



LEGAL NOTICE:

© Copyright 2007 - 2019 NVM Express, Inc. ALL RIGHTS RESERVED.

This NVM Express revision 1.3 technical proposal is proprietary to the NVM Express, Inc. (also referred to as "Company") and/or its successors and assigns.

NOTICE TO USERS WHO ARE NVM EXPRESS, INC. MEMBERS: Members of NVM Express, Inc. have the right to use and implement this NVM Express revision 1.3 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 - 2019 NVM Express, Inc. ALL RIGHTS RESERVED." When making any such citations or references to this document you are not permitted to revise, alter, modify, make any derivatives of, or otherwise amend the referenced portion of this document in any way without the prior express written permission of NVM Express, Inc. Nothing contained in this document shall be deemed as granting you any kind of license to implement or use this document or the specification described therein, or any of its contents, either expressly or impliedly, or to any intellectual property owned or controlled by NVM Express, Inc., including, without limitation, any trademarks of NVM Express, Inc.

LEGAL DISCLAIMER:

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

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

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

NVM Express Technical Proposal for New Feature

Technical Proposal ID	TP 4004b
Change Date	6/12/2019
Builds on Specification	NVM Express 1.3 or later; does not apply to versions earlier than 1.3.

Technical Proposal Author(s)

Name	Company
Fred Knight	NetApp
David Black	Dell EMC
Curtis Ballard	HPE
Christoph Hellwig	WDC

This technical proposal defines capabilities that enable reporting of Asymmetric Namespace Access characteristics to the host. This TP is also associated with and builds on changes described in TP 4028 (Error codes and Pathing errors). This TP is also related to TP 4009 (Domains and Partitioning), ~~and~~ TP 4003 (IO Determinism), and TP 4018a (NVM Sets and Read Recovery Level).

TP 4004a adds clarifications to the ANA protocol. As implementations progress, these clarifications were found to be necessary. Existing text indicated that ANA groups that contained namespaces were reported, new text was added to clarify that empty ANA groups are not reported.

TP4004b corrects a typo that created a conflict in two descriptions for the NCAP field (section 6.1.6 vs. Figure 114). The NCAP in section 6.1.6 should have been NVMCAP.

Revision History

Revision Date	Change Description
4/27/2017	Move into the standard TP template.
5/1/2017	Sub-group con-call Edits. Added Failed state; NS management interactions; and state machine. Try to cleanup capacity issues re: inaccessible/failed states.
5/3/2017	Sub-group con-call Edits. Change Asymmetric Access Failure to Asymmetric Access Persistent Loss. Integrate with NS management. Change "Asymmetric Access" statements to "Asymmetric Namespace Access"
5/14/17	Add Command effects table; clarify private namespace interactions; clarify INACCESSIBLE and PERSISTENT LOSS error implications relate controller to individual NS groups.
5/17/17	Sub-group con-call Edits.
5/18/17	Group con-call Edits, change Group ID to 2 bytes from 4. Clarify command processing in INACCESSIBLE and PERSISTENT LOSS (again). Added reviewer's notes section, and other e-mail feedback.

5/30/17	Add Subsystem/Node/partition concepts.
5/31/17	Break into 2 parts – ANA (this part) and Partition Model. ANA now makes NO statements about how or why ANA states happen, it just explains what to do when it does happen.
6/5/17 & 6/14/17	Changes from con-call
8/16/17	Limit to 1.3 or later; updates based on NVM Sets
9/19/17	Updates based on discussion to restore ANA specific grouping and not overload NVM Sets; add path and transport type error status
9/21/17	Increase ANA Group ID to 4 bytes so it matches the NSID size.
9/28/17	Change “transport type error” to “pathing error” and other con-call changes.
11/28/17	Resolve Figure 9 (Effects) questions – FFFFFFFFh case already covered in existing text (see ECN-002 or 1.3a)
11/30/17	Move AEN exception text to Figure 49, so it’s all in one place.
12/06/17	Update Log specific field for ANA log page
12/07/17	Add connections in the description text to other related TPs. Update log specific field definition for ANA log page.
12/13/17	Split ANA log page into 2 pages; for Attached NS and Allocated NS.
12/17/17	Remove allocated log page, defer allocated. Added informative section on host action descriptions for ANA states. Rework definition of Change Count field in ANA Group Descriptor to allow 0h to indicate that a change count is not reported (due to configs with lots of small groups).
12/26/17	Address comments from Mike Allison (e.g., remove duplicate material that is in TP4028)
12/29/17	Address comments from John Geldman (e.g., add several cross-references)
1/2/18	Updates during sub-group call (change asymmetric operation to asymmetric controller behavior, change i.e. in the GROUP definition, add explicit definition of valid ANAGRPID values).
1/4/18	Editorial changes from group call, and remove zeroing of NSZE and NCAP while inaccessible, fix MNN field offset (to use “reserved” space rather than previously allocated space).
1/9/18	Change MNN to MNAN and modify description
1/16/18	Remove namespace attribute change events that resulted from an ANA change.
1/18/18	Editorial changes from Curtis Ballard and Kevin Marks
1/30/18	Group changes based on Kevin’s comments (w/change tracking turned on).
2/26/18	Fix offsets that overlapped with other TPs - found in 30-day review
7/12/18	Add clarifications found necessary during implementation.
10/31/18	Integration
11/13/18	Ratified
3/29/19	Fix conflict on NCAP description (change the incorrect NCAP in section 6.1.6 to NVMCAP).
6/12/19	Ratified

