



LEGAL NOTICE:

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

This NVM Express revision 2.0 technical proposal is proprietary to the NVM Express, Inc. (also referred to as “Company”) and/or its successors and assigns.

NOTICE TO USERS WHO ARE NVM EXPRESS, INC. MEMBERS: Members of NVM Express, Inc. have the right to use and implement this NVM Express revision 2.0 technical proposal subject, however, to the Member’s continued compliance with the Company’s Intellectual Property Policy and Bylaws and the Member’s Participation Agreement.

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

LEGAL DISCLAIMER:

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

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

The NVM Express® design mark is a registered trademark of NVM Express, Inc.

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	TP 4095 Namespace Capability Reporting
Change Date	2021-06-23
Builds on Specification	NVM Express Base Specification 2.0 NVM Express NVM Command Set 1.0 NVM Express Key Value Command Set 1.0 NVM Express Zone Namespace Command Set 1.1
References Specification	

Technical Proposal Author(s)

Name	Company
Mike Allison, Ambika Krishnamoorthy, Jonathan Hughes, Kapil Karkra	Intel
Mike Allison	Samsung

This proposal intends to:

- Define the common namespace capabilities.;
- Define a mechanism to support LBA formats that are not part of the common namespace capabilities and to report the namespace capabilities of that LBA format without having to create the namespace.

Revision History

Revision Date	Change Description
2021-01-21	Initial rough version
2021-02-24	Aligned to latest NVMe 2.0 specifications.
2021-02-25	Made all new CNS require the CSI field. Removed comments where WG approved text.
2021-02-26	Made the new CNS values optional. Modified the CNS Specific Identifier to be patterned after the Log Specific Field in the Get Log Page command in that the description of the identifier used is defined by the definition section of a CNS value.
2021-03-11	Accepted all changes and resolved comment. Ready for Phase 2 exit.
2021-03-25	Aligned to latest member review documents. Editorial changes due to Phase 3 review in Technical WG meeting.
2021-04-01	Corrected the data structure names. Minor edited during WG discussion.
2021-04-08	Changes the “may” to a “shall” on the format of the returned namespace data structures per WG request.
2021-04-22	Accepted all changes and converted all references/cross-references to text for the 30 day member review.
2021-04-27	Corrected the summing of the NLBAF field and the NULBAF field to account for one being 0-based and one being 1-based as reported by Edward Hseih.
2021-04-29	Modified figure X to refer to fields as the diagram to reference the value represented by the field and not just the field value.
2021-05-27	Correct pluralization and removed redundant words.

2021-06-03	Addressed comments from Gerry Houlder (Seagate). Approved in NVMe Technical WG
2021-06-03	Aligned to the released versions of NVMe 2.0 family of specifications. Removed large chunks of unchanged text. Removed ZNS Command Spec section 4.1.2 as change already applied. Corrected the order of section 5.17.2.4 in the NVMe Base Spec.
2021-06-21	Integrated into the NVMe Command Set Specification, revision 1.0; the NVMe Key Value Command Set specification, revision 1.0; the NVMe Zoned Namespace Command Set Specification, revision 1.1; and the NVMe Base Specification, revision 2.0.
2021-06-23	Removed all comments, accepted all changes, and converted references/cross-references to text.

Description for NVM Express NVM Command Set Specification 1.0 Changes Document

Feature Enhancement:

- Added CNS value **TBDh** to the Identify command to return an NVM Command Set Identify Namespace data structure containing the capability fields per the specified User Data Formal index. The capability fields are used to create and format namespaces.
- Added CNS value **TBD+1h** to the Identify command to return an Identify I/O Command Set specific Namespace data structure for the NVM Command Set containing the capability fields per the specified User Data Formal index and Command Set Identifier. The capability fields are used to create and format namespaces.
- **New requirement / incompatible change:**
 - The Identify command for CNS values 00h, 05h and 1Dh removed the term “common” and clearly specified the requirements on all fields when the specified NSID is set to FFFFFFFFh
 - The requirement for an Identify command with CNS values that returned the capabilities in a namespace data structure if the NSID specified is FFFFFFFFh to abort the command if the Namespace Management has been changed to be optional.
 - The LBA Format list structure is defined to have a group of LBA Formats that have the same capabilities for creating and formatting a namespace and a group of LBA Formats that each may have a unique set of capabilities for creating and formatting a namespace.
- References:
 - Sections 4.1.2, 4.1.3.4, 4.1.5, 5.TBD

Description for NVM Express Key Value Command Set Specification 1.0 Changes Document

Feature Enhancement:

- Added CNS value **TBD+1h** to the Identify command to return an Identify I/O Command Set specific Namespace data structure for the Key Value Command Set containing the capability fields per the specified User Data Formal index and Command Set Identifier. The capability fields are used to create and format namespaces.
- References:
 - Sections 4.1.5, 4.1.5.1, and 4.1.5.TBD1

Description for NVM Express Zoned Namespaces Command Set Specification 1.0 Changes Document

Feature Enhancement:

- Added CNS value **TBDh** to the Identify command to return an NVM Command Set Identify Namespace data structure containing the capability fields per the specified User Data Formal index. The capability fields are used to create and format namespaces.
- Added CNS value **TBD+1h** to the Identify command to return an Identify I/O Command Set specific Namespace data structure for the Zoned Namespace Command Set containing the capability fields per the specified User Data Formal index and Command Set Identifier. The capability fields are used to create and format namespaces.
- **New requirement / incompatible change:**
 - The Identify command for CNS values 05h removed the term “common” and clearly specified the requirements on all fields when the specified NSID is set to FFFFFFFFh
 - The requirement for an Identify command with CNS values that returned the capabilities in a namespace data structure if the NSID specified is FFFFFFFFh to abort the command if the Namespace Management has been changed to be optional.
- References:
 - Sections 4.1.2, 4.1.5, 4.1.5.2, and 4.1.5.TBD

Description for NVM Express Base Specification 2.0 Changes Document

Feature Enhancement:

- Added CNS value **TBDh** to the Identify command to return an Identify Namespace data structure containing the capability fields per the specified User Data Formal index. The capability fields are used to create and format namespaces.
- Added CNS value **TBD+1h** to the Identify command to return an Identify I/O Command Set specific Namespace data structure containing the capability fields per the specified User Data Formal index and Command Set Identifier. The capability fields are used to create and format namespaces.
- **New requirement / incompatible change:**
 - The Identify command for CNS values 05h and 1Dh removed the term “common” and clearly specified the requirements on all fields when the specified NSID is set to FFFFFFFFh.
 - The requirement for an Identify command with CNS values that returned the capabilities in a namespace data structure if the NSID specified is FFFFFFFFh to abort the command if the Namespace Management has been changed to be optional.
- References:
 - Sections 5.7, 5.17.1, 5.17.2.5, 5.17.2.21

Markup Conventions:

Black:	Unchanged (however, hot links are removed)
Red Strikethrough:	Deleted
Blue:	New
Blue Highlighted:	TBD values, anchors, and links to be inserted in new text.
<Green Bracketed>:	Notes to editor

Modify portions of NVM Command Set Specification 1.0 (used the member review version dated 3/4/2021) as shown below:

Modify section 4.1.2 as shown below:

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 (refer to section 5.TBD). Other NVM Command Set specific fields are defined in Figure 78.

...

Modify a portion of figure 86 in section 4.1.4.3 as shown below:

4.1.3.4 Host Behavior Support (Feature Identifier 16h)

The Host Behavior Support feature operates as defined in the NVMe Base Specification. In addition to the requirements in the NVMe Base Specification, this specification provides NVM Command Set specific definitions.

Figure 86: Host Behavior Support – Data Structure

Bytes	Description
02	<p>LBA Format Extension Enable (LBAFEE): This field allows the host to specify support for the extended LBA formats (refer to the EBLAS field in the Identify Controller data structure in the NVMe Base Specification). If this field is set to 1h and the ELBAS field is set to '1', then the controller:</p> <ol style="list-style-type: none"> 1) shall report a maximum number that is less than or equal to 64 for: <ol style="list-style-type: none"> a. the number of LBA formats supported (refer to the NLBAF field and the NULBAF field in the Identify Namespace data structure in Figure 97); and b. the number of namespace granularity descriptors (refer to Figure 104); and 2) is enabled to create, format, and perform I/O commands on namespaces formatted with (refer to section 5.2.1): <ol style="list-style-type: none"> a. 16b Guard Protection Information with the STS field set to a non-zero value; b. 32b Guard Protection Information; and c. 64b Guard Protection Information, where the extended LBA formats (refer to Figure 101) define the actual protection information formats supported. <p>If this field is cleared to 0h, then the controller:</p> <ol style="list-style-type: none"> 1) shall report a maximum that is less than or equal to 16 for: <ol style="list-style-type: none"> a. the number of LBA formats supported; and b. the number of namespace granularity descriptors; 2) shall not create, format, and perform I/O commands on namespaces formatted with (refer to section 5.2.1): <ol style="list-style-type: none"> a. 16b Guard Protection Information with the STS field set to a non-zero value; b. 32b Guard Protection Information; and c. 64b Guard Protection Information, and commands requesting these restrictions shall be aborted with a status code of Invalid Namespace or Format. <p>All values other than 0h and 1h are reserved.</p>

Modify figure 96 in section 4.1.5 as shown below:

4.1.5 Identify Command

...

Figure 96: Identify – CNS Values

CNS Value	O/M ¹	Definition	NSID ²	CNTID ³	CSI ⁴	Reference Section
Active Namespace Management						
...						

CNS Value	O/M ¹	Definition	NSID ²	CNTID ³	CSI ⁴	Reference Section
06h	M	Identify I/O Command Set specific Controller data structure for the controller processing the command. ⁶	Y	N	Y	4.1.5.4
09h	O	Identify Namespace data structure for the specified LBA Format index containing the namespace capabilities for the NVM Command Set. ⁶	N	N	Y	4.1.5.TBD
0Ah	O	I/O Command Set specific Identify Namespace data structure for the specified LBA Format index for the I/O Command Set specified in the CSI field. ⁶	N	N	Y	4.1.5.TBD1
11h	O	Identify Namespace data structure for the specified allocated NSID.	Y	N	N	4.1.5.5
...						
NOTES:						
<ol style="list-style-type: none"> O/M definition: O = Optional, M = Mandatory. The NSID field is used: Y = Yes, N = No. The CDW10.CNTID field is used: Y = Yes, N = No. The CDW11.CSI field is used: Y = Yes, N = No. Mandatory for controllers that support the Namespace Management capability (refer to the NVMe Base Specification). Selection of a UUID may be supported. Refer to the Universally Unique Identifiers (UUIDs) for Vendor Specific Information section in the NVMe Base Specification. 						

Modify portions of section 4.1.5.1 as shown below:

4.1.5.1 NVM Command Set Identify Namespace data structure (CNS 00h)

If the Namespace Identifier (NSID) field specifies an active NSID, then the NVM Command Set Identify Namespace data structure (refer to Figure 97) is returned to the host for that specified namespace. If that specified namespace is an inactive NSID, then the controller returns a zero filled data structure. If the specified namespace is not associated with an I/O Command Set that supports this data structure, then the controller shall abort the command with the status code of Invalid I/O Command Set.

