FFFFFFFFh, then the controller aborts the command with a status code of Invalid Namespace or Format. If bits 2:1 are cleared to 00b in the VWC field, then the controller behavior if the specified NSID is FFFFFFFFh is not indicated. Controllers compliant with NVM Express Base Specification revision 1.4 and later shall not set bits 2:1 in the VWC field to the value of 00b.

If a namespace exists in an Endurance Group that has:

- Flexible Data Placement (refer to section 8.TBD) enabled; and
- the Volatile Write Cache Not Present (VWCNP) bit set to '1' in the I/O Command Set Independent Identify Namespace data structure (refer to Figure 280),

then, even though the Volatile Write Cache field in the Identify Controller data structure (refer to [Figure 275](#)) indicates the presence of a volatile write cache in the controller, this namespace does not have a volatile write cache present.

A host may use the Volatile Write Cache Not Present (VWCNP) bit in the I/O Command Set Independent Identify Namespace data structure to determine if a Volatile Write Cache is not present for a namespace.

If a volatile write cache is not present or not enabled, then Flush commands shall have no effect and:

- a) shall complete successfully if a sanitize operation is not in progress; and
- b) may complete successfully if a sanitize operation is in progress.

All command specific fields are reserved.

7.TBD_{IOMR} I/O Management Receive command

The I/O Management Receive command is used to receive information from the controller used by the host to manage I/O. The behavior of the command is dependent on the specified operation as defined in the Management Operation field in [Figure IOMR2](#).

The command uses the Data Pointer, Command Dword 10, and Command Dword 11 fields. All other command specific fields are reserved. If the command uses PRPs for the data transfer, then the PRP Entry 1 and PRP Entry 2 fields are used. If the command uses SGLs for the data transfer, then the SGL Entry 1 field is used. All other command specific fields are reserved.

If the Number of Dwords (NUMD) field corresponds to a length that is less than the size of the data structure to be returned, then only that specified portion of the data structure is transferred. If the NUMD field corresponds to a length that is greater than the size of the associated data structure, then the entire contents of the data structure are transferred and no additional data is transferred.

Figure IOMR1: I/O Management Receive – Data Pointer

Bits	Description
127:00	Data Pointer (DPTR): This field specifies the data to use for the command. Refer to Figure 86 for the definition of this field.

Figure IOMR2: I/O Management Receive – Command Dword 10

Bits	Description										
31:16	Management Operation Specific (MOS): This definition of this field is specific the value of the MO field. If this field is not defined for the management operation specified by the MO field, then this field is reserved.										
15:08	Reserved										
07:00	Management Operation (MO): This field specifies the management operation to perform.										
	<table><tr><th>Value</th><th>Description</th></tr><tr><td>00h</td><td>No action</td></tr><tr><td>01h</td><td>Reclaim Unit Handle Status: For each Placement Handle of the namespace, the controller shall return a Reclaim Unit Handle Status Descriptor for each Reclaim Group.</td></tr><tr><td>FFh</td><td>Vendor specific</td></tr><tr><td>All others</td><td>Reserved</td></tr></table>	Value	Description	00h	No action	01h	Reclaim Unit Handle Status: For each Placement Handle of the namespace, the controller shall return a Reclaim Unit Handle Status Descriptor for each Reclaim Group.	FFh	Vendor specific	All others	Reserved
	Value	Description									
	00h	No action									
	01h	Reclaim Unit Handle Status: For each Placement Handle of the namespace, the controller shall return a Reclaim Unit Handle Status Descriptor for each Reclaim Group.									
	FFh	Vendor specific									
All others	Reserved										

Figure IOMR3: I/O Management Receive – Command Dword 11

Bits	Description
31:00	Number of Dwords (NUMD): This field specifies the number of dwords to transfer. This is a 0's based value.

7.TBD_{F DPR}.1 I/O Management Receive Operations

7.TBD_{F DPR}.1.1 Reclaim Unit Handle Status (Management Operation 01h)

The Reclaim Unit Handle Status management operation is used to provide information about Reclaim Unit Handles (refer to section 1.5.TBD_{SS}) that are accessible by the specified namespace. The returned data contains one or more Reclaim Unit Handle Status Descriptor data structures (refer to the applicable I/O Command Set specification). The information contained in each Reclaim Unit Handle Status Descriptor:

- is the information at the time the controller processes that Reclaim Unit Handle Status Descriptor; and
- may or may not contain the information reflecting any outstanding command that affects the reported Reclaim Unit Handle Status Descriptor.

If Flexible Data Placement is disabled in the Endurance Group containing the specified namespace, then the command shall be aborted with a status code of FDP Disabled.

If the NSID field is set to 0h or FFFFFFFFh, then the controller shall abort the command with a status code of Invalid Namespace or Format.

If the host reads beyond the size of the Reclaim Unit Handle Status data structure (refer to Figure FigLogRUHS), zeroes are returned.

Figure FigLogRUHS: Reclaim Unit Handle Status

Bytes	Description
13:00	Reserved
15:14	Number of Reclaim Unit Handle Status Descriptors (NRUHSD): This field indicates the number of Reclaim Unit Handle Status Descriptors in the Reclaim Unit Handle Status Descriptor list.
Reclaim Unit Handle Status Descriptor List	
47:16	Reclaim Unit Handle Status Descriptor 1 (RUHSD 1): The first Reclaim Unit Handle Status Descriptor (refer to the applicable I/O Command Set specification).
79:48	Reclaim Unit Handle Status Descriptor 2 (RUHSD 2): The second Reclaim Unit Handle Status Descriptor (refer to the applicable I/O Command Set specification), if any.
...	...
(NRUHSD *32)+15: (NRUHSD *32)–16	Reclaim Unit Handle Status Descriptor NRUHS (RUHSD NRUHD): The last Reclaim Unit Handle Status Descriptor (refer to the applicable I/O Command Set specification), if any.

The Reclaim Unit Handle Status Descriptors are defined in the I/O Command Set specifications, if supported.

7.TBD_{F DPR}.2 Command Completion

When the command is completed with success or failure, the controller shall post a completion queue entry to the associated I/O Completion Queue indicating the status for the command.

7.TBD_{IOMS} I/O Management Send command

The I/O Management Send command is used to manage I/O and the behavior of the command is dependent on the specified operation as defined in the Management Operation field in [Figure IOMS2](#).

The command uses the Data Pointer and Command Dword 10 field. All other command specific fields are reserved. If the command uses PRPs for the data transfer, then the PRP Entry 1 and PRP Entry 2 fields are used. If the command uses SGLs for the data transfer, then the SGL Entry 1 field is used. All other command specific fields are reserved.

Figure IOMS1: I/O Management Send – Data Pointer

Bits	Description
127:00	Data Pointer (DPTR): This field specifies the data to use for the command. Refer to Figure 86 for the definition of this field.