Description of Specification Changes

Make changes to the following sections as shown below:

...

1.4.1 Multi-Path I/O and Namespace Sharing

This section provides an overview of multi-path I/O and namespace sharing. Multi-path I/O refers to two or more completely independent ~~PCI Express~~ paths between a single host and a namespace while namespace sharing refers to the ability for two or more hosts to access a common shared namespace using different NVM Express controllers. Both multi-path I/O and namespace sharing require that the NVM subsystem contain two or more controllers. ~~NVM subsystems that support Multi-Path I/O and Namespace Sharing may also support asymmetric controller behavior (refer to section 1.4.2).~~ Concurrent access to a shared namespace by two or more hosts requires some form of coordination between hosts. The procedure used to coordinate these hosts is outside the scope of this specification.

Figure 3 shows an NVM subsystem that contains a single NVM Express controller and a single PCI Express port. Since this is a single Function PCI Express device, the NVM Express controller shall be associated

...

1.4.2 Asymmetric Controller Behavior

Asymmetric controller behavior occurs in NVM subsystems where namespace access characteristics (e.g., performance) may vary based on:

- the internal configuration of the NVM subsystem; or
- which controller is used to access a namespace (e.g., Fabrics).

NVM subsystems that provide asymmetric controller behavior may support Asymmetric Namespace Access Reporting as described in section 8.18.

Make changes to section 4.6.1.2.1 (Generic Command Status Definition) as shown below:

4.6.1.2.1 Generic Command Status Definition

...

Figure 32 Status Code – Generic Command Status Values, NVM Command Set

Value	Description
80h	LBA Out of Range: The command references an LBA that exceeds the size of the namespace.
81h	Capacity Exceeded: Execution of the command has caused the capacity of the namespace to be exceeded. This error occurs when the Namespace Utilization exceeds the Namespace Capacity, as reported in Figure 114.
82h	Namespace Not Ready: The namespace is not ready to be accessed as a result of a condition other than a condition that is reported as an Asymmetric Namespace Access condition. The Do Not Retry bit indicates whether re-issuing the command at a later time may succeed.
83h	Reservation Conflict: The command was aborted due to a conflict with a reservation held on the accessed namespace. Refer to section 8.8.
84h	Format In Progress: A Format NVM command is in progress on the namespace. The Do Not Retry bit shall be cleared to '0' to indicate that the command may succeed if it is resubmitted.
85h – BFh	Reserved

Make changes to section 4.6.1.2.2 (Command Specific Errors Definition) as shown below:

4.6.1.2.2 Command Specific ~~Errors~~-Status Definition

Completion queue entries with a Status Code Type of Command Specific Errors indicate an error that is specific to a particular command opcode. Status codes of 0h to 7Fh are for Admin command errors. Status codes of 80h – BFh are specific to the selected I/O command set.

Figure 33: Status Code – Command Specific Status Values

Value	Description	Commands Affected
...		
15h	Namespace Insufficient Capacity	Namespace Management
16h	Namespace Identifier Unavailable	Namespace Management
17h	Reserved	
18h	Namespace Already Attached	Namespace Attachment
19h	Namespace Is Private	Namespace Attachment
...		
21h	Invalid Number of Controller Resources	Virtualization Management
22h	Invalid Resource Identifier	Virtualization Management
23h	...	
24h	ANA Group Identifier Invalid	Namespace Management
25h	ANA Attach Failed	Namespace Attachment
26h – 6Fh	Reserved	
70h – 7Fh	Directive Specific	NOTE 1
80h – BFh	I/O Command Set Specific	NOTE 2
C0h – FFh	Vendor Specific	
NOTES:		
1. The Directives Specific range defines Directives specific status values. Refer to section 9.		
2. The I/O Command Set Specific range in NVMe over Fabrics defines Fabrics command specific status values.		

Make changes to section 4.6.1.2.4 (Path Related Status Definition) as shown below (these changes are based on additions made in TP4028):

4.6.1.2.4 Path Related Status Definition

...