~~If the controller supports the Namespace Management capability (refer to the Namespace Management section in the NVMe Base Specification) and the NSID field is set to FFFFFFFFh, then the controller returns an Identify Namespace data structure that specifies NVM Command Set capabilities that are common across namespaces for the controller. If the controller does not support the Namespace Management capability and the NSID field is set to FFFFFFFFh, then the controller shall abort the command with a status code of Invalid Namespace or Format.~~

Figure 97 specifies fields in the NVM Command Set Identify Namespace data structure that define capabilities used by a host to format or create a namespace. If the NSID field is set to FFFFFFFFh, then the controller shall return an NVM Command Set Identify Namespace data structure that has:

- the capability fields as specified by Figure 97 that requires a value that is the same value for all namespaces formatted using the LBA formats associated with the Number of LBA Formats field (refer to section 5.TBD); and
- fields that are not capability fields set to the value 0h.

If the controller supports the Namespace Management capability (refer to the Namespace Management section in the NVMe Base Specification) and the NSID field is set to FFFFFFFFh, then the controller shall return an NVM Command Set Identify Namespace data structure. If the controller does not support the Namespace Management capability and the NSID field is set to FFFFFFFFh, then the controller may abort the command with a status code of Invalid Namespace or Format.

<Note to readers: This is about adding the right most column

Figure 97: Identify – Identify Namespace Data Structure, NVM Command Set

Bytes	O/M 1	Description	Capability Field
07:00	M	Namespace Size (NSZE): This field indicates the total size of the namespace in logical blocks. A namespace of size n consists of LBA 0 through $(n - 1)$. The number of logical blocks is based on the formatted LBA size.	No
15:08	M	Namespace Capacity (NCAP): This field indicates the maximum number of logical blocks that may be allocated in the namespace at any point in time. The number of logical blocks is based on the formatted LBA size. Spare LBAs are not reported as part of this field. Refer to section 2.1.1 for details on the usage of this field.	No
23:16	M	Namespace Utilization (NUSE): This field indicates the current number of logical blocks allocated in the namespace. This field is smaller than or equal to the Namespace Capacity. The number of logical blocks is based on the formatted LBA size. Refer to section 2.1.1 for details on the usage of this field.	No
24	M	Namespace Features (NSFEAT): This field defines features of the namespace. Bits 7:5 are reserved. Bit 4 (OPTPERF) if set to '1' indicates that the fields NPWG, NPWA, NPDG, NPDA, and NOWS are defined for this namespace and should be used by the host for I/O optimization (refer to the NVM Set List section in the NVMe Base Specification). If cleared to '0', then the controller does not support the fields NPWG, NPWA, NPDG, NPDA, and NOWS for this namespace. Bit 3 (UIDREUSE) This bit is as defined in the UIDREUSE bit in the I/O Command Set Independent Identify Namespace data structure (refer to the I/O Command Set Independent Identify Namespace data structure section in the NVMe Base Specification). Bit 2 (DAE) if set to '1' indicates that the controller supports the Deallocated or Unwritten Logical Block error for this namespace. If cleared to '0', then the controller does not support the Deallocated or Unwritten Logical Block error for this namespace. Refer to section 3.2.3.2.1. Bit 1 (NSABP) if set to '1' indicates that the fields NAWUN, NAWUPF, and NACWU are defined for this namespace and should be used by the host for this namespace instead of the AWUN, AWUPF, and ACWU fields in the Identify Controller data structure. If cleared to '0', then the controller does not support the fields NAWUN, NAWUPF, and NACWU for this namespace. In this case, the host should use the AWUN, AWUPF, and ACWU fields defined in the Identify Controller data structure in the NVMe Base Specification. Refer to section 2.1.4. Bit 0 (THINP) if set to '1' indicates that the namespace supports thin provisioning. If cleared to '0' indicates that thin provisioning is not supported Refer to section 2.1.1 for details on the usage of this bit.	No

Figure 97: Identify – Identify Namespace Data Structure, NVM Command Set

Bytes	O/M 1	Description	Capability Field
25	M	<p>Number of LBA Formats (NLBAF): This field defines the number of supported LBA data size and metadata size combinations supported by the namespace that share the same capabilities. LBA formats shall be allocated in order (starting with 0) and packed sequentially. This is a 0's based value.</p> <p>Refer to section 5.TBD for the structure of the LBA formats, the association to the NULBAF field, and the maximum values of this field.</p> <p>The maximum number of LBA formats that may be indicated as supported is:</p> <ul style="list-style-type: none"> • 16 if the host has cleared the LBA Format Extension Enable (LBAFEE) field to 0h in the Host Behavior Support feature (refer to the Host Behavior Support section in the NVMe Base Specification); or • 64 if the host has set the LBAFEE field to 1h in the Host Behavior Support feature (refer to the Host Behavior Support section in the NVMe Base Specification). <p>The supported LBA formats are indicated in bytes 128 to 383 in this data structure. The LBA Format fields with an index beyond the value defined by section 5.TBD set in this field are invalid and not supported. LBA Formats that are valid, but not currently available may be indicated by setting the LBA Data Size for that LBA Format to 0h.</p> <p>The metadata may be either transferred as part of the LBA (creating an extended LBA which is a larger LBA size that is exposed to the application) or may be transferred as a separate contiguous buffer of data. The metadata shall not be split between the LBA and a separate metadata buffer.</p> <p>It is recommended that software and controllers transition to an LBA size that is 4 KiB or larger for ECC efficiency at the controller. If providing metadata, it is recommended that at least 8 bytes are provided per logical block to enable use with end-to-end data protection, refer to section 5.8.3.</p>	Yes
26	M	<p>Formatted LBA Size (FLBAS): This field indicates the LBA data size & metadata size combination that the namespace has been formatted with (refer to section 4.1.2).</p> <p>Bits 7 is reserved.</p> <p>Bits 6:5 indicate the most significant 2 bits of the supported LBA Formats indicated in this data structure. These bits are ignored if the number of supported LBA Formats is less than or equal to 16.</p> <p>Bit 4 if set to '1' indicates that the metadata is transferred at the end of the data LBA, creating an extended data LBA. Bit 4 if cleared to '0' indicates that all of the metadata for a command is transferred as a separate contiguous buffer of data. Bit 4 is not applicable when there is no metadata.</p> <p>Bits 3:0 indicates one of the 16 supported LBA Formats indicated in this data structure.</p>	No

Figure 97: Identify – Identify Namespace Data Structure, NVM Command Set

Bytes	O/M 1	Description	Capability Field														
27	M	<p>Metadata Capabilities (MC): This field indicates the capabilities for metadata.</p> <p>Bits 7:2 are reserved.</p> <p>Bit 1 if set to '1' indicates the namespace supports the metadata being transferred as part of a separate buffer that is specified in the Metadata Pointer. Bit 1 if cleared to '0' indicates that the namespace does not support the metadata being transferred as part of a separate buffer.</p> <p>Bit 0 if set to '1' indicates that the namespace supports the metadata being transferred as part of an extended data LBA. Bit 0 if cleared to '0' indicates that the namespace does not support the metadata being transferred as part of an extended data LBA.</p>	Yes														
28	M	<p>End-to-end Data Protection Capabilities (DPC): This field indicates the capabilities for the end-to-end data protection feature. Multiple bits may be set in this field. Refer to section 5.1.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:5</td> <td>Reserved</td> </tr> <tr> <td>4</td> <td>Protection Information In Last Bytes (PIILB): If set to '1' indicates that the namespace supports protection information transferred as the last bytes of metadata. If cleared to '0' indicates that the namespace does not support protection information transferred as the last bytes of metadata.</td> </tr> <tr> <td>3</td> <td>Protection Information In First Bytes (PIIFB): If set to '1' indicates that the namespace supports protection information transferred as the first bytes of metadata. If cleared to '0' indicates that the namespace does not support protection information transferred as the first bytes of metadata. For versions later than version 1.4 of this specification, this bit shall be cleared to '0'.</td> </tr> <tr> <td>2</td> <td>Protection Information Type 3 Supported (PIT3S): If set to '1' indicates that the namespace supports Protection Information Type 3. If cleared to '0' indicates that the namespace does not support Protection Information Type 3.</td> </tr> <tr> <td>1</td> <td>Protection Information Type 2 Supported (PIT2S): If set to '1' indicates that the namespace supports Protection Information Type 2. If cleared to '0' indicates that the namespace does not support Protection Information Type 2.</td> </tr> <tr> <td>0</td> <td>Protection Information Type 1 Supported (PIT1S): If set to '1' indicates that the namespace supports Protection Information Type 1. If cleared to '0' indicates that the namespace does not support Protection Information Type 1.</td> </tr> </tbody> </table>	Bits	Description	7:5	Reserved	4	Protection Information In Last Bytes (PIILB): If set to '1' indicates that the namespace supports protection information transferred as the last bytes of metadata. If cleared to '0' indicates that the namespace does not support protection information transferred as the last bytes of metadata.	3	Protection Information In First Bytes (PIIFB): If set to '1' indicates that the namespace supports protection information transferred as the first bytes of metadata. If cleared to '0' indicates that the namespace does not support protection information transferred as the first bytes of metadata. For versions later than version 1.4 of this specification, this bit shall be cleared to '0'.	2	Protection Information Type 3 Supported (PIT3S): If set to '1' indicates that the namespace supports Protection Information Type 3. If cleared to '0' indicates that the namespace does not support Protection Information Type 3.	1	Protection Information Type 2 Supported (PIT2S): If set to '1' indicates that the namespace supports Protection Information Type 2. If cleared to '0' indicates that the namespace does not support Protection Information Type 2.	0	Protection Information Type 1 Supported (PIT1S): If set to '1' indicates that the namespace supports Protection Information Type 1. If cleared to '0' indicates that the namespace does not support Protection Information Type 1.	Yes
Bits	Description																
7:5	Reserved																
4	Protection Information In Last Bytes (PIILB): If set to '1' indicates that the namespace supports protection information transferred as the last bytes of metadata. If cleared to '0' indicates that the namespace does not support protection information transferred as the last bytes of metadata.																
3	Protection Information In First Bytes (PIIFB): If set to '1' indicates that the namespace supports protection information transferred as the first bytes of metadata. If cleared to '0' indicates that the namespace does not support protection information transferred as the first bytes of metadata. For versions later than version 1.4 of this specification, this bit shall be cleared to '0'.																
2	Protection Information Type 3 Supported (PIT3S): If set to '1' indicates that the namespace supports Protection Information Type 3. If cleared to '0' indicates that the namespace does not support Protection Information Type 3.																
1	Protection Information Type 2 Supported (PIT2S): If set to '1' indicates that the namespace supports Protection Information Type 2. If cleared to '0' indicates that the namespace does not support Protection Information Type 2.																
0	Protection Information Type 1 Supported (PIT1S): If set to '1' indicates that the namespace supports Protection Information Type 1. If cleared to '0' indicates that the namespace does not support Protection Information Type 1.																