Figure IOMS2: I/O Management Send – Command Dword 10

Bits	Description										
31:16	Management Operation Specific (MOS): The definition of this field is specific to the value of the MO field. If this field is not defined for the management operation specified by the MO field, then this field is reserved.										
15:08	Reserved										
07:00	Management Operation (MO): This field specifies the management operation to perform.										
	<table><tr><th>Value</th><th>Description</th></tr><tr><td>00h</td><td>No action</td></tr><tr><td>01h</td><td>Reclaim Unit Handle Update: Update the reference to the Reclaim Unit specified by each entry in the Placement Identifier list to reference an empty Reclaim Unit (refer to section 3.2.TBD).</td></tr><tr><td>FFh</td><td>Vendor specific</td></tr><tr><td>All others</td><td>Reserved</td></tr></table>	Value	Description	00h	No action	01h	Reclaim Unit Handle Update: Update the reference to the Reclaim Unit specified by each entry in the Placement Identifier list to reference an empty Reclaim Unit (refer to section 3.2.TBD).	FFh	Vendor specific	All others	Reserved
	Value	Description									
	00h	No action									
	01h	Reclaim Unit Handle Update: Update the reference to the Reclaim Unit specified by each entry in the Placement Identifier list to reference an empty Reclaim Unit (refer to section 3.2.TBD).									
	FFh	Vendor specific									
All others	Reserved										

7.TBD_{FDPS.1} I/O Management Send Operations

7.TBD_{FDPS.1.1} Reclaim Unit Handle Update (Management Operation 01h)

The Reclaim Unit Handle Update management operation for the I/O Management Send command provides a list of Placement Identifiers (refer to [Figure PIDL](#)). The number of Placement Identifiers is defined in the Management Operation Specific field defined in [Figure PIU](#). For each Placement Identifier in the list:

- If the currently referenced Reclaim Unit has been written with user data, then the Placement Identifier shall be modified to reference a different Reclaim Unit that is empty (refer to section [3.2.TBD](#)); and
- If the currently referenced Reclaim Unit has not been written with any user data (i.e., is already empty), then the Placement Identifier may be modified to reference a different Reclaim Unit that is empty.

Figure PIU: Management Operation Specific – Reclaim Unit Handle Update Operation

Bits	Description
15:00	<p>Number of Placement Identifiers (NPID): Indicates the number of Placement Identifiers that are specified in the command. This is a 0's based value.</p> <p>This field shall not exceed the minimum of:</p> <ul style="list-style-type: none"> the value in the Max Placement Identifiers (MAXPIDS) field of the enabled FDP configuration (refer to Figure FDPD); and the product of the Number of Reclaim Groups (NRG) field and the Number of Reclaim Unit Handles (NRUH) field of the enabled FDP configuration (refer to Figure FDPD);

If a specified Placement Identifier is invalid due to:

- the value of the Reclaim Group Identifier field being greater than or equal to the Number of Reclaim Groups field of the FDP Configuration Descriptor (refer to [Figure FDPD](#)) for the Endurance Group associated with the specified namespace; or
- the specified Placement Handle field being greater than or equal to the Number of Placement Handles field specified when the namespace was created,

then the controller shall abort the command with a status code of Invalid Field in Command.

If the value represented by the Number of Placement Identifiers (NPID) field is greater than the Max Placement Identifiers (MAXPIDS) field (refer to [Figure FDPD](#)) in the current FDP configuration, then the controller shall abort the command with a status code of Invalid Field in Command.

If the command is aborted, then Placement Identifiers may or may not have been updated.

While processing an I/O Management Send command that specifies the Reclaim Unit Handle Update operation, if the controller processes a write command that utilizes a Placement Identifier specified in the Placement Identifier List of that I/O Management Send command, then the controller may write the user data for that write command to the referenced Reclaim Unit:

- prior to processing that I/O Management Send command; or
- upon the completion of that I/O Management Send command.

Figure PIDL: I/O Management Send – Placement Identifier List

Bytes	Description
01:00	Placement Identifier 1: This field specifies the first Placement Identifier that indicates a Placement Handle and a Reclaim Group Identifier. Refer to Figure PIDRG and Figure PIDRG_NORGI .
03:02	Placement Identifier 2: This field specifies the second Placement Identifier that indicates a Placement Handle and a Reclaim Group Identifier, if any. Refer to Figure PIDRG and Figure PIDRG_NORGI .
...	
(NPID*2)+1: (NPID*2)	Placement Identifier NPID: This field specifies the last Placement Identifier that indicates a Placement Handle and a Reclaim Group Identifier, if any. Refer to Figure PIDRG and Figure PIDRG_NORGI .

If Flexible Data Placement is disabled in the Endurance Group containing the specified namespace, then the controller shall abort the command with a status code of FDP Disabled.

7.TBD_{FDPs}.2 Command Completion

When the command is completed with success or failure, the controller shall post a completion queue entry to the associated I/O Completion Queue indicating the status for the command.

...

8 Extended Capabilities

...

8.7 Directives

Directives is a mechanism to enable host and NVM subsystem or controller information exchange. The Directive Receive command (refer to section 5.10) is used to transfer data related to a specific Directive Type from the controller to the host. The Directive Send command (refer to section 5.11) is used to transfer data related to a specific Directive Type from the host to the controller. Other commands may include a Directive Specific value specific for a given Directive Type (e.g., the Write command in the NVM command set).

Support for Directives is optional and is indicated in the Optional Admin Command Support (OACS) field in the Identify Controller data structure (refer to Figure 275).

If a controller supports Directives, then the controller shall:

- Indicate support for Directives in the Optional Admin Command Support (OACS) field in the Identify Controller data structure;
- Support the Directive Receive command;
- Support the Directive Send command; and
- Support the Identify Directive (i.e., Type 00h).

The Directive Types that may be supported by a controller are ~~define~~ defined in ~~(refer to Figure 416)~~ are:

- ~~the Identify Directive (refer to section 8.7.2);, and~~
- ~~the Streams Directive (refer to section 8.7.3).~~

The Directive Specific field and Directive Operation field are dependent on the Directive Type specified in the command (e.g., Directive Send, Directive Receive, or I/O command).

Figure 416: Directive Types

Directive	Directive Type Value	Definition	I/O Command Directive
Identify	00h	Section 8.7.2	No
Streams	01h	Section 8.7.3	Yes
Data Placement	02h	Section 8.7.TBD	Yes

If a Directive is not supported or is supported and disabled, then all Directive Send commands and Directive Receive commands with that Directive Type shall be aborted with a status code of Invalid Field in Command.

Support for a specific directive type is indicated using the Return Parameters operation of the Identify Directive. A specific directive may be enabled or disabled using the Enable operation of the Identify Directive. Before using a specific directive, the host should determine if that directive is supported and should enable that directive using the Identify Directive.

8.7.1 Directive Use in I/O Commands

I/O Command Directives are the subset of Directive Types that may be used as part of I/O commands. For example, a Write command in the NVM command set may specify a Directive Type and an associated Directive Specific value. I/O Command Directives shall have a Directive Type value that is less than or equal to 0Fh due to the size of the Directive Type field in I/O commands. When a Directive Type is specified in an I/O command, the most significant four bits are assumed to be 0h. A Directive Type of 00h in an I/O command specifies that the I/O command is not using Directives.

In an I/O command, if the Directive Type (DTYPE) field is set to an I/O Command Directive, then the Directive Specific (DSPEC) field includes additional information for the associated I/O command (refer to Figure 417).

Figure 417: Directive Specific Field Interpretation

Directive Type Value	Directive Specific Field Definition
00h (Directives not in use)	Field not used.
01h (Streams)	Specifies the identifier of the stream associated with the data.
02h (Data Placement)	Specifies the Placement Identifier used to determine where to write the user data within non-volatile storage (refer to Figure PIDRG and Figure PIDRG NORGI) of the Endurance Group associated with the namespace.
03h to 0Fh	Reserved

In an I/O command:

- if no I/O Command Directive is enabled or the DTYPE field is cleared to 00h, then the DTYPE field and the DSPEC field are ignored; and
- if one or more I/O Command Directives is enabled and the DTYPE field is set to a value that is not supported or not enabled, then the controller shall abort the command with a status code of Invalid Field in Command.

For the Streams Directive (i.e., DTYPE field set to 01h), if the DSPEC field is cleared to 0h in an I/O command that supports the Streams Directive, then that I/O command shall be processed normally (i.e., as if DTYPE field is cleared to 00h).

8.7.2 Identify (Directive Type 00h)

The Identify Directive is used to determine the Directive Types that the controller supports and to enable use of the supported Directives. If Directives are supported, then this Directive Type shall be supported.

The Directive operations that shall be supported for the Identify Directive are listed in Figure 418.