Figure 36a: Status Code – Path Related Status Values

Value	Description
...	
01h	Asymmetric Namespace Access Persistent Loss: The requested function (e.g., command) is not able to be performed as a result of the relationship between the controller and the namespace being in the ANA Persistent Loss state (refer to section 8.18.3.4). The command should not be re-submitted to the same controller.
02h	Asymmetric Namespace Access Inaccessible: The requested function (e.g., command) is not able to be performed as a result of the relationship between the controller and the namespace being in the ANA Inaccessible state (refer to section 8.18.3.3). The command should not be re-submitted to the same controller.
03h	Asymmetric Namespace Access Transition: The requested function (e.g., command) is not able to be performed as a result of the relationship between the controller and the namespace transitioning between Asymmetric Namespace Access states (refer to section 8.18.3.5). The requested function should be retried after the transition is complete.
...	

Make changes to section 5.2 (Asynchronous Event Request command) as shown below:

5.2 Asynchronous Event Request command

Asynchronous events are used to notify host software of status, error, and health information as these ...

5.2.1 Command Completion

A completion queue entry is posted to the Admin Completion Queue if there is an asynchronous event to ...

Figure 49: Asynchronous Event Information - Notice

Value	Description
00h	<p>Namespace Attribute Changed: The Identify Namespace data structure for one or more namespaces, as well as the Namespace List returned when the Identify command is issued with the CNS field set to 02h, have changed. Host software may use this event as an indication that it should read the Identify Namespace data structures for each namespace to determine what has changed.</p> <p>Alternatively, host software may request the Changed Namespace List (Log Identifier 04h) to determine which namespaces in this controller have changed Identify Namespace information since the last time the log page was read.</p> <p>A controller shall not send this event if:</p> <ul style="list-style-type: none">a) when Namespace Utilization (refer to Figure 114) has changed, as this is a frequent event that does not require action by the host;b) the ANAGRPID field (refer to Figure 114) has changed; orc) capacity information (i.e., the NUSE field and the NVMCAP field) returned in the Identify Namespace Data Structure (refer to Figure 114) changed as a result of an ANA state change. <p>A controller shall only send this event for changes to the Format Progress Indicator field when bits 6:0 of that field transition from a non-zero value to zero, or from a zero value to a non-zero value.</p>
01h	<p>Firmware Activation Starting: The controller is starting a firmware activation process during which command processing is paused. Host software may use CSTS.PP to determine when command processing has resumed. To clear this event, host software reads the Firmware Slot Information log page.</p>
02h	<p>Telemetry Log Changed: The controller has saved the controller internal state in the Telemetry Controller-Initiated log page and set the Telemetry Controller-Initiated Data Available field to 1h in that log page. To clear this event, the host issues a Get Log Page with Retain Asynchronous Event cleared to '0' for the Telemetry Controller-Initiated Log.</p>
03h	<p>Asymmetric Namespace Access Change: The Asymmetric Namespace Access information (refer to section 5.14.1.12) related to an ANA Group that contains namespaces attached to this controller has changed (e.g., an ANA state has changed, an ANAGRPID has changed). The current Asymmetric Namespace Access information for attached namespaces is indicated in the Asymmetric Namespace Access log page (refer to section 5.14.1.12). To clear this event, the host issues a Get Log Page with Retain Asynchronous Event cleared to '0' for the Asymmetric Namespace Access Log.</p> <p>A controller shall not send this event if:</p> <ul style="list-style-type: none">a) the change is due to the creation of a namespace (refer to section 5.20); orb) the change is due to the deletion of a namespace (refer to section 5.20), as the Namespace Attribute Changed event is sent for these changes.
04h 3h – FFh	Reserved

Make changes to section 5.14 (Get Log Page command) as shown below (based on changes in TP4018):

5.14 Get Log Page command

...

5.14.1 Log Specific Information

Figure 90 and Figure 91 define the Log pages that may be retrieved with the Get Log Page command.

Figure 90: Get Log Page – Log Page Identifiers

Log Identifier	O/M	Scope	Description	Reference
00h			Reserved	
...				
08h	O	Controller	Telemetry Controller-Initiated	5.14.1.8
...				
0Ch	O	Controller	Asymmetric Namespace Access	5.14.1.12
0Dh-6Fh			Reserved	
70h			Discovery (refer to the NVMe over Fabrics specification)	
71h – 7Fh			Reserved for NVMe over Fabrics	
80h – BFh			I/O Command Set Specific	
C0h – FFh			Vendor specific	

O/M: O = Optional, M = Mandatory

...

Add the following new sections as shown below (this goes before NVM Command Set Specific log page identifiers.):

5.14.1.12 Asymmetric Namespace Access (Log Identifier 0Ch)

This log consists of a header describing the log and descriptors containing the asymmetric namespace access information for ANA Groups (refer to section 8.18.2) that contain namespaces that are attached to the controller processing the command. If ANA Reporting (refer to section 8.18) is supported, this log page is supported. ANA