Figure 97: Identify – Identify Namespace Data Structure, NVM Command Set

Bytes	O/M 1	Description	Capability Field																				
29	M	<p>End-to-end Data Protection Type Settings (DPS): This field indicates the protection information Type settings for the end-to-end data protection feature. Refer to section 5.1.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:4</td> <td>Reserved</td> </tr> <tr> <td>3</td> <td>Protection Information Position (PIP): This bit indicates that the protection information, if enabled, is transferred as the first bytes of metadata. Bit 3 if cleared to '0' indicates that the protection information, if enabled, is transferred as the last bytes of metadata. For versions later than version 1.4 of this specification this bit shall be cleared to '0'.</td> </tr> <tr> <td>2:0</td> <td> <p>Protection Information Type (PIT): This field indicates whether protection information is enabled and the type of protection information enabled. The values for this field have the following meanings:</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Protection information is not enabled</td> </tr> <tr> <td>001b</td> <td>Protection information is enabled, Type 1</td> </tr> <tr> <td>010b</td> <td>Protection information is enabled, Type 2</td> </tr> <tr> <td>011b</td> <td>Protection information is enabled, Type 3</td> </tr> <tr> <td>100b to 111b</td> <td>Reserved</td> </tr> </tbody> </table> </td> </tr> </tbody> </table>	Bits	Description	7:4	Reserved	3	Protection Information Position (PIP): This bit indicates that the protection information, if enabled, is transferred as the first bytes of metadata. Bit 3 if cleared to '0' indicates that the protection information, if enabled, is transferred as the last bytes of metadata. For versions later than version 1.4 of this specification this bit shall be cleared to '0'.	2:0	<p>Protection Information Type (PIT): This field indicates whether protection information is enabled and the type of protection information enabled. The values for this field have the following meanings:</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Protection information is not enabled</td> </tr> <tr> <td>001b</td> <td>Protection information is enabled, Type 1</td> </tr> <tr> <td>010b</td> <td>Protection information is enabled, Type 2</td> </tr> <tr> <td>011b</td> <td>Protection information is enabled, Type 3</td> </tr> <tr> <td>100b to 111b</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Definition	000b	Protection information is not enabled	001b	Protection information is enabled, Type 1	010b	Protection information is enabled, Type 2	011b	Protection information is enabled, Type 3	100b to 111b	Reserved	No
Bits	Description																						
7:4	Reserved																						
3	Protection Information Position (PIP): This bit indicates that the protection information, if enabled, is transferred as the first bytes of metadata. Bit 3 if cleared to '0' indicates that the protection information, if enabled, is transferred as the last bytes of metadata. For versions later than version 1.4 of this specification this bit shall be cleared to '0'.																						
2:0	<p>Protection Information Type (PIT): This field indicates whether protection information is enabled and the type of protection information enabled. The values for this field have the following meanings:</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Protection information is not enabled</td> </tr> <tr> <td>001b</td> <td>Protection information is enabled, Type 1</td> </tr> <tr> <td>010b</td> <td>Protection information is enabled, Type 2</td> </tr> <tr> <td>011b</td> <td>Protection information is enabled, Type 3</td> </tr> <tr> <td>100b to 111b</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Definition	000b	Protection information is not enabled	001b	Protection information is enabled, Type 1	010b	Protection information is enabled, Type 2	011b	Protection information is enabled, Type 3	100b to 111b	Reserved										
Value	Definition																						
000b	Protection information is not enabled																						
001b	Protection information is enabled, Type 1																						
010b	Protection information is enabled, Type 2																						
011b	Protection information is enabled, Type 3																						
100b to 111b	Reserved																						
30	O	Namespace Multi-path I/O and Namespace Sharing Capabilities (NMIC): This field is as defined in the I/O Command Set Independent Identify Namespace data structure (refer to the I/O Command Set Independent Identify Namespace data structure section in the NVMe Base Specification).	No																				
31	O	Reservation Capabilities (RESCAP): This field is as defined in the I/O Command Set Independent Identify Namespace data structure (refer to the I/O Command Set Independent Identify Namespace data structure section in the NVMe Base Specification).	No																				
32	O	Format Progress Indicator (FPI): This field is as defined in the I/O Command Set Independent Identify Namespace data structure (refer to the I/O Command Set Independent Identify Namespace data structure section in the NVMe Base Specification).	No																				

Figure 97: Identify – Identify Namespace Data Structure, NVM Command Set

Bytes	O/M 1	Description	Capability Field										
33	O	<p>Deallocate Logical Block Features (DLFEAT): This field indicates information about features that affect deallocating logical blocks for this namespace.</p> <p>Bits 7:5 are reserved.</p> <p>Bit 4 if set to '1' indicates that the Guard field for deallocated logical blocks that contain protection information is set to the CRC for the value read from the deallocated logical block and its metadata (excluding protection information). If cleared to '0' indicates that each byte in the Guard field for the deallocated logical blocks that contain protection information is set to FFh.</p> <p>Bit 3 if set to '1' indicates that the controller supports the Deallocate bit in the Write Zeroes command for this namespace. If cleared to '0' indicates that the controller does not support the Deallocate bit in the Write Zeroes command for this namespace. This bit shall be set to the same value for all namespaces in the NVM subsystem.</p> <p>Bits 2:0 indicate deallocated logical block read behavior. For a logical block that is deallocated, this field indicates the values read from that deallocated logical block and its metadata (excluding protection information). The values for this field have the following meanings:</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>The read behavior is not reported</td> </tr> <tr> <td>001b</td> <td>A deallocated logical block returns all bytes cleared to 0h</td> </tr> <tr> <td>010b</td> <td>A deallocated logical block returns all bytes set to FFh</td> </tr> <tr> <td>011b to 111b</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Definition	000b	The read behavior is not reported	001b	A deallocated logical block returns all bytes cleared to 0h	010b	A deallocated logical block returns all bytes set to FFh	011b to 111b	Reserved	No
Value	Definition												
000b	The read behavior is not reported												
001b	A deallocated logical block returns all bytes cleared to 0h												
010b	A deallocated logical block returns all bytes set to FFh												
011b to 111b	Reserved												
35:34	O	<p>Namespace Atomic Write Unit Normal (NAWUN): This field indicates the namespace specific size of the write operation guaranteed to be written atomically to the NVM during normal operation. If the NSABP bit is cleared to '0', then this field is reserved.</p> <p>A value of 0h indicates that the size for this namespace is the same size as that reported in the AWUN field of the Identify Controller data structure. All other values specify a size in terms of logical blocks using the same encoding as the AWUN field. Refer to section 2.1.4.</p>	No										
37:36	O	<p>Namespace Atomic Write Unit Power Fail (NAWUPF): This field indicates the namespace specific size of the write operation guaranteed to be written atomically to the NVM during a power fail or error condition. If the NSABP bit is cleared to '0', then this field is reserved.</p> <p>A value of 0h indicates that the size for this namespace is the same size as that reported in the AWUPF field of the Identify Controller data structure. All other values specify a size in terms of logical blocks using the same encoding as the AWUPF field. Refer to section 2.1.4.</p>	No										
39:38	O	<p>Namespace Atomic Compare & Write Unit (NACWU): This field indicates the namespace specific size of the write operation guaranteed to be written atomically to the NVM for a Compare and Write fused command. If the NSABP bit is cleared to '0', then this field is reserved.</p> <p>A value of 0h indicates that the size for this namespace is the same size as that reported in the ACWU field of the Identify Controller data structure. All other values specify a size in terms of logical blocks using the same encoding as the ACWU field. Refer to section 2.1.4.</p>	No										

Figure 97: Identify – Identify Namespace Data Structure, NVM Command Set

Bytes	O/M 1	Description	Capability Field
41:40	O	<p>Namespace Atomic Boundary Size Normal (NABSN): This field indicates the atomic boundary size for this namespace for the NAWUN value. This field is specified in logical blocks. Writes to this namespace that cross atomic boundaries are not guaranteed to be atomic to the NVM with respect to other read or write commands.</p> <p>A value of 0h indicates that there are no atomic boundaries for normal write operations. All other values specify a size in terms of logical blocks using the same encoding as the AWUN field. Refer to section 2.1.4.</p> <p>Refer to section 5.8.2 for how this field is utilized.</p>	No
43:42	O	<p>Namespace Atomic Boundary Offset (NABO): This field indicates the LBA on this namespace where the first atomic boundary starts.</p> <p>If the NABSN and NABSPF fields are cleared to 0h, then the NABO field shall be cleared to 0h. NABO shall be less than or equal to NABSN and NABSPF. Refer to section 2.1.4.</p> <p>Refer to section 5.8.2 for how this field is utilized.</p>	No
45:44	O	<p>Namespace Atomic Boundary Size Power Fail (NABSPF): This field indicates the atomic boundary size for this namespace specific to the Namespace Atomic Write Unit Power Fail value. This field is specified in logical blocks. Writes to this namespace that cross atomic boundaries are not guaranteed to be atomic with respect to other read or write commands and there is no guarantee of data returned on subsequent reads of the associated logical blocks.</p> <p>A value of 0h indicates that there are no atomic boundaries for power fail or error conditions. All other values specify a size in terms of logical blocks using the same encoding as the AWUPF field. Refer to section 2.1.4.</p>	No
47:46	O	<p>Namespace Optimal I/O Boundary (NOIOB): This field indicates the optimal I/O boundary for this namespace. This field is specified in logical blocks. The host should construct Read and Write commands that do not cross the I/O boundary to achieve optimal performance. A value of 0h indicates that no optimal I/O boundary is reported.</p> <p>Refer to section 5.8.2 for how this field is utilized to improve performance and endurance.</p>	No
63:48	O	<p>NVM Capacity (NVMCAP): This field indicates the total size of the NVM allocated to this namespace. The value is in bytes. This field shall be supported if the Namespace Management capability (refer to section 5.3) is supported.</p> <p>Note: This field may not correspond to the logical block size multiplied by the Namespace Size field. Due to thin provisioning or other settings (e.g., endurance), this field may be larger or smaller than the product of the logical block size and the Namespace Size reported.</p> <p>If the controller supports Asymmetric Namespace Access Reporting (refer to the CMIC field), and the relationship between the controller and the namespace is in the ANA Inaccessible state (refer to the ANA Inaccessible state section in the NVMe Base Specification) or the ANA Persistent Loss state (refer to the ANA Persistent Loss state section in the NVMe Base Specification), then this field shall be cleared to 0h.</p>	No