Figure 418: Identify Directive – Directive Operations

Directive Command	Directive Operation Name	Directive Operation Value	Definition
Directive Receive	Return Parameters	01h	Section 8.7.2.1.1
	Reserved	All other values	
Directive Send	Enable Directive	01h	Section 8.7.2.2.1
	Reserved	All other values	

8.7.2.1 Directive Receive

This section defines operations used with the Directive Receive command for the Identify Directive.

8.7.2.1.1 Return Parameters (Directive Operation 01h)

This operation returns a data structure that contains a bit vector specifying the Directive Types supported by the controller and a bit vector specifying the Directive Types enabled for the namespace. The data structure returned is defined in Figure 419. If an NSID value of FFFFFFFFh is specified, then the controller shall abort the command with a status code of Invalid Field in Command. The DSPEC field in command Dword 11 is not used for this operation.

Figure 419: Identify Directive – Return Parameters Data Structure

Bytes	Bits	Description
		Directives Supported
31:00	255:0203	Reserved
	02	Data Placement Directive: This bit is set to '1' if the Data Placement Directive is supported. This bit is cleared to '0' if the Data Placement Directive is not supported.
	01	Streams Directive: This bit is set to '1' if the Streams Directive is supported. This bit is cleared to '0' if the Streams Directive is not supported.
	00	Identify Directive: This bit shall be set to '1' to indicate that the Identify Directive is supported.
		Directives Enabled
63:32	255:0203	Reserved
	02	Data Placement Directive: This bit is set to '1' if the Data Placement Directive is enabled. This bit is cleared to '0' if the Data Placement Directive is not enabled.
	01	Streams Directive: This bit is set to '1' if the Streams Directive is enabled. This bit is cleared to '0' if the Streams Directive is not enabled.
	00	Identify Directive: This bit shall be set to '1' to indicate that the Identify Directive is enabled.
		Directive Persistent Across Controller Level Resets
95:64	255:03	Reserved
	02	Data Placement Directive: If the Data Placement Directive is supported, then this bit shall be set to '1' to indicate that the host specified Data Placement Directive state is preserved across Controller Level Resets. If the Data Placement Directive is not supported, then this bit shall be cleared to '0'.
	01	Streams Directive: This bit shall be cleared to '0' to indicate that the Streams Directive state is not preserved across Controller Level Resets.
	00	Identify Directive: This bit shall be set cleared to '0' as the host is not able to change the state of the Identify Directive.
4095:6496	n/a	Reserved

8.7.2.2 Directive Send

This section defines operations used with the Directive Send command for the Identify Directive.

8.7.2.2.1 Enable Directive (Directive Operation 01h)

The Enable Directive operation is used to enable a specific Directive for use within a namespace by all controllers that are associated with the same Host Identifier. The DSPEC field in command Dword 11 is not used for this operation. The Identify Directive is always enabled. The enable state of each Directive on each shared namespace attached to enabled controllers associated with the same non-zero Host Identifier is the same. If the Directive is not the Data Placement Directive and an NSID value of FFFFFFFFh is specified, then the Enable Directive operation applies to the NVM subsystem (i.e., all namespaces and all controllers associated with the NVM subsystem). If the Directive is the Data Placement Directive and an NSID value of FFFFFFFFh is specified, then the controller shall abort the command with a status code of Invalid

Namespace or Format. ~~On an NVM Subsystem Reset, all Directives other than the Identify Directive are disabled in the Domains impacted by that reset (refer to section 3.7.1).~~

On a Controller Level Reset:

- all Directives other than the Identify Directive [that have the Directive Persistent Across Controller Level Resets bit cleared to '0'](#) are disabled for that controller; and
- if there is an enabled controller associated with the Host Identifier for the controller that was reset, then for namespaces attached to enabled controllers associated with that Host Identifier, Directives are not disabled.

If a host sets the Host Identifier of a controller to the same non-zero Host Identifier as one or more other controllers in the NVM subsystem, then setting that Host Identifier shall result in each shared namespace attached to that controller having the same enable state for each Directive as the enable state for each Directive for that namespace attached to other controllers associated with that Host Identifier.

If a host enables a controller that has the same non-zero Host Identifier as one or more other controllers in the NVM subsystem, then enabling that controller shall result in each shared namespace attached to that controller having the same enable state for each Directive as the enable state for each Directive for that namespace attached to other controllers associated with that Host Identifier.

For all controllers in an NVM subsystem that have the same non-zero Host Identifier, if a host changes the enable state of any Directive for a shared namespace attached to a controller by a means other than a Controller Level Reset, then that change shall be made to the enable state of that Directive for that namespace attached to any other controller associated with that Host Identifier.

[Refer to the sections defining each Directive Type for restrictions on enabling that directive.](#)

Figure 420: Enable Directive – Command Dword 12

Bits	Description
31:16	Reserved
15:08	Directive Type (DTYPE): This field specifies the Directive Type to enable or disable. If this field specifies the Identify Directive (i.e., 00h), then a status code of Invalid Field in Command shall be returned.
07:01	Reserved
00	Enable Directive (ENDIR): If set to '1' and the Directive Type is supported, then the Directive is enabled unless otherwise specified . If cleared to '0', then the Directive is disabled. If this bit is set to '1' for a Directive that is not supported, then a status code of Invalid Field in Command shall be returned.

8.7.3 Streams (Directive Type 01h, Optional)

...

Streams Directive defines the command specific status values specified in Figure 424.

Figure 424: Streams Directive – Command Specific Status Values

Value	Description
7Fh	Stream Resource Allocation Failed: The controller was not able to allocate stream resources for exclusive use of the specified namespace and no NVM subsystem stream resources are available.

[The Streams Directive is not allowed to be enabled in any namespace that is contained in an Endurance Group with Flexible Data Placement enabled. If the specified namespace is contained in an Endurance](#)

Group that has Flexible Data Placement enabled, then the controller shall not enable the Streams Directive. If:

- a Directive Send command specifies the Streams Directive in the DTYPE field and the Directive Operation field is set to 01h (i.e., Enable Directive); and
- the namespace specified by the NSID field is contained in an Endurance Group that has Flexible Data Placement enabled,

then the controller shall:

- abort that Directive Send command with a status code of Invalid Field in Command; and
- indicate the Directive Operation field in the Parameter Error Location field (refer to Figure 206) if an entry is added to the Error Information log page due to aborting that Directive Send command.

...

8.7.TBD Data Placement (Directive Type 02h, Optional)

The Data Placement Directive enables the host to specify to the controller the Reclaim Unit (refer section 8.TBD) to place the user data in I/O commands specified by the appropriate I/O Command Set specification.

The Data Placement Directive has no Directive Operations defined. Any Directive Receive command or Directive Send command that specifies a Data Placement Directive Type shall be aborted by the controller with a status code of Invalid Field in Command.

If the specified namespace is not contained in an Endurance Group with Flexible Data Placement enabled, then the controller shall abort the Directive Send command with a status code of FDP Disabled.

...

8.TBD Flexible Data Placement

8.TBD.1 Flexible Data Placement Overview

The Flexible Data Placement (FDP) capability is an optional capability which allows the host to reduce write amplification by aligning user data usage to physical media. The controller supports log pages which indicate the status of FDP, statistics about the FDP operation, and information the host uses to detect and correct usage patterns that increase write amplification.

The scope of the Flexible Data Placement capability is an Endurance Group.

The host enables or disables Flexible Data Placement by issuing a Set Features command specifying the Flexible Data Placement feature (refer to section 5.27.1.TBD) and the FDP configuration (refer to section 5.16.1.TBD_{Log}) to be applied to an Endurance Group. The host is required to delete all namespaces associated with the specified Endurance Group before modifying the value of the Flexible Data Placement feature (e.g., transitioning from disabled to enabled).

If a Set Features command specifies:

- the Flexible Data Placement feature;
- an Endurance Group in which one or more namespaces exist; and
- a different value from the current value of that feature for that Endurance Group,

then the controller shall abort the command with a status code of Command Sequence Error.