Figure 97: Identify – Identify Namespace Data Structure, NVM Command Set

Bytes	O/M 1	Description	Capability Field
65:64	O	<p>Namespace Preferred Write Granularity (NPWG): This field indicates the smallest recommended write granularity in logical blocks for this namespace. This is a 0's based value. If the NSABP bit is cleared to '0', then this field is reserved.</p> <p>The size indicated should be less than or equal to Maximum Data Transfer Size (MDTS) that is specified in units of minimum memory page size. The value of this field may change if the namespace is reformatted. The size should be a multiple of Namespace Preferred Write Alignment (NPWA).</p> <p>Refer to section 5.8.2 for how this field is utilized to improve performance and endurance.</p>	No
67:66	O	<p>Namespace Preferred Write Alignment (NPWA): This field indicates the recommended write alignment in logical blocks for this namespace. This is a 0's based value. If the NSABP bit is cleared to '0', then this field is reserved.</p> <p>The value of this field may change if the namespace is reformatted.</p> <p>Refer to section 5.8.2 for how this field is utilized to improve performance and endurance.</p>	No
69:68	O	<p>Namespace Preferred Deallocate Granularity (NPDG): This field indicates the recommended granularity in logical blocks for the Dataset Management command with the Attribute – Deallocate bit set to '1' in Dword 11. This is a 0's based value. If the NSABP bit is cleared to '0', then this field is reserved.</p> <p>The value of this field may change if the namespace is reformatted. The size should be a multiple of Namespace Preferred Deallocate Alignment (NPDA).</p> <p>Refer to section 5.8.2 for how this field is utilized to improve performance and endurance.</p>	No
71:70	O	<p>Namespace Preferred Deallocate Alignment (NPDA): This field indicates the recommended alignment in logical blocks for the Dataset Management command with the Attribute – Deallocate bit set to '1' in Dword 11. This is a 0's based value. If the NSABP bit is cleared to '0', then this field is reserved.</p> <p>The value of this field may change if the namespace is reformatted.</p> <p>Refer to section 5.8.2 for how this field is utilized to improve performance and endurance.</p>	No
73:72	O	<p>Namespace Optimal Write Size (NOWS): This field indicates the size in logical blocks for optimal write performance for this namespace. This is a 0's based value. If the NSABP bit is cleared to '0', then this field is reserved.</p> <p>The size indicated should be less than or equal to Maximum Data Transfer Size (MDTS) that is specified in units of minimum memory page size. The value of this field may change if the namespace is reformatted. The value of this field should be a multiple of Namespace Preferred Write Granularity (NPWG).</p> <p>If the namespace is associated with an NVM set, NOWS defined for this namespace shall be set to the Optimal Write Size field setting defined in NVM Set Attributes Entry (refer to the Namespace Identification Descriptor in the NVMe Base Specification) for the NVM Set with which this namespace is associated. If NOWS is not supported, the Optimal Write Size field in NVM Sets Attributes Entry (refer to the Namespace Identification Descriptor in the NVMe Base Specification) for the NVM Set with which this namespace is associated should be used by the host for I/O optimization.</p> <p>Refer to section 5.8.2 for how this field is utilized to improve performance and endurance.</p>	No

Figure 97: Identify – Identify Namespace Data Structure, NVM Command Set

Bytes	O/M 1	Description	Capability Field
75:74	O	Maximum Single Source Range Length (MSSRL): This field indicates the maximum number of logical blocks that may be specified in each valid Source Range field (refer to 3) of a Copy command. If the controller supports the Copy command, then this field shall be set to a non-zero value.	No
79:76	O	Maximum Copy Length (MCL): This field indicates the maximum number of logical blocks that may be specified in a Copy command (i.e., the sum of the number of logical blocks specified in all Source Range entries). If the controller supports the Copy command, then this field shall be set to a non-zero value.	No
80	O	Maximum Source Range Count (MSRC): This field indicates the maximum number of Source Range entries that may be used to specify source data in a Copy command. This is a 0's based value.	No
82	M	Number of Unique Capability LBA Formats (NULBAF): This field defines the number of supported user data size and metadata size combinations supported by the namespace that may not share the same capabilities. LBA formats shall be allocated in order (starting at the first index after the LBA formats defined by the NLBAF field) and packed sequentially (refer to section 5.TBD). Refer to section 5.TBD for the structure of the LBA formats, the association to the NLBAF field, and the maximum value of this field.	Yes
...
95:92	O	ANA Group Identifier (ANAGRPID): This field is as defined in the I/O Command Set Independent Identify Namespace data structure (refer to the I/O Command Set Independent Identify Namespace data structure section in the NVMe Base Specification).	No
98:96		Reserved	
99	O	Namespace Attributes (NSATTR): This field is as defined in the I/O Command Set Independent Identify Namespace data structure (refer to the I/O Command Set Independent Identify Namespace data structure section in the NVMe Base Specification).	No
101:100	O	NVM Set Identifier (NVMSETID): This field is as defined in the I/O Command Set Independent Identify Namespace data structure (refer to the I/O Command Set Independent Identify Namespace data structure section in the NVMe Base Specification).	No
103:102	O	Endurance Group Identifier (ENDGID): This field is as defined in the I/O Command Set Independent Identify Namespace data structure (refer to the I/O Command Set Independent Identify Namespace data structure section in the NVMe Base Specification).	No

Figure 97: Identify – Identify Namespace Data Structure, NVM Command Set

Bytes	O/M 1	Description	Capability Field
119:104	O	<p>Namespace Globally Unique Identifier (NGUID): This field contains a 128-bit value that is globally unique and assigned to the namespace when the namespace is created. This field remains fixed throughout the life of the namespace and is preserved across namespace and controller operations (e.g., Controller Level Reset, namespace format, etc.).</p> <p>This field uses the EUI-64 based 16-byte designator format. Bytes 114:112 contain the 24-bit Organizationally Unique Identifier (OUI) value assigned by the IEEE Registration Authority. Bytes 119:115 contain an extension identifier assigned by the corresponding organization. Bytes 111:104 contain the vendor specific extension identifier assigned by the corresponding organization. Refer to the IEEE EUI-64 guidelines for more information. This field is big endian (refer to the Namespace Globally Unique Identifier section in the NVMe Base Specification).</p> <p>The controller shall specify a globally unique namespace identifier in this field, the EUI64 field, or a Namespace UUID in the Namespace Identification Descriptor (refer to the Namespace Identification Descriptor figure in the NVMe Base Specification) when the namespace is created. If the controller is not able to provide a globally unique identifier in this field, then this field shall be cleared to 0h. Refer to the Unique Identifier section in the NVMe Base Specification.</p>	No
127:120	O	<p>IEEE Extended Unique Identifier (EUI64): This field contains a 64-bit IEEE Extended Unique Identifier (EUI-64) that is globally unique and assigned to the namespace when the namespace is created. This field remains fixed throughout the life of the namespace and is preserved across namespace and controller operations (e.g., Controller Level Reset, namespace format, etc.).</p> <p>The EUI-64 is a concatenation of a 24-bit or 36-bit Organizationally Unique Identifier (OUI or OUI-36) value assigned by the IEEE Registration Authority and an extension identifier assigned by the corresponding organization. Refer to the IEEE EUI-64 guidelines for more information. This field is big endian (refer to the IEEE Extended Unique Identifier section in the NVMe Base Specification).</p> <p>The controller shall specify a globally unique namespace identifier in this field, the NGUID field, or a Namespace UUID in the Namespace Identification Descriptor (refer to the Namespace Identification Descriptor figure in the NVMe Base Specification) when the namespace is created. If the controller is not able to provide a globally unique 64-bit identifier in this field, then this field shall be cleared to 0h. Refer to the Unique Identifier section in the NVMe Base Specification.</p>	No
LBA Format List (refer to section 5.TBD)			
131:128	M	<p>LBA Format 0 Support (LBAF0): This field indicates the LBA format 0 that is supported by the controller. The LBA format field is defined in Figure 98.</p> <p>Additional information may be provided in the ELBAF0 field (refer to Figure 100).</p>	Yes
135:132	O	<p>LBA Format 1 Support (LBAF1): This field indicates the LBA format 1 that is supported by the controller. The LBA format field is defined in Figure 98.</p> <p>Additional information may be provided in the ELBAF1 field (refer to Figure 100).</p>	Yes
...			
383:380	O	<p>LBA Format 63 Support (LBAF63): This field indicates the LBA format 63 that is supported by the controller. The LBA format field is defined in Figure 98.</p> <p>Additional information may be provided in the ELBAF63 field (refer to Figure 100).</p>	Yes

Figure 97: Identify – Identify Namespace Data Structure, NVM Command Set

Bytes	O/M ¹	Description	Capability Field
383:192		Reserved	
4095:384	O	Vendor Specific	No
NOTES:			
1. O/M definition: O = Optional, M = Mandatory.			

Modify portions of section 4.1.5.3 as shown below:

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

Figure 100 defines the I/O Command Set specific Identify Namespace data structure for the NVM Command Set.

Figure 100 specifies fields in the NVM Command Set I/O Command Set Specific Identify Namespace data structure that define capabilities used by a host to format or create a namespace. If the NSID field is set to FFFFFFFFh, then the controller shall return an NVM Command Set I/O Command Set Specific Identify Namespace data structure that has:

- the capability fields as specified by Figure 100 that requires a value that is the same value for all namespaces formatted using the LBA formats associated with the Number of LBA Formats field; and
- fields that are not capability fields set to the value 0h.

If the controller supports the Namespace Management capability (refer to the Namespace Management section in the NVMe Base Specification) and the NSID field is set to FFFFFFFFh, then the controller shall return an NVM Command Set I/O Command Set Specific Identify Namespace data structure. If the controller does not support the Namespace Management capability and the NSID field is set to FFFFFFFFh, then the controller may abort the command with a status code of Invalid Namespace or Format.

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

Bytes	O/M ¹	Description	Capabilities Field
7:0	O	<p>Logical Block Storage Tag Mask (LBSTM): Identifies the mask for the Storage Tag field for the protection information (refer to section 5.1). The size of this field is defined by the STS field. If the size of this field is less than 64 bits, the mask is contained in the least significant bits of this field.</p> <p>If end-to-end protection is not enabled in the namespace, then this field is ignored.</p> <p>If:</p> <ul style="list-style-type: none"> • end-to-end protection is enabled; • 16b Guard Protection Information format is used; and • The 16BPISTM bit is set to '1', <p>then all bits in the mask shall be set to '1'.</p>	No

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

Bytes	O/M ¹	Description	Capabilities Field								
8	O	Protection Information Capabilities (PIC): This field indicates the capabilities for the protection information formats.	Yes								
		<table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:2</td> <td>Reserved</td> </tr> <tr> <td>1</td> <td>16b Guard Protection Information Storage Tag Mask (16BPISM): If set to '1', then the LBSTM field shall have all bits set to '1' for the 16b Guard Protection Information. If cleared to '0' then the Logical Block Storage Tag Mask field is allowed to have any bits set to '1' for the 16b Guard Protection Information.</td> </tr> <tr> <td>0</td> <td>16b Guard Protection Information Storage Tag Support (16BPISTS): If set to '1', then the end-to-end protection 16b Guard Protection Information format (refer to section 5.2.1.1) supports a non-zero value in the STS field. If cleared to '0', then the end-to-end protection 16b Guard Protection Information format support requires that the STS field be cleared to 0h (i.e., the Storage Tag field is not supported). If the 32b Guard Protection Information or 64b Guard Protection Information is supported in any LBA format (refer to Figure 97 and Figure 100), then this bit shall be set to '1'.</td> </tr> </tbody> </table>		Bits	Description	7:2	Reserved	1	16b Guard Protection Information Storage Tag Mask (16BPISM): If set to '1', then the LBSTM field shall have all bits set to '1' for the 16b Guard Protection Information. If cleared to '0' then the Logical Block Storage Tag Mask field is allowed to have any bits set to '1' for the 16b Guard Protection Information.	0	16b Guard Protection Information Storage Tag Support (16BPISTS): If set to '1', then the end-to-end protection 16b Guard Protection Information format (refer to section 5.2.1.1) supports a non-zero value in the STS field. If cleared to '0', then the end-to-end protection 16b Guard Protection Information format support requires that the STS field be cleared to 0h (i.e., the Storage Tag field is not supported). If the 32b Guard Protection Information or 64b Guard Protection Information is supported in any LBA format (refer to Figure 97 and Figure 100), then this bit shall be set to '1'.
		Bits		Description							
		7:2		Reserved							
1	16b Guard Protection Information Storage Tag Mask (16BPISM): If set to '1', then the LBSTM field shall have all bits set to '1' for the 16b Guard Protection Information. If cleared to '0' then the Logical Block Storage Tag Mask field is allowed to have any bits set to '1' for the 16b Guard Protection Information.										
0	16b Guard Protection Information Storage Tag Support (16BPISTS): If set to '1', then the end-to-end protection 16b Guard Protection Information format (refer to section 5.2.1.1) supports a non-zero value in the STS field. If cleared to '0', then the end-to-end protection 16b Guard Protection Information format support requires that the STS field be cleared to 0h (i.e., the Storage Tag field is not supported). If the 32b Guard Protection Information or 64b Guard Protection Information is supported in any LBA format (refer to Figure 97 and Figure 100), then this bit shall be set to '1'.										
11:9	Reserved										
LBA Format List (refer to section 5.TBD)											
15:12	O	Extended LBA Format 0 Support (ELBAF0): This field indicates additional LBA Format 0 information related to the LBA Format 0 Support (LBAF0) field in the Identify Namespace data structure. The Extended LBA format field is defined in Figure 101.	Yes								
19:16	O	Extended LBA Format 1 Support (ELBAF1): This field indicates additional LBA Format 1 information related to the LBA Format 1 Support (LBAF1) field in the Identify Namespace data structure. The Extended LBA format field is defined in Figure 101..	Yes								
...									
267:264	O	Extended LBA Format 63 Support (ELBAF63): This field indicates additional LBA Format 63 information related to the LBA Format 63 Support (LBAF63) field in the Identify Namespace data structure. The Extended LBA format field is defined in Figure 101.	Yes								
4095:268	O	Reserved									
NOTES:											
1. O/M definition: O = Optional, M = Mandatory.											

Add section 4.1.5.TBD as shown below:

4.1.5.TBD Identify I/O Command Set Specific Namespace data structure (CNS 09h, CSI 00h)

An NVM Command Set Identify Namespace data structure (refer to Figure 97) is returned to the host for the specified LBA Format Index specified by the CNS Specific Identifier field as defined in Figure CNS_TBDH. The returned NVM Command Set Identify Namespace data structure specifies fields that define capabilities used by a host to format or create a namespace. If the LBA Format index specified is

valid (refer to section 5.TBD), then the controller shall return an NVM Command Set Identify Namespace data structure that has:

- the capability fields as specified by the NVM Command Set Identify Namespace data structure that requires a value that is the same value for all namespaces formatted using the LBA formats associated with the Number of LBA Formats field; and
- fields that are not capability fields set to the value 0h.

Figure CNS_TBDH: Command Dword 11 – CNS Specific Identifier

Bits	Description
15:0	LBA Format Index: This field specifies the index into the LBA Format list identifying the LBA Format capabilities that are to be returned. Refer to section 5.TBD.

Add section 4.1.5.TBD1 as shown below:

4.1.5.TBD1 Identify I/O Command Set specific Namespace data structure (CNS 0Ah, CSI 00h)

An I/O Command Set specific Identify Namespace data structure for the NVM Command Set (refer to Figure 100) is returned to the host for the for the specified LBA Format Index specified by the CNS Specific Identifier field as defined in Figure CNS_TBDH. The returned I/O Command Set specific Identify Namespace data structure specifies fields that define capabilities used by a host to format or create a namespace. If the LBA Format index specified is valid (refer to section 5.TBD), then the controller shall return an I/O Command Set specific Identify Namespace data structure that has:

- the capability fields as specified by the I/O Command Set specific Identify Namespace data structure for the NVM Command Set that requires a value that is the same value for all namespaces formatted using the LBA formats associated with the Number of LBA Formats field; and
- fields that are not capability fields set to the value 0h.

Modify figure 103 in section 4.1.5.6 as shown below:

4.1.5.6 Namespace Granularity List (CNS 16h)

...

Figure 103: Namespace Granularity List

Bytes	Description						
03:00	<p>Namespace Granularity Attributes: This field indicates attributes of the Namespace Granularity List.</p> <p>Bits 31:1 are reserved.</p> <p>Bit 0 (Granularity Descriptor Mapping): If set to '1', then each valid namespace granularity descriptor applies to the LBA format having the same index and the Number of Descriptors field shall be equal to the Number of LBA Formats field in the Identify Namespace data structure (refer to Figure 97). If cleared to '0', then NG Descriptor 0 shall apply to all LBA formats and the Number of Descriptors field shall be cleared to 0h.</p>						
	<table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>31:1</td> <td>Reserved</td> </tr> <tr> <td>0</td> <td> <p>Granularity Descriptor Mapping: If set to '1', then each valid namespace granularity descriptor applies to the LBA format having the same index and the Number of Descriptors field shall be equal to the sum of the values represented by the Number of LBA Formats field and the Number of Unique Capability LBA Formats field in the Identify Namespace data structure (refer to Figure 97 and section 5.TBD). If cleared to '0', then NG Descriptor 0 shall apply to all LBA formats and the Number of Descriptors field shall be cleared to 0h.</p> </td> </tr> </tbody> </table>	Bits	Description	31:1	Reserved	0	<p>Granularity Descriptor Mapping: If set to '1', then each valid namespace granularity descriptor applies to the LBA format having the same index and the Number of Descriptors field shall be equal to the sum of the values represented by the Number of LBA Formats field and the Number of Unique Capability LBA Formats field in the Identify Namespace data structure (refer to Figure 97 and section 5.TBD). If cleared to '0', then NG Descriptor 0 shall apply to all LBA formats and the Number of Descriptors field shall be cleared to 0h.</p>
	Bits	Description					
31:1	Reserved						
0	<p>Granularity Descriptor Mapping: If set to '1', then each valid namespace granularity descriptor applies to the LBA format having the same index and the Number of Descriptors field shall be equal to the sum of the values represented by the Number of LBA Formats field and the Number of Unique Capability LBA Formats field in the Identify Namespace data structure (refer to Figure 97 and section 5.TBD). If cleared to '0', then NG Descriptor 0 shall apply to all LBA formats and the Number of Descriptors field shall be cleared to 0h.</p>						
...							

Add section 5.TBD as shown below:

5.TBD LBA Format List Structure

To create or format a namespace, a host specifies an LBA Format (i.e., an index) referencing the LBA Format List in the NVM Command Set Identify Namespace data structure. The LBA Format List has a structure as is illustrated in Figure X. The NLBAF field, a 0-based number (i.e. at least one is required to be defined), identifies the number of LBA Formats that have the same capabilities used to format and create a namespace (i.e., the green LBA Formats). The Identify command provides the ability to access these same capabilities in a namespace data structure by specifying an NSID of FFFFFFFh.

The NULBAF field, a 1-based number (i.e., none may be defined), identifies the number of LBA Formats that have unique capabilities (i.e., that may not be the same capabilities as other LBA Formats) used to format and create a namespace (i.e., the orange LBA Formats). A host should use the Identify command with a CNS value of 09h to access the capabilities of a specific LBA Format for an NVM Command Set Identify Namespace data structure. A host should use the Identify command with a CNS value of 0Ah to access the capabilities of a specific LBA Format for a I/O Command Set specific Identify Namespace data structure for the NVM Command Set.

The maximum number of LBA formats allowed to be supported is:

- 16 if the host has cleared the LBA Format Extension Enable (LBAFEE) field to 0h in the Host Behavior Support feature (refer to the Host Behavior Support section in the NVMe Base Specification); or
- 64 if the host has set the LBAFEE field to 1h in the Host Behavior Support feature (refer to the Host Behavior Support section in the NVMe Base Specification).

The number of LBA formats supported is the sum of the values represented by the NLBAF field and the NULBAF field. A LBA Format index is valid if the index is less than the sum of the NLBAF field and the NULBAF field.

Figure X: LBA Format List Structure



Modify portions of Key Value Command Set Specification 1.0 as shown below:

Modify section 4.1.5 as shown below:

4.1.5 Identify Command

This specification implements the Identify Command and associated Identify data structures defined in the NVMe Base specification. Additionally, the Key Value Command Set specifies the data structures defined in this section.

Each I/O Command Set is assigned a specific Command Set Identifier (CSI) value by the NVMe Base Specification. The Key Value Command Set is assigned a CSI value of 01h.

Figure 34: Identify – CNS Values

CNS Value	O/M ¹	Definition	NSID ²	CNTID ³	CSI ⁴	Reference Section
Active Namespace Management						
05h	M ⁵	Identify I/O Command Set specific Namespace data structure for the specified NSID for the I/O Command Set specified in the CSI field.	Y	N	Y	4.1.5.1
06h	M	Identify I/O Command Set Specific Controller data structure for the controller processing the command.	Y	N	Y	4.1.5.2
0Ah	O	I/O Command Set specific Identify Namespace data structure for the specified KV Format index for the I/O Command Set specified in the CSI field. ⁶	N	N	Y	4.1.5.TBD

NOTES:

- O/M definition: O = Optional, M = Mandatory.
- The NSID field is used: Y = Yes, N = No.

CNS Value	O/M ¹	Definition	NSID ²	CNTID ³	CSI ⁴	Reference Section
3.		The CDW10.CNTID field is used: Y = Yes, N = No.				
4.		The CDW11.CSI field is used: Y = Yes, N = No.				
5.		Mandatory for controllers that support the Namespace Management capability (refer to the NVMe Base Specification).				
6.		Selection of a UUID may be supported. Refer to the Universally Unique Identifiers (UUIDs) for Vendor Specific Information section in the NVMe Base Specification.				

Modify section 4.1.5.1 as shown below:

4.1.5.1 I/O Command Set specific Identify Namespace data structure (CNS 05h, CSI 01h)

The I/O Command Set specific Identify Namespace data structure (i.e., CNS 05h) for the Key Value Command Set is defined in Figure 35.