If a Set Features command is successful and changes the Feature value, then:

- all events in the FDP Events log page (refer to section 5.16.1.TBD_{Log3}) are cleared; and
- all fields in the FDP Statistics log page (refer to section 5.16.1.TBD_{Log2}) are cleared to 0h.

Refer to section 3.2.TBD for the logical view of the non-volatile storage capacity for the Endurance Group with Flexible Data Placement enabled.

The Namespace Management command (refer to section 8.11) is used to create a namespace within the Endurance Group. All namespaces in an Endurance Group with Flexible Data Placement enabled shall be associated with the NVM Command Set (refer to Figure 274) (i.e., the Command Set Identifier (CSI) field in the Namespace Management command shall be cleared to 00h). The capacity for the user data stored by a write command to a namespace is allocated from the Reclaim Unit referenced by the Reclaim Unit Handle and Reclaim Group specified by that write command. The user data for a namespace is allowed to be in any stored Reclaim Unit within any Reclaim Group within the Endurance Group (refer to Figure FDPEX3).

The Namespace Management command specifies a Placement Handle List (refer to the Namespace Management command in the appropriate I/O Command Set specification). Each entry in that list specifies the Reclaim Unit Handle associated with the Placement Handle for that entry. The Placement Handles are numbered from 0h to the number of entries in the list minus 1. Figure ZZ shows an example where Namespace A associates Reclaim Unit Handle 1 with Placement Handle 0. A specific Reclaim Unit Handle is not allowed to be associated with more than one Placement Handle per namespace. There are no other requirements on which Reclaim Unit Handle is associated with which Placement Handle.

Namespaces are allowed to utilize the same (i.e., share) Reclaim Unit Handle but may have restrictions defined in the applicable I/O Command Set specification.

The host is required to enable the Data Placement Directive (refer to section 8.7.TBD) to submit write commands that specify a Placement Identifier (refer to Figure PIDRG and Figure PIDRG_NORGI) in the DSPEC field. That Placement Identifier uniquely specifies the Reclaim Unit in which to place the user data (refer to Figure ZZ). The controller determines the Reclaim Unit by using the Reclaim Unit Handle associated with the Placement Handle that was specified by the host when the namespace was created. For example, if the write command in Figure ZZ specifies a Reclaim Group Identifier value of 2h and a Placement Handle value of 1h, then the host is targeting the user data for that write command to be placed into the Reclaim Unit associated with Reclaim Group 2h and Reclaim Unit Handle *NRUH-1*.

A Placement Identifier is invalid if:

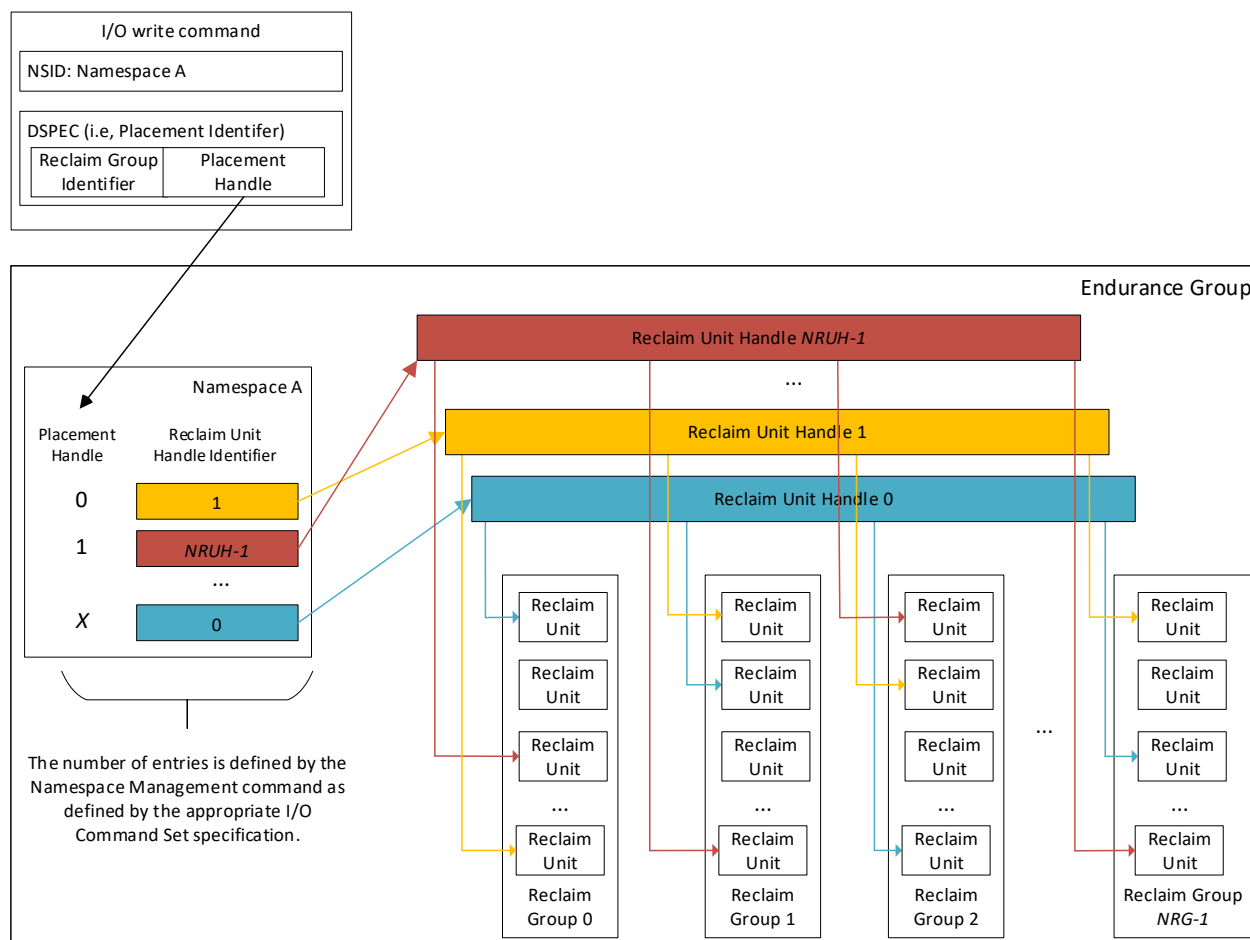
- the Reclaim Group Identifier value is greater than or equal to the number of Reclaim Groups supported by the enabled FDP configuration (refer to the NRG field in Figure FDPD); or
- the Placement Handle is greater than or equal to the number of Placement Handles supported by the Namespace Management command of the applicable I/O Command Set specification.

If the Placement Identifier in the DSPEC field of a write command is invalid, then the controller shall:

- select the Reclaim Group and Reclaim Unit Handle accessible by the specified namespace to determine the placement of the user data for that write command; and
- generate an Invalid Placement Identifier FDP Event if that event is enabled for the selected Reclaim Unit Handle. To receive the FDP event, the host should enable the Invalid Placement Identifier FDP Event on all Reclaim Unit Handles.

Any write command to a namespace that exists in an Endurance Group with Flexible Data Placement enabled that does not specify the Data Placement Directive uses the Placement Handle value 0h and the controller selects the Reclaim Group for that command. For example, if the write command does not specify the Data Placement Directive, then, as illustrated in [Figure ZZ](#), Placement Handle value of 0h is associated with Reclaim Unit Handle 1h and the controller may place the user data into the Reclaim Unit associated any Reclaim Group using Reclaim Unit Handle 1h.

Figure ZZ: Flexible Data Placement Model



A write command that places user data into a Reclaim Unit causes reduction of the remaining capacity of that Reclaim Unit. If that Reclaim Unit is written to capacity, then the controller modifies the Reclaim Unit Handle referencing that Reclaim Unit to reference a different Reclaim Unit within the same Reclaim Group.

A controller may modify a Reclaim Unit Handle to reference a different Reclaim Unit as part of performing a sanitize operation (refer to section [8.21](#)).

A Reclaim Unit shall only be referenced by one Reclaim Unit Handle. Therefore, a write command that utilizes a Reclaim Unit Handle results in the user data from that write command being initially isolated from the user data from any other write command that utilizes a different Reclaim Unit Handle. Each Reclaim Unit Handle has a type as defined in [Figure FDPRUHD](#) that identifies the isolation requirements of the user data moved by the controller to a different Reclaim Unit due to vendor specific controller operations (e.g., garbage collection).

A controller is allowed to move user data that was written by the host that utilized a Reclaim Unit Handle with an Initially Isolated type (Figure FDP RUHD) to a different Reclaim Unit within the same Reclaim Group. In addition, the controller is allowed to move user data written by the host that utilized different Reclaim Unit Handles with an Initially Isolated type to that different Reclaim Unit within the same Reclaim Group as shown in Figure Y.

A controller is allowed to move user data that was written by the host that utilized a Reclaim Unit Handle with a Persistently Isolated type (Figure FDP RUHD) to a different Reclaim Unit within the same Reclaim Group. However, the controller shall only move user data written by the host that utilized the same Reclaim Unit Handles with a Persistently Isolated type to the different Reclaim Unit within the same Reclaim Group as shown in Figure Z.

Figure Y: Initially Isolated Reclaim Unit Handles

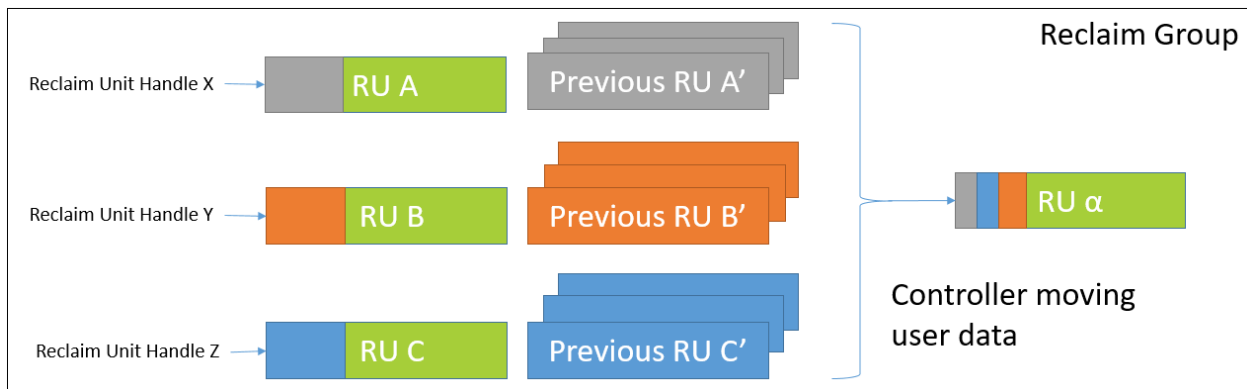
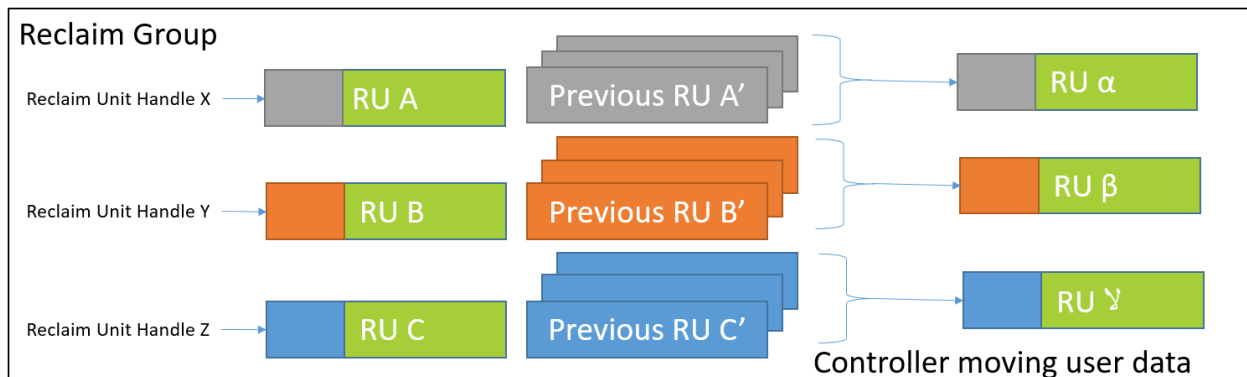


Figure Z: Persistently Isolated Reclaim Unit Handle



The result of a Namespace Management command that creates a namespace within the Endurance Group is that at least one Placement Handle is defined for that namespace (refer to the Namespace Management section in applicable I/O Command Set specification).

For any namespace created in an Endurance Group that has Flexible Data Placement enabled, the host may issue an I/O Management Receive command to obtain a list of Placement Handles and the associated Reclaim Unit Handles that are used by the Data Placement Directive in write commands.

If the Flexible Data Placement capability is supported, then the controller shall support the following:

- a) the Feature Identifiers Supported and Effects log page (refer to section 5.16.1.18);
- b) the I/O Management Receive command (refer to section 7.TBD_{IOMR}) and the Reclaim Unit Handle Status Management Operation in that command (refer to Figure IOMR2);
- c) the I/O Management Send command (refer to section 7.TBD_{IOMS}) and the Placement Identifier Update Management Operation in that command (refer to Figure IOMS2);
- d) the Flexible Data Placement feature (refer to section 5.27.1.TBD);
- e) the FDP Configurations log page (refer to section 5.16.1.TBD_{Log});
- f) the Reclaim Unit Handle Usage log page (refer to section 5.16.1.TBD_{Log1});
- g) the FDP Statistics log page (refer to section 5.16.1.TBD_{Log2});
- h) the Flexible Data Placement Events feature (refer to section 5.27.1.TBD1);
- i) the FDP Events log page (refer to section 5.16.1.TBD_{Log3}); and
- j) the Data Placement Directive (refer to section 8.7.TBD).

8.TBD.2 Enabling Flexible Data Placement (Informative)

The host prepares an Endurance Group for operation in Flexible Data Placement using the following process:

- 1) Validate that the Flexible Data Placement capability is supported by issuing an Identify command to access the Identify Controller data structure and checking that the Flexible Data Placement Support (FDPS) bit is set to '1'.
- 2) Delete any existing namespaces that exist in the Endurance Group where Flexible Data Placement is to be enabled.
- 3) Issue a Get Log Page command specifying the FDP Configurations log page (refer to section 5.16.1.TBD_{Log}). Parse the FDP Configurations List in the returned log page to determine the desired FDP configuration.
- 4) Enable Flexible Data Placement utilizing that desired FDP configuration by issuing a Set Features command specifying:
 - a. the Flexible Data Placement feature;
 - b. the Endurance Group Identifier field set to the ENDGID of the Endurance Group in which Flexible Data Placement is to be enabled;
 - c. the FDPE bit set to '1' (i.e., enabling Flexible Data Placement); and
 - d. the Flexible Data Placement Configuration Index field set to the index of the desired FDP configuration from the FDP Configurations List in the FDP Configurations log page.
- 5) Issue Get Features commands for the FDP Events feature (refer to section 5.27.1.TBD1) to acquire the list of supported FDP events and the enabled state of each supported FDP event.
- 6) Issue Set Features commands for the FDP Events feature (refer to section 5.27.1.TBD1) to specify if FDP events are required for Reclaim Unit Handles.
- 7) Issue an Identify command specifying the Identify Namespace data structure (i.e., CNS 00h as defined in Figure 273) to identify which LBA formats are supported.
- 8) For each namespace to be created in the Endurance Group, issue a Namespace Management command specifying:
 - a. the Select field set to the Create operation (refer to the NVM Express Base Specification);
 - b. the User Data Format; and
 - c. the Placement Handle List used to define the Reclaim Unit Handle associated with each Placement Handle of the namespace (refer to the Namespace Management section of the applicable I/O Command Set specification).