Figure 35: Identify – Identify Namespace Data Structure, Key Value Type Specific

Bytes	O/M ¹	Description	Capability Field
07:00	M	Namespace Size (NSZE): This field indicates the total size of the namespace in bytes. This is the space to store KV keys and KV values. This field is undefined prior to the namespace being formatted.	No
15:08		Reserved	
23:16	M	Namespace Utilization (NUSE): This field indicates the current number of bytes allocated in the namespace. This is the space to store KV keys and KV values. This field is smaller than or equal to the Namespace Size. A key value pair is allocated when it is written with a Store command. A key value pair is deallocated using the Delete command. If the controller supports Asymmetric Namespace Access Reporting (refer to the CMIC field), and the relationship between the controller and the namespace is in the ANA Inaccessible state (refer to the NVMe Base Specification) or the ANA Persistent Loss state (refer to the NVMe Base Specification), then this field shall be cleared to 0h.	No
24	M	Namespace Features (NSFEAT): This field defines features of the namespace. Bits 7:4 are reserved. Bit 3 if set to '1' indicates that the non-zero NGUID and non-zero EUI64 fields for this namespace are never reused by the controller. If cleared to '0', then the NGUID and EUI64 values 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 the NVMe Base Specification. Bits 2:0 are reserved.	No
25	M	Number of KV Formats (NKVF): This field defines the number of KV format descriptors supported by the namespace. KV formats shall be allocated in order (starting with 0) and packed sequentially. This is a 0's based value. The maximum number of KV formats that may be indicated as supported is 16. The supported KV formats are indicated in bytes 72 to 327 in this data structure. The KV Format fields with an index beyond the value set in this field are invalid and not supported. KV Formats that are valid, but not currently available may be indicated by setting the KV Key Max and KV Value Max both to 0000h for that KV Format.	Yes
26	O	Namespace Multi-path I/O and Namespace Sharing Capabilities (NMIC): Refer to the NMIC field in the Identify Namespace data structure in the NVMe Base Specification.	No

Figure 35: Identify – Identify Namespace Data Structure, Key Value Type Specific

Bytes	O/M ¹	Description	Capability Field
27	O	Reservation Capabilities (RESCAP): Refer to NVMe Base Specification.	No
28	O	Format Progress Indicator (FPI): Refer to NVMe Base Specification.	No
31:29		Reserved	
35:32	O	Namespace Optimal Value Granularity (NOVG): This field indicates the optimal value granularity for this namespace. This field is specified in bytes. The host should construct Store commands that store multiples of NOVG bytes to achieve optimal performance. A value of 0h indicates that no optimal value granularity is reported.	No
39:36	O	ANA Group Identifier (ANAGRPID): Refer to NVMe Base Specification.	No
42:40		Reserved	
43	O	Namespace Attributes (NSATTR): Refer to NVMe Base Specification.	No
45:44	O	NVM Set Identifier (NVMSETID): Refer to NVMe Base Specification.	No
47:46	O	Endurance Group Identifier (ENDGID): Refer to NVMe Base Specification.	No
63:48	O	Namespace Globally Unique Identifier (NGUID): Refer to NVMe Base Specification.	No
71:64	O	IEEE Extended Unique Identifier (EUI64): Refer to NVMe Base Specification.	No
87:72	M	KV Format 0 Support (KVF0): This field indicates the KV format 0 that is supported by the controller. The KV format field is defined in Figure 37.	Yes
103:88	O	KV Format 1 Support (KVF1): This field indicates the KV format 1 that is supported by the controller. The KV format field is defined in Figure 37.	Yes
119:104	○	KV Format 2 Support (KVF2): This field indicates the KV format 2 that is supported by the controller. The KV format field is defined in Figure 37.	Yes
135:120	○	KV Format 3 Support (KVF3): This field indicates the KV format 3 that is supported by the controller. The KV format field is defined in Figure 37.	Yes
151:136	○	KV Format 4 Support (KVF4): This field indicates the KV format 4 that is supported by the controller. The KV format field is defined in Figure 37.	Yes
167:152	○	KV Format 5 Support (KVF5): This field indicates the KV format 5 that is supported by the controller. The KV format field is defined in Figure 37.	Yes
183:168	○	KV Format 6 Support (KVF6): This field indicates the KV format 6 that is supported by the controller. The KV format field is defined in Figure 37.	Yes
199:184	○	KV Format 7 Support (KVF7): This field indicates the KV format 7 that is supported by the controller. The KV format field is defined in Figure 37.	Yes
215:200	○	KV Format 8 Support (KVF8): This field indicates the KV format 8 that is supported by the controller. The KV format field is defined in Figure 37.	Yes
231:216	○	KV Format 9 Support (KVF9): This field indicates the KV format 9 that is supported by the controller. The KV format field is defined in Figure 37.	Yes
247:232	○	KV Format 10 Support (KVF10): This field indicates the KV format 10 that is supported by the controller. The KV format field is defined in Figure 37.	Yes
263:248	○	KV Format 11 Support (KVF11): This field indicates the KV format 11 that is supported by the controller. The KV format field is defined in Figure 37.	Yes
279:264	○	KV Format 12 Support (KVF12): This field indicates the KV format 12 that is supported by the controller. The KV format field is defined in Figure 37.	Yes
295:280	○	KV Format 13 Support (KVF13): This field indicates the KV format 13 that is supported by the controller. The KV format field is defined in Figure 37.	Yes

Figure 35: Identify – Identify Namespace Data Structure, Key Value Type Specific

Bytes	O/M ¹	Description	Capability Field
311:296	⊖	KV Format 14 Support (KVF14): This field indicates the KV format 14 that is supported by the controller. The KV format field is defined in Figure 37.	Yes
...	
327:312	O	KV Format 15 Support (KVF15): This field indicates the KV format 15 that is supported by the controller. The KV format field is defined in Figure 37.	Yes
3839:328		Reserved	
4095:3840	O	Vendor Specific	No
NOTES:			
1. O/M definition: O = Optional, M = Mandatory.			

Add section 4.1.5.TBD1 as shown below:

4.1.5.TBD1 Identify I/O Command Set specific Namespace data structure (CNS 0Ah, CSI 01h)

An I/O Command Set specific Identify Namespace data structure for the Key Value Command Set (refer to Figure 9) is returned to the host for the for the specified KV Format Index specified by the CNS Specific Identifier field as defined in Figure CNS_TBDH. The returned I/O Command Set specific Identify Namespace data structure specifies fields that define capabilities used by a host to format or create a namespace. If the KV Format index specified is valid, then the controller shall return an I/O Command Set specific Identify Namespace data structure that has:

- the capability fields as specified by the I/O Command Set specific Identify Namespace data structure for the Key Value Command Set that requires a value that is the same value for all namespaces formatted using the KV formats associated with the Number of KV Formats field; and
- fields that are not capability fields set to the value 0h.

Figure CNS_TBDH: Command Dword 11 – CNS Specific Identifier

Bits	Description
15:0	KV Format Index Identifier: This field specifies the index into the KV Format list (refer to Figure 36) identifying the KV Format capabilities that are to be returned.

Modify portions of Zoned Namespace Command Set Specification 1.1 as shown below:

Modify section 4.1.5 as shown below:

4.1.5 Identify Command

...

Figure 45: CNS Values

CNS Value	O/M ¹	Definition	NSID ²	CNTID ³	CSI ⁴	Reference Section
Active Namespace Management						
...						
06h	M	Identify I/O Command Set Specific Controller data structure for the controller processing the command. ⁵	Y	N	Y	4.1.5.2
09h	O	Identify Namespace data structure for the specified LBA Format index containing the namespace capabilities for the NVM Command Set. ⁵	N	N	Y	NVM Command Set Specification
0Ah	O	I/O Command Set specific Identify Namespace data structure for the specified LBA Format index for the I/O Command Set specified in the CSI field. ⁵	N	N	Y	4.1.5.TBD
16h	O	A Namespace Granularity List is returned to the host for up to sixteen Namespace Granularity Entries.	N	N	N ⁶	NVM Command Set Specification
NOTES: 1. O/M definition: O = Optional, M = Mandatory. 2. The NSID field is used: Y = Yes, N = No. 3. The CDW10.CNTID field is used: Y = Yes, N = No. 4. The CDW11.CSI field is used: Y = Yes, N = No. 5. Selection of a UUID may be supported. Refer to the UUIDs for Vendor Specific Information section of the NVM Command Set Specification. 6. This Identify data structure applies to namespace that are associated with command sets that specify logical blocks (i.e., Command Set Identifier 0h or 02h).						

Modify section 4.1.5.1 as shown below:

4.1.5.1 Identify I/O Command Set Specific Namespace Data Structure for the Zoned Namespace Command Set (CNS 05h, CSI 02h)

Figure 48 defines the Identify I/O Command Set specific Namespace data structure for the Zoned Namespace Command Set.

Figure 48 specifies fields in the Zoned Namespace Command Set I/O Command Set Specific Identify Namespace data structure that define capabilities used by a host to format or create a namespace. If the NSID field is set to FFFFFFFFh, then the controller shall return a Zoned Namespace Command Set I/O Command Set Specific Identify Namespace data structure that has:

- the capability fields as specified by Figure 48 that requires a value that is the same value for all namespaces formatted using the LBA formats associated with the Number of LBA Formats field; and
- fields that are not capability fields set to the value 0h.

If the controller supports the Namespace Management capability (refer to the Namespace Management section in the NVMe Base Specification) and the NSID field is set to FFFFFFFFh, then the controller shall return a Zoned Namespace Command Set I/O Command Set Specific Identify Namespace data structure. If the controller does not support the Namespace Management capability and the NSID field is set to FFFFFFFFh, then the controller may abort the command with a status code of Invalid Namespace or Format.

Figure 48: Identify I/O Command Set Specific Namespace Data Structure for the Zoned Namespace Command Set

Bytes	O/M ¹	Description	Capability Field								
01:00	O	<p>Zone Operation Characteristics (ZOC): This field indicates the zone operation characteristics of the zoned namespace.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>15:2</td> <td>Reserved</td> </tr> <tr> <td>1</td> <td>Zone Active Excursions: If set to '1', then a controller may transition a zone in the ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, or the ZSC:Closed state to the ZSF:Full state due to a vendor specific excursion event. If cleared to '0', then a controller shall not transition a zone due to a vendor specific excursion event. Refer to section 5.6.</td> </tr> <tr> <td>0</td> <td>Variable Zone Capacity: if set to '1', then the capacity for a zone may change without a change to the format of the zoned namespace. If cleared to '0', then the capacity for a zone does not change without a change to the format of the zoned namespace. Refer to Figure 37.</td> </tr> </tbody> </table>	Bits	Description	15:2	Reserved	1	Zone Active Excursions: If set to '1', then a controller may transition a zone in the ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, or the ZSC:Closed state to the ZSF:Full state due to a vendor specific excursion event. If cleared to '0', then a controller shall not transition a zone due to a vendor specific excursion event. Refer to section 5.6.	0	Variable Zone Capacity: if set to '1', then the capacity for a zone may change without a change to the format of the zoned namespace. If cleared to '0', then the capacity for a zone does not change without a change to the format of the zoned namespace. Refer to Figure 37.	No
Bits	Description										
15:2	Reserved										
1	Zone Active Excursions: If set to '1', then a controller may transition a zone in the ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, or the ZSC:Closed state to the ZSF:Full state due to a vendor specific excursion event. If cleared to '0', then a controller shall not transition a zone due to a vendor specific excursion event. Refer to section 5.6.										
0	Variable Zone Capacity: if set to '1', then the capacity for a zone may change without a change to the format of the zoned namespace. If cleared to '0', then the capacity for a zone does not change without a change to the format of the zoned namespace. Refer to Figure 37.										
03:02	O	<p>Optional Zoned Command Support (OZCS): This field defines optional features of the zoned namespace.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>15:1</td> <td>Reserved</td> </tr> <tr> <td>0</td> <td>Read Across Zone Boundaries: If set to '1', then any User Data Read Access Command is allowed to perform read operations that specify an LBA range containing logical blocks in more than one zone. If cleared to '0', then any command that performs a read operation that specifies an LBA range containing logical blocks in more than one zone is aborted as defined in section 2.1.1.2.1.2.</td> </tr> </tbody> </table>	Bits	Description	15:1	Reserved	0	Read Across Zone Boundaries: If set to '1', then any User Data Read Access Command is allowed to perform read operations that specify an LBA range containing logical blocks in more than one zone. If cleared to '0', then any command that performs a read operation that specifies an LBA range containing logical blocks in more than one zone is aborted as defined in section 2.1.1.2.1.2.	No		
Bits	Description										
15:1	Reserved										
0	Read Across Zone Boundaries: If set to '1', then any User Data Read Access Command is allowed to perform read operations that specify an LBA range containing logical blocks in more than one zone. If cleared to '0', then any command that performs a read operation that specifies an LBA range containing logical blocks in more than one zone is aborted as defined in section 2.1.1.2.1.2.										
07:04	M	Maximum Active Resources (MAR): This field defines the maximum number of zones that are allowed to be in the ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, or the ZSC:Closed state. A value of FFFFFFFh indicates that there is no limit. This is a 0's based value	No								
11:08	M	Maximum Open Resources (MOR): This field defines the maximum number of zones that are allowed to be in the ZSIO:Implicitly Opened state or the ZSEO:Explicitly Opened state. The Maximum Open Resource field shall be less than or equal to the Maximum Active Resources field. A value of FFFFFFFh indicates that there is no limit. This is a 0's based value	No								
15:12	O	Reset Recommended Limit (RRL): If the zone attribute Reset Zone Recommended Time Limit field is set to 00b, then this field indicates the number of seconds before the NVM subsystem may perform a vendor specific action on a zone after the Reset Zone Recommended zone attribute is set to '1' for that zone. If this field is cleared to 0h, then no Reset Recommended Limit is reported. Refer to section 5.4.	No								
19:16	O	Finish Recommended Limit (FRL): If the zone attribute Finish Zone Recommended Time Limit field is set to 00b, then this field indicates the number of seconds before the NVM subsystem may perform the vendor specific action on a zone after the Finish Zone Recommended zone attribute is set to '1' for that zone. If this field is cleared to 0h, then no Finish Recommended Limit is reported. Refer to section 5.5.	No								
23:20	O	Reset Recommended Limit 1 (RRL 1): If the zone attribute Reset Zone Recommended Time Limit field is set to 01b, then this field indicates the number of seconds before the NVM subsystem may perform a vendor specific action on a zone after the Reset Zone Recommended zone attribute is set to '1' for that zone. If this field is cleared to 0h, then no Reset Recommended Limit is reported. Refer to section 5.4.	No								