- 9) For each namespace created in an Endurance Group in which the Data Placement Directive is to be utilized by the host by write commands specifying that namespace:
 - a. Issue a Directive Send command specifying:
 - i. the Directive Operation field set to Enable Directive (i.e., 01h);
 - ii. the NSID of that namespace;
 - iii. the DTYPE field in Command Dword 11 set to Identify (i.e., 00h);
 - iv. the DTYPE field in Command Dword 12 set to Data Placement (i.e., 02h); and
 - v. the Enable Directive bit set to '1' (i.e., to enable the Data Placement Directive).
 - b. Issue an Identify command with CNS value 08h (i.e., the I/O Command Set Independent Identify Namespace data structure) to determine if a volatile write cache is present for the namespace.

Following a Controller Level Reset, the host performs the following actions before resuming use of Flexible Data Placement in an Endurance Group:

- 1) Determine the FDP configuration:
 - a. Issue a Get Feature command specifying the Flexible Data Placement feature (i.e., 1Dh) and the ENDGID for the Endurance Group to determine if Flexible Data Placement is enabled and the Flexible Data Placement Configuration Index.
 - b. Issue a Get Log Page command specifying the FDP Configurations log page (i.e., Log Page Identifier 20h).
- 2) Issue Get Features commands for the FDP Events feature (refer to section 5.27.1.TBD1) to acquire the list of supported FDP events and the enabled state of each supported FDP event.
- 3) Issue Set Features commands for the FDP Events feature (refer to section 5.27.1.TBD1) to specify if supported FDP events are required for Reclaim Unit Handles.
- 4) For each namespace in the Endurance Group that the Data Placement Directive is to be utilized by the host by write commands specifying that namespace:
 - a. Issue an I/O Management Receive command to determine the Placement Handles for that namespace; and
 - b. Issue an Identify command with CNS value 08h (i.e., the I/O Command Set Independent Identify Namespace data structure) to determine if a volatile write cache is present for the namespace.

8.TBD.3 Write Commands using Flexible Data Placement (Informative)

Hosts writing user data to a namespace created in an Endurance Group that has Flexible Data Placement enabled have two mechanisms for specifying the placement of the user data with the Endurance Group to the controller.

1. Each namespace has a default Placement Handle. If any write command:
 - a. has the DTYPE field being ignored by the controller (refer to section 8.7.1); or
 - b. has the DTYPE field not being ignored by the controller and specifies a DTYPE field value cleared to 00h,

then the Placement Handle value of 0h for the specified namespace is combined with a Reclaim Group selected by the controller to create the Placement Identifier (refer to Figure PIDRG and Figure PIDRG_NORGI) utilized for that write command.

2. If any write command has the DTYPE field not being ignored by the controller and specifies a DTYPE field set to Data Placement (i.e., 02h) and specifies a namespace in which the Data Placement Directive is enabled, then the Placement Identifier specified by the DSPEC field specifies the Reclaim Group Identifier and Placement Handle used for the write command.

To reduce write amplification, hosts track the user data written to an entire Reclaim Unit and deallocate that user data by issuing one or more Dataset Management commands specifying the AD bit set to '1' and one or more Range fields specifying the LBA ranges that were written to that Reclaim Unit.

Modify portions of the NVM Command Set Specification 1.0b as follows:

3 I/O Commands for the NVM Command Set

...

<Editor Note: Add the following new section before section 3.2>

3.TBD I/O Command behavior for the NVM Command Set

This section defines specific behavior for I/O commands defined in the NVM Express Base Specification for the NVM Command Set.

3.TBD.1 I/O Management Receive command

3.TBD.1.1 Reclaim Unit Handle Status (Management Operation 01h)

The Reclaim Unit Handle Status Descriptor utilized for this management operation for the NVM Command Set is defined in **Figure FigRUHSD**. The Reclaim Unit Handle Status Descriptor in the Reclaim Unit Handle Status Descriptor List shall be listed first in ascending order of Placement Handle and second in ascending order of Reclaim Group Identifier.

Figure FigRUHSD: Reclaim Unit Handle Status Descriptor

Bytes	Description
1:0	Placement Identifier (PID): This field indicates the Placement Identifier (refer to the Reclaim Unit Handle Status section in the NVM Express Base Specification) containing the Placement Handle and Reclaim Group Identifier for this Reclaim Unit Handle Status Descriptor.
3:2	Reclaim Unit Handle Identifier: This field indicates the Reclaim Unit Handle for the Placement Identifier field.
7:4	Estimated Active Reclaim Unit Time Remaining (EARUTR): This field indicates an estimate of the time in seconds that the Reclaim Unit currently referenced by the Reclaim Unit Handle is allowed to remain referenced by that Reclaim Unit Handle (refer to the Flexible Data Placement section in the NVM Express Base Specification) before the controller may modify the Reclaim Unit Handle to reference a different Reclaim Unit. This value is the remaining time at the time the I/O Management Receive command is processed by the controller. If this field is cleared to 0h, then no time is reported.
15:08	Reclaim Unit Available Media Writes (RUAMW): This field indicates the number of logical blocks which are currently able to be written to the media associated with the Reclaim Unit currently referenced by the Placement Identifier field. The product of this field and the Formatted LBA Size field is able to be less than the nominal size (refer to the RUNS field in the FDP Configurations log page defined in the NVM Express Base Specification) of the Reclaim Unit (e.g., the Reclaim Unit has not been written, but excess defects prevent writing some of the media). The product of this field and the Formatted LBA Size field is able to be greater than the nominal size (refer to the RUNS field in the FDP Configurations log page defined in the NVM Express Base Specification) of the Reclaim Unit (e.g., there is over-provisioned capacity to support the Reclaim Unit). The value of this field may or may not be modified by a Controller Level Reset or the processing of a Flush command.
31:16	Reserved

3.2 NVM Command Set Commands

...

3.2.7 Write Uncorrectable command

...

Figure 69: Write Uncorrectable – Command Dword 12

Bits	Description
31:24	Reserved
23:20	Directive Type (DTYPE) : Specifies the Directive Type associated with the Directive Specific field (refer to the Directives section in the NVM Express Base Specification).
19:16	Reserved
15:00	Number of Logical Blocks (NLB) : This field specifies the number of logical blocks to become uncorrectable. This is a 0's based value.

Figure TBDWU: Write Uncorrectable– Command Dword 13

Bits	Description
31:16	Directive Specific (DSPEC) : Specifies the Directive Specific value associated with the Directive Type field (refer to the Directives section in the NVM Express Base Specification).
15:00	Reserved

...

3.2.8 Write Zeroes command

...

Figure 73: Write Zeroes – Command Dword 12

Bits	Description
31	Limited Retry (LR) : If set to '1', the controller should apply limited retry efforts. If cleared to '0', the controller should apply all available error recovery means to write the data to the NVM.
30	Force Unit Access (FUA) : If set to '1', then the controller shall write the data, and metadata, if any, to non-volatile media before indicating command completion. There is no implied ordering with other commands. If cleared to '0', then this bit has no effect.
29:26	Protection Information Field (PRINFO) : Specifies the protection information action and check field, as defined in Figure 9. The Protection Information Check (PRCHK) field shall be cleared to 000b.
25	Deallocate (DEAC) : If set to '1', then the host is requesting that the controller deallocate the specified logical blocks. If cleared to '0', then the host is not requesting that the controller deallocate the specified logical blocks.
24	Storage Tag Check (STC) : This bit specifies the Storage Tag field shall be checked as part of end-to-end data protection processing as defined in Figure 10. This bit shall be cleared to '0'.
23:20	Directive Type (DTYPE) : Specifies the Directive Type associated with the Directive Specific field (refer to the Directives section in the NVM Express Base Specification).
19:16	Reserved
15:00	Number of Logical Blocks (NLB) : This field indicates the number of logical blocks to be written. This is a 0's based value.