Figure 48: Identify I/O Command Set Specific Namespace Data Structure for the Zoned Namespace Command Set

Bytes	O/M ¹	Description	Capability Field
27:24	O	Reset Recommended Limit 2 (RRL 2): If the zone attribute Reset Zone Recommended Time Limit field is set to 10b, then this field indicates the number of seconds before the NVM subsystem may perform a vendor specific action on a zone after the Reset Zone Recommended zone attribute is set to '1' for that zone. If this field is cleared to 0h, then no Reset Recommended Limit is reported. Refer to section 5.4.	No
31:28	O	Reset Recommended Limit 3 (RRL 3): If the zone attribute Reset Zone Recommended Time Limit field is set to 11b, then this field indicates the number of seconds before the NVM subsystem may perform a vendor specific action on a zone after the Reset Zone Recommended zone attribute is set to '1' for that zone. If this field is cleared to 0h, then no Reset Recommended Limit is reported. Refer to section 5.4.	No
35:32	O	Finish Recommended Limit 1 (FRL1): If the zone attribute Finish Zone Recommended Time Limit field is set to 01b, then this field indicates the number of seconds before the NVM subsystem may perform the vendor specific action on a zone after the Finish Zone Recommended zone attribute is set to '1' for that zone. If this field is cleared to 0h, then no Finish Recommended Limit is reported. Refer to section 5.5.	No
39:36	O	Finish Recommended Limit 2 (FRL2): If the zone attribute Finish Zone Recommended Time Limit field is set to 10b, then this field indicates the number of seconds before the NVM subsystem may perform the vendor specific action on a zone after the Finish Zone Recommended zone attribute is set to '1' for that zone. If this field is cleared to 0h, then no Finish Recommended Limit is reported. Refer to section 5.5.	No
43:40	O	Finish Recommended Limit 3 (FRL3): If the zone attribute Finish Zone Recommended Time Limit field is set to 11b, then this field indicates the number of seconds before the NVM subsystem may perform the vendor specific action on a zone after the Finish Zone Recommended zone attribute is set to '1' for that zone. If this field is cleared to 0h, then no Finish Recommended Limit is reported. Refer to section 5.5.	No
2815:44		Reserved	
2831:2816	M	LBA Format 0 Extension (LBAFE0): This field indicates the LBA format Extension 0 that is supported by the controller. The Zone format field is defined in Figure 49.	Yes
2847:2832	O	LBA Format 1 Extension (LBAFE1): This field indicates the LBA format Extension 1 that is supported by the controller. The LBA Format Extension field is defined in Figure 49.	Yes
...			
3839:3824	O	LBA Format 63 Extension (LBAFE63): This field indicates the LBA format Extension 63 that is supported by the controller. The LBA Format Extension field is defined in Figure 49.	Yes
4095:3840	O	Vendor Specific	No
NOTES:			
1. O/M definition: O = Optional, M = Mandatory.			

Add section 4.1.5.TBD1 as shown below:

4.1.5.TBD Identify I/O Command Set specific Namespace data structure (CNS 0Ah, CSI 02h)

An I/O Command Set specific Identify Namespace data structure for the Zoned Namespace Command Set (refer to Figure 48) is returned to the host for the for the specified LBA Format Index specified by the CNS Specific Identifier field as defined in Figure CNS_TBDH. The returned I/O Command Set specific Identify Namespace data structure specifies fields that define capabilities used by a host to format or create a

namespace. If the User Data format index specified is valid, then the controller shall return an I/O Command Set specific Identify Namespace data structure that has:

- the capability fields as specified by the I/O Command Set specific Identify Namespace data structure for the Zoned Namespace Command Set that requires a value that is the same value for all namespaces formatted using the LBA formats associated with the Number of LBA Formats field; and
- fields that are not capability fields set to the value 0h.

Figure CNS_TBDH: Command Dword 11 – CNS Specific Identifier

Bits	Description
15:0	LBA Format Index Identifier: This field specifies the index into the LBA Format list identifying the LBA Format capabilities that are to be returned. Refer to the NVM Command Set Specification.

Modify portions of NVMe Base Specification 2.0:

Modify section 5.17 as shown below:

5.17 Identify command

5.17.1 Identify command overview

....

Figure 271: Identify – Command Dword 11

Bits	Description								
31:24	Command Set Identifier (CSI): This field is CNS value specific. This field specifies the I/O Command Set to be used by the command for CNS values that require a Command Set Identifier. Refer to Figure 273 for Identify command CNS values that use this field. This field shall be cleared to 0h for Identify operations with CNS values that do not use this field. Values for this field are defined by Figure 274.								
23:16	Reserved								
15:00	CNS Specific Identifier: This field specifies an identifier that is required for a particular CNS value. If not defined for the CNS specified by the CNS field, then this field is reserved. The CNS values that require a CNS specific identifier are indicated in the table below. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>CNS Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>NVM Set List (04h)</td> <td>NVM Set Identifier (NVMSETID) (refer to section 3.2.2)</td> </tr> <tr> <td>Domain List (18h)</td> <td>Domain Identifier (DOMID) (refer to section 3.2.4.3)</td> </tr> <tr> <td>Endurance Group List (19h)</td> <td>Endurance Group Identifier (ENDGID) (refer to section 3.2.3)</td> </tr> </tbody> </table>	CNS Value	Definition	NVM Set List (04h)	NVM Set Identifier (NVMSETID) (refer to section 3.2.2)	Domain List (18h)	Domain Identifier (DOMID) (refer to section 3.2.4.3)	Endurance Group List (19h)	Endurance Group Identifier (ENDGID) (refer to section 3.2.3)
CNS Value	Definition								
NVM Set List (04h)	NVM Set Identifier (NVMSETID) (refer to section 3.2.2)								
Domain List (18h)	Domain Identifier (DOMID) (refer to section 3.2.4.3)								
Endurance Group List (19h)	Endurance Group Identifier (ENDGID) (refer to section 3.2.3)								

...

Figure 273: Identify – CNS Values

CNS Value	O/M ¹	Definition	NSID ²	CNTID ³	CSI ⁴	Reference Section
Active Namespace Management						
00h	M	Identify Namespace data structure for the specified NSID or the common namespace capabilities for the NVM Command Set. ⁷	Y	N	N ⁸	NVM Command Set Specification
...						
07h	M	Active Namespace ID list associated with the specified I/O Command Set.	Y	N	Y	5.17.2.7

Figure 273: Identify – CNS Values

CNS Value	O/M ¹	Definition	NSID ²	CNTID ³	CSI ⁴	Reference Section
09h	O	Identify Namespace data structure for the specified User Data Format index containing the namespace capabilities for the NVM Command Set. ⁷	N	N	Y	NVM Command Set Specification
0Ah	O	I/O Command Set specific Identify Namespace data structure for the specified User Data Format index containing the namespace capabilities for the I/O Command Set specified in the CSI field. ⁷	N	N	Y	I/O Command Set Specification
0E08h to 0Fh		Reserved				
Controller and Namespace Management						
...						
18h to 1Fh		Reserved				
Future Definition						
20h to FFh		Reserved				
NOTES:						
1. O/M definition: O = Optional, M = Mandatory.						
2. The NSID field is used: Y = Yes, N = No.						
3. The CDW10.CNTID field is used: Y = Yes, N = No.						
4. The CDW11.CSI field is used: Y = Yes, N = No.						
5. Mandatory for controllers that support the Namespace Management capability (refer to section 8.11).						
6. Mandatory for controllers that support Virtualization Enhancements (refer to section 8.26).						
7. Selection of a UUID may be supported (refer to section 8.25).						
8. This Identify data structure applies to namespaces that are associated with command sets that specify logical blocks (i.e., Command Set Identifier 0h or 2h).						
9. Mandatory for controllers that support Variable Capacity Management (refer to section 8.3.3).						
10. Mandatory for controllers that support Capacity Management (refer to section 8.3) in an NVM subsystem that supports multiple domains (refer to section 3.2.4).						
11. Only applicable for the NVM Command Set and I/O Command Sets based on the NVM Command Set. Prohibited for all other I/O Command Sets.						

Modify section 5.17.2.4 as shown below:

5.17.2.4 NVM Set List (CNS 04h)

Figure 278 defines an NVM Set List. The data structure is an ordered list of NVM Set Attribute Entry data structures, sorted by NVM Set Identifier, starting with the first NVM Set Identifier supported by the NVM subsystem that is equal to or greater than the NVM Set Identifier specified by the CNS Specific Identifier field as defined in Figure CNS_04H indicated in CDW11.NVMSETID and are accessible by the controller processing the command. The NVM Set List describes the attributes for each NVM Set in the list based on the NVM Set Attributes Entry in Figure 278.

Figure CNS_04H: Command Dword 11 – CNS Specific Identifier

Bits	Description
15:0	NVM Set Identifier (NVMSETID): This field specifies the NVM Set Identifier of the first NVM Set of the ordered list of NVM Set Attribute Entry data structures to be returned.

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

...

Modify section 5.17.2.5 as shown below:

5.17.2.5 I/O Command Set specific Identify Namespace data structure (CNS 05h)

An I/O Command Set specific Identify Namespace data structure (refer to the applicable I/O Command Set specification) is returned to the host for the namespace specified in the Namespace Identifier (NSID) field if the NSID is active. If the specified namespace is an inactive NSID, then the controller returns a zero filled data structure.

The specific Identify Namespace data structure that is returned by this command is specified by the Command Set Identifier (CSI) field (refer to Figure 274). If the I/O Command Set associated with the namespace identified by the NSID field does not support the Identify Namespace data structure specified by the CSI field, the controller shall abort the command with a status code of Invalid Field in Command.

~~If the controller supports the Namespace Management capability (refer to section 8.11), the I/O Command Set requested in the CSI field has been enabled with the I/O Command Set Profile feature (refer to section 5.27.1.21), and the NSID field is set to FFFFFFFFh, then the controller returns an I/O Command Set specific Identify Namespaces data structure that specifies capabilities that are common across namespaces for the I/O Command Set specified in the CSI field (refer to Figure 274). If the controller does not support the Namespace Management capability and the NSID field is set to FFFFFFFFh, then the controller shall abort the command with a status code of Invalid Namespace or Format.~~

Each I/O Command Set specific Identify Namespace data structure specifies fields that define capabilities used by a host to format or create a namespace. If the NSID field is set to FFFFFFFFh, then the controller shall return an I/O Command Set specific Identify Namespace data structure for the I/O Command Set specified in the CSI field (refer to Figure 274) that has:

- the capability fields as specified by the I/O Command Set specific Identify Namespace data structure that requires a value that is the same value for all namespaces formatted using the User Data Formats associated with the Number of LBA Formats field; and
- fields that are not capability fields set to the value 0h.

If the controller supports the Namespace Management capability (refer to section 8.11) and the NSID field is set to FFFFFFFFh, then the controller shall return an I/O Command Set specific Identify Namespace data structure. If the controller does not support the Namespace Management capability and the NSID field is set to FFFFFFFFh, then the controller may abort the command with a status code of Invalid Namespace or Format.

Modify section 5.17.2.8 as shown below:

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

If the Namespace Identifier (NSID) field specifies an active NSID, then the I/O Command Set Independent Identify Namespace data structure (refer to Figure 280) is returned to the host for that specified namespace. If that specified namespace is an inactive NSID, then the controller returns a zero filled data structure.

~~If the controller supports the Namespace Management capability (refer to section 8.11) and the NSID field is set to FFFFFFFFh, then the controller returns an I/O Command Set Independent Identify Namespace data structure that specifies capabilities that are common for the controller. If the controller does not support the Namespace Management capability and the NSID field is set to FFFFFFFFh, then the controller shall abort the command with a status code of Invalid Namespace or Format.~~

Figure 280 specifies fields in the I/O Command Set Independent Identify Namespace data structure that define capabilities used by a host to format or create a namespace. If the NSID field is set to FFFFFFFFh, then the controller shall return an I/O Command Set Independent Identify Namespace data structure that has:

- the capability fields as specified by Figure 280 that requires a value that is the same value for all namespaces formatted using the User Data Formats associated with the Number of LBA Formats field; and
- fields that are not capability fields set to the value 0h.

If the controller supports the Namespace Management capability (refer to section 8.11) and the NSID field is set to FFFFFFFFh, then the controller shall return an I/O Command Set Independent Identify Namespace data structure. If the controller does not support the Namespace Management capability and the NSID field is set to FFFFFFFFh, then the controller may abort the command with a status code of Invalid Namespace or Format.

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

Bytes	O/M ¹	Description	Capability Field
00	M	<p>Common Namespace Features (NSFEAT): This field defines features of the namespace.</p> <p>Bits 7:5 are reserved.</p> <p>Bit 4 (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.</p> <p>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.</p> <p>Bit 2:1 are reserved.</p> <p>Bit 0 (Rotational Media) if set to '1' indicates that the namespace is associated with rotational media (refer to section 8.20). If cleared to '0', then the namespace is not associated with rotational media.</p>	No
01	O	<p>Namespace Multi-path I/O and Namespace Sharing Capabilities (NMIC): This field specifies multi-path I/O and namespace sharing capabilities of the namespace.</p> <p>Bits 7:1 are reserved.</p> <p>Bit 0: If set to '1', then the namespace may be attached to two or more controllers in the NVM subsystem concurrently (i.e., may be a shared namespace). If cleared to '0', then the namespace is a private namespace and is able to be attached to only one controller at a time.</p>	No

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

Bytes	O/M ¹	Description	Capability Field
02	O	<p>Reservation Capabilities (RESCAP): This field indicates the reservation capabilities of the namespace. A value of 0h in this field indicates that reservations are not supported by this namespace. Refer to section 8.19 for more details.</p> <p>Bit 7 if set to '1' indicates that Ignore Existing Key is used as defined in revision 1.3 or later of this specification. Bit 7 if cleared to '0' indicates that Ignore Existing Key is used as defined in revision 1.2.1 or earlier of this specification. This bit shall be set to '1' if the controller supports revision 1.3 or later as indicated in the Version register.</p> <p>Bit 6 if set to '1' indicates that the namespace supports the Exclusive Access – All Registrants reservation type. If this bit is cleared to '0', then the namespace does not support the Exclusive Access – All Registrants reservation type.</p> <p>Bit 5 if set to '1' indicates that the namespace supports the Write Exclusive – All Registrants reservation type. If this bit is cleared to '0', then the namespace does not support the Write Exclusive – All Registrants reservation type.</p> <p>Bit 4 if set to '1' indicates that the namespace supports the Exclusive Access – Registrants Only reservation type. If this bit is cleared to '0', then the namespace does not support the Exclusive Access – Registrants Only reservation type.</p> <p>Bit 3 if set to '1' indicates that the namespace supports the Write Exclusive – Registrants Only reservation type. If this bit is cleared to '0', then the namespace does not support the Write Exclusive – Registrants Only reservation type.</p> <p>Bit 2 if set to '1' indicates that the namespace supports the Exclusive Access reservation type. If this bit is cleared to '0', then the namespace does not support the Exclusive Access reservation type.</p> <p>Bit 1 if set to '1' indicates that the namespace supports the Write Exclusive reservation type. If this bit is cleared to '0', then the namespace does not support the Write Exclusive reservation type.</p> <p>Bit 0 if set to '1' indicates that the namespace supports the Persist Through Power Loss capability. If this bit is cleared to '0', then the namespace does not support the Persist Through Power Loss Capability.</p>	No
03	O	<p>Format Progress Indicator (FPI): If a format operation is in progress, this field indicates the percentage of the namespace that remains to be formatted.</p> <p>Bit 7 if set to '1' indicates that the namespace supports the Format Progress Indicator defined by bits 6:0 in this field. If this bit is cleared to '0', then the namespace does not support the Format Progress Indicator and bits 6:0 in this field shall be cleared to 0h.</p> <p>Bits 6:0 indicate the percentage of the Format NVM command that remains to be completed (e.g., a value of 25 indicates that 75% of the Format NVM command has been completed and 25% remains to be completed). If bit 7 is set to '1', then a value of 0h indicates that the namespace is formatted with the format specified by the FLBAS and DPS fields in this data structure and there is no Format NVM command in progress.</p>	No

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

Bytes	O/M ¹	Description	Capability Field						
07:04	O	<p>ANA Group Identifier (ANAGRPID): For NSID other than FFFFFFFFh, this field indicates the ANA Group Identifier of the ANA group (refer to section 8.1.2) of which the namespace is a member. Each namespace that is attached to a controller that supports Asymmetric Namespace Access Reporting (refer to the CMIC field) shall report a valid ANAGRPID. If the controller does not support Asymmetric Namespace Access Reporting, then this field shall be cleared to 0h.</p> <p>If the value in this field changes and Asymmetric Namespace Access Change Notices are supported and enabled, then the controller shall issue an Asymmetric Namespace Access Change Notice.</p>	No						
08	O	<p>Namespace Attributes (NSATTR): This field specifies attributes of the namespace.</p> <p>Bits 7:1 are reserved.</p> <p>Bit 0: If set to '1', then the namespace is currently write protected due to any condition (e.g., namespace write protection set for the namespace, media errors) and all write access to the namespace shall fail. If cleared to '0', then the namespace is not currently write protected.</p>	No						
09		Reserved							
11:10	O	<p>NVM Set Identifier (NVMSETID): For NSID other than FFFFFFFFh, this field indicates the NVM Set with which this namespace is associated. If NVM Sets are not supported by the controller, then this field shall be cleared to 0h.</p>	No						
13:12	O	<p>Endurance Group Identifier (ENDGID): For NSID other than FFFFFFFFh, this field indicates the Endurance Group with which this namespace is associated. If Endurance Groups are not supported by the controller, then this field shall be cleared to 0h.</p>	No						
14	M	<p>Namespace Status (NSTAT): This field indicates the status of the namespace with the specified NSID.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:1</td> <td>Reserved</td> </tr> <tr> <td>0</td> <td>Namespace Ready (NRDY): A value of '1' indicates that the namespace is ready (refer to section 3.5.3). A value of '0' indicates that the namespace is not ready.</td> </tr> </tbody> </table>	Bits	Description	7:1	Reserved	0	Namespace Ready (NRDY): A value of '1' indicates that the namespace is ready (refer to section 3.5.3). A value of '0' indicates that the namespace is not ready.	No
Bits	Description								
7:1	Reserved								
0	Namespace Ready (NRDY): A value of '1' indicates that the namespace is ready (refer to section 3.5.3). A value of '0' indicates that the namespace is not ready.								
4095:15		Reserved							
NOTES:									
1. O/M definition: O = Optional, M = Mandatory.									

Modify section 5.17.2.17 as shown below:

5.17.2.17 Domain List (CNS 18h)

Figure 286 defines a Domain List. The data structure is an ordered list by Domain Identifier, starting with the first Domain Identifier that is greater than or equal to ~~or greater than~~ the Domain Identifier specified by the CNS Specific Identifier field as defined in **Figure CNS_18H** ~~CDW41.NVMSETID~~ and is accessible by the controller processing the command. The Domain List describes the attributes for each Domain in the list based on the Domain Attributes Entry in 6.

Figure CNS_18H: Command Dword 11 – CNS Specific Identifier

Bits	Description
15:0	Domain Identifier (DOMID): This field specifies the Domain Identifier of the first Domain of the ordered list of Domain Attribute Entry data structures to be returned.

...

Modify section 5.17.2.18 as shown below:

5.17.2.18 Endurance Group List (19h)

An Endurance Group List (refer to 8) of up to 2,047 Endurance Group Identifiers in increasing order is returned containing an Endurance Group Identifier greater than or equal to the ~~value specified in the~~ Endurance Group Identifier ~~(CDW11.ENDGID)~~ specified by the CNS Specific Identifier field as defined in **Figure CNS_19H**. The list contains Endurance Group Identifiers of Endurance Groups that are accessible by the controller processing the command. If the value specified in the Endurance Group Identifier is greater than ENDGIDMAX, then the controller shall complete the command with a status of Successful Completion and return an Endurance Group List containing no Endurance Group Identifiers.

Figure CNS_19H: Command Dword 11 – CNS Specific Identifier

Bits	Description
15:0	Endurance Group Identifier (ENGGID): This field specifies the Endurance Group Identifier of the first Endurance Group of the ordered list of Endurance Group Identifiers to be returned.

...