Figure TBDWZ: Write Zeroes– Command Dword 13

Bits	Description
31:16	Directive Specific (DSPEC): Specifies the Directive Specific value associated with the Directive Type field (refer to the Directives section in the NVM Express Base Specification).
15:08	Reserved

Figure 74: Write Zeroes – Command Dword 14

Bits	Description
31:00	This field and bits 47:00 of Command Dword 2 and Dword 3 specify the variable sized Logical Block Storage Tag (LBST) and Initial Logical Block Reference Tag (ILBRT) fields, which are defined in section 5.2.1.4.1. If the namespace is not formatted to use end-to-end protection information, then this field is ignored.

...

4 Admin Commands for the NVM Command Set

...

4.1 Admin Command behavior for the NVM Command Set

...

4.1.2 Format NVM command

The Format NVM command operates as defined in the NVMe Base Specification. The Format Index indicates a valid User Data Format from the LBA Format field in the Identify Namespace data structure. Other NVM Command Set specific fields are defined in Figure 78.

For the NVM Command Set, if the Format NVM command results in a change of the logical block size for the namespace, then the resulting namespace size (i.e., NSZE) (refer to Figure 97) and the namespace capacity (i.e., NCAP) (refer to Figure 97) may differ from the values indicated prior to the processing of the Format NVM command.

If the LBA Format Extension Enable (LBAFEE) field is not set to 1h in the Host Behavior Support feature (refer to the Host Behavior Support section in the NVMe Base Specification), then the controller aborts a Format NVM command with a status code of Invalid Namespace or Format that specifies a format (refer to section 5.2.1) of:

- 16b Guard Protection Information with the STS field set to a non-zero value;
- 32b Guard Protection Information; or
- 64b Guard Protection Information.

If Flexible Data Placement (refer to the NVM Express Base Specification) is enabled in the Endurance Group associated with the specified namespace, then:

- If any Reclaim Unit Handle utilized by the specified namespace is shared by other namespaces and the specified Format Index does not match the Format Index of the other namespaces, then the command shall be aborted with a status code of Invalid Format. Refer to section 4.1.6.

...

4.1.4 Get Log Page command

4.1.4.TBD Flexible Data Placement (FDP) Statistics (Log Page Identifier 22h)

The Flexible Data Placement Statistics log page is defined in the NVM Express Base Specification. The I/O commands that are specifically used by the NVM Command Set for the Host Bytes with Metadata Written (HBMW) field and the Media Bytes with Metadata Written (MBMW) field are the User Data Out Commands, the Write Zeroes command, and the Write Uncorrectable command.

4.1.4.TBD1 Flexible Data Placement (FDP) Events (Log Page Identifier 23h)

4.1.4.TBD1.1 Controller Events

4.1.4.TBD1.1.1 Media Reallocated (Event Type 80h)

The Flexible Data Placement Events log page is defined in the NVM Express Base Specification. **Figure ETS_MR** describes the NVM Command Set specific definition of the Media Reallocated event.

Figure ETS_MR: Media Reallocated - Event Type Specific Data Structure

Bytes	Description						
00	Specific Event Flags: This field indicates specific attributes of the event.						
	<table><tr><th>Bit</th><th>Description</th></tr><tr><td>07:01</td><td>Reserved</td></tr><tr><td>00</td><td>LBA Valid (LBAV): If set to '1', the LBA field contains a valid value. If cleared to '0', then the LBA field does not contain a valid value.</td></tr></table>	Bit	Description	07:01	Reserved	00	LBA Valid (LBAV): If set to '1', the LBA field contains a valid value. If cleared to '0', then the LBA field does not contain a valid value.
	Bit	Description					
07:01	Reserved						
00	LBA Valid (LBAV): If set to '1', the LBA field contains a valid value. If cleared to '0', then the LBA field does not contain a valid value.						
01	Reserved.						
03:02	Number of LBAs Moved (NLBAM): This field indicates the number of LBAs moved by the controller from a Reclaim Unit written by the host. If the PIV bit is set to '1', then the Placement Identifier field indicates the Reclaim Unit Handle used to initially write the LBAs. A value of 0h in this field indicates that no number is being reported. A value of FFFFh means that FFFFh or more LBAs were moved by the controller.						
11:04	Logical Block Address (LBA): This field indicates one of the LBAs moved by the controller. If the PIV bit is set to '1', then the Placement Identifier field indicates the Reclaim Unit Handle used to initially write the LBAs. If the LBAV bit is cleared to '0', then this field shall be cleared to 0h and the host should ignore this field.						
15:12	Reserved						

...

4.1.5 Identify Command

...

4.1.5.3 I/O Command Set Specific Identify Namespace Data Structure (CNS 05h)

...

< **Note to Editor:** The NPDGL field was added by TP 4090. The Reported column is from TP4095a>

Figure 100: NVM Command Set I/O Command Set Specific Identify Namespace Data Structure (CSI 00h)

Bytes	O/M ¹	Description	Reported ²
...	
267:264	O	Extended LBA Format 63 Support (ELBAF63): This field indicates additional LBA format information to the LBA Format 63 Support (ELBAF63) field in the Identify Namespace data structure. The Extended LBA format field is defined in Figure 100.	Yes
271:268	O	Namespace Preferred Deallocate Granularity Dword (NPDGL): This field indicates the recommended granularity in logical blocks for the Dataset Management command with the Attribute – Deallocate bit set to '1' in Command Dword 11. This is a 1's based value. If this field is not supported as defined by the OPTPERF field (refer to Figure 97), then this field is reserved. If this field is cleared to 0h, then this field does not indicate a recommended granularity. The value of this field may change if the namespace is reformatted. The size should be a multiple of the Namespace Preferred Deallocate Alignment (NPDA) field (refer to Figure 97). Refer to section 5.8.2 for how this field is utilized to improve performance and endurance.	No
291:288	O	LBA Format Placement Shard Size: This field indicates the optimal number of LBAs to be written to a Reclaim Group and then written to the other Reclaim Groups in order to maximize I/O write performance. A value of 0h means that no Placement Shard Size value is reported.	No
4095: 268 292	O	Reserved	
NOTES: 1. O/M definition: O = Optional, M = Mandatory. 2. Identifies fields that report information for the Identify command when querying the capabilities of LBA formats.			

The Extended LBA format data structure is described in Figure 101.

...

4.1.6 Namespace Management command

...

If Flexible Data Placement (refer to the NVM Express Base Specification) is enabled in the specified Endurance Group and the Select field is set to Create (i.e., 0h):

- The Placement Handle List allows the host to specify the Reclaim Unit Handle associated with each specified Placement Handle. The number of Placement Handles in the list is specified by the NPHNDLS field which is limited to a value less than or equal to the lesser value of:
 - the number of Reclaim Unit Handles supported by the FDP configuration (refer to the NVM Express Base Specification); and
 - 128.
- If the NPHNDLS field is cleared to 0h, then:

- if no namespace exists that was created with a Namespace Management command that specified the NPHNDLS field cleared to 0h in the specified Endurance Group, then the controller shall select a Reclaim Unit Handle for the Placement Handle 0 Associated RUH field that is not utilized by any namespace in the same Endurance Group that was created with a Namespace Management command that specified a non-zero NPHNDLS field; and
 - if any namespace exists that was created with a Namespace Management command that specified the NPHNDLS field cleared to 0h in the specified Endurance Group, then the controller shall utilize the same Reclaim Unit Handle for the Placement Handle 0 Associated RUH field that is utilized by those namespaces in the same Endurance Group that were created with a Namespace Management command that specified the NPHNDLS field cleared to 0h.
- If:
 - the NPHNDLS field is cleared to 0h; and
 - all Reclaim Unit Handle Identifiers accessible to the namespace are allocated to namespaces created by Namespace Management command with the NPHNDLS field set to a non-zero value,then the controller shall abort the command with a status code of Invalid Placement Handle List.
- If:
 - the NPHNDLS field is non-zero; and
 - a Reclaim Unit Handle Identifier specified by the host is the same as the Reclaim Unit Handle Identifier selected by the controller due to an existing namespace being created by Namespace Management command with the NPHNDLS field cleared to 0h,then the controller shall abort the command with a status code of Invalid Placement Handle List.
- If the NPHNDLS field is greater than the lesser value of:
 - the number of Reclaim Unit Handles supported by the FDP configuration (refer to the NVM Express Base Specification); and
 - 128,then the controller shall abort the command with a status code of Invalid Placement Handle List.
- If a Reclaim Unit Handle Identifier value in any entry in the Placement Handle List is greater than or equal to the number of Reclaim Unit Handles supported by the FDP configuration for the Endurance Group (refer to the NVM Express Base Specification), then controller shall abort the command with a status code of Invalid Placement Handle List.
- If the same Reclaim Unit Handle Identifier value is in two or more entries in the Placement Handle List, then controller shall abort the command with a status code of Invalid Placement Handle List.
- Namespaces that exist in the specified Endurance Group that utilize the same (i.e., share) Reclaim Unit Handle shall have the same user data format (i.e., report the same Format Index). If a Reclaim Unit Handle specified in the Placement Handle List is utilized by another namespace and the Format Index for that namespace does not match the specified Format Index for the namespace to be created, then the controller shall abort the command with a status code of Invalid Format.

Figure 105: Namespace Management – Host Software Specified Fields

Bytes	Description	Host Specified
Fields that are a subset of the Identify Namespace data structure (refer to Figure 97)		
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) ^{1,3}	Yes
103:102	Endurance Group Identifier (ENDGID)	Yes
383:104	Reserved	
Fields that are not a subset of the Identify Namespace data structure.		
391:384	Logical Block Storage Tag Mask (LBSTM)	Yes
393:392	Number of Placement Handles ² (NPHNDLS): This field specifies the number of Placement Handles included in the Placement Handle List. If the Flexible Data Placement capability (refer to the NVM Express Base Specification) is not supported or not enabled in specified Endurance Group, then the controller shall ignore this field.	Yes
498:392394	Reserved	
511:499	Reserved for I/O Command Sets that extend this specification. Refer to the applicable I/O Command Set specification (e.g., Zoned Namespace Command Set Specification).	
Placement Handle List		
513:512	Placement Handle 0 Associated RUH²: This field specifies the Reclaim Unit Handle Identifier to be associated with the Placement Handle value 0h, if any. This Reclaim Unit Handle Identifier is used by the controller for any write commands that do not specify the Data Placement Directive. If the Flexible Data Placement capability (refer to the NVM Express Base Specification) is not supported or not enabled in specified Endurance Group, then the controller shall ignore this field.	Yes
515:514	Placement Handle 1 Associated RUH²: This field specifies the Reclaim Unit Handle Identifier to be associated with the Placement Handle value 1h, if any. If the Flexible Data Placement capability (refer to the NVM Express Base Specification) is not supported or not enabled in specified Endurance Group, then the controller shall ignore this field.	Yes
...		
767:766	Placement Handle 127 Associated RUH²: This field specifies the Reclaim Unit Handle Identifier to be associated with the Placement Handle value 127, if any. If the Flexible Data Placement capability (refer to the NVM Express Base Specification) is not supported or not enabled in specified Endurance Group, then the controller shall ignore this field.	Yes
4095514:392768	Reserved	

Figure 105: Namespace Management – Host Software Specified Fields

Bytes	Description	Host Specified
Fields that are a subset of the Identify Namespace data structure (refer to Figure 97)		
Notes: 1. A value of 0h specifies that the controller determines the value to use (refer to the Namespace Management section of the NVMe Base Specification). If the associated feature is not supported, then this field is ignored by the controller. 2. Refer to the Flexible Data Placement section in the NVM Express Base Specification for requirements and use of this field. These fields are reserved if Flexible Data Placement is disabled in the specified Endurance Group. 3. NVM Sets are not supported if FDP is enabled in the specified Endurance Group as defined in the NVM Express Base Specification.		

...

Modify portions of the NVMe Zoned Namespace Command Set Specification 1.1b as follows:

4 Admin Commands for the Zoned Namespace Command Set

4.1 Admin Command behavior for the Zoned Namespace Command Set

...

4.1.6 Namespace Management command

...

< **Note to Editor:** The following figure was added to the ZNS Command Set spec by TP4115. The text from that TP is in black. The effect of this change is to add one byte to the size of the row “Refer to the NVM Command Set Specification”. That final byte range must match Figure 105 in the NVM Command Set section of this TP. To distinguish between TBDs from that TP and this TP, the TP number is prepended to the TBDs. >

Figure TP4115_TBDX: Namespace Management – Host Software Specified Fields

Bytes	Description	Host Specified
Fields that are a subset of the Identify Namespace data structure		
383:00	Refer to the NVM Command Set Specification	
Fields that are not a subset of the Identify Namespace data structure		
498391:384	Refer to the NVM Command Set Specification	Yes
498:392	Reserved	
499	Zoned Namespace Create Options	
	Bits	Description
	07:01	Reserved
	00	Allocate ZRWA Resources (AZR): If set to ‘1’, then the namespace is to be created with the number of ZRWA resource specified in the RNUMZRWA field of this data structure. If cleared to ‘0’, then no ZRWA resources are allocated to the namespace to be created. If the ZRWASUP bit is cleared to ‘0’ (refer to Figure 48), then this field shall be ignored by the controller.
...		

Figure TP4115_TBDX: Namespace Management – Host Software Specified Fields

Bytes	Description	Host Specified
511:508	<p>Requested Number of ZRWA Resources (RNUMZRWA)¹: This field specifies the number of ZRWA resources to be allocated to the created namespace.</p> <p>A value of 00000000h specifies the number of ZRWA resources requested to be allocated to the namespace is vendor specific and that number is reflected in the NUMZRWA field in the I/O Command Set Specific Identify Namespace data structure.</p> <p>Values of 00000001h to FFFFFFFEh specify the number of ZRWA resources requested to be allocated is as defined in section 2.1.1.4 and that number is reflected in the NUMZRWA field in the I/O Command Set Specific Identify Namespace data structure.</p> <p>A value of FFFFFFFFh specifies the number of ZRWA resources requested to be allocated to the namespace is not limited and that number is reflected in the NUMZRWA field in the I/O Command Set Specific Identify Namespace data structure.</p>	Yes
4095:512	Reserved	

...

Modify portions of the NVM Express Management Interface Specification 1.2b as follows:

6 NVM Express Admin Command Set

...

6.3 Get Log Page

Figure 122 defines the log pages that are mandatory, optional, and prohibited for SMBus/I2C and PCIe VDM Management Endpoint on NVMe Storage Devices and NVMe Enclosures.

Figure 122: Management Endpoint - Log Page Support

Log Page Name ³	Log Identifier	Requirements ¹	
		NVMe Storage Device	NVMe Enclosure
...			
FDP Configurations	20h	O	O
Reclaim Unit Handle Usage	21h	O	O
FDP Statistics	22h	O	O
FDP Events	23h	O	O
...			

...

6.5 Set Features and Get Features

...

Figure 126: Management Endpoint - Feature Support

Feature Name ²	Feature Identifier	SMBus/I2C Feature Support Requirements ¹	
		NVMe Storage Device	NVMe Enclosure
...			
Flexible Data Placement	1Dh	P	P
Flexible Data Placement Events	1Eh	P	P
...			

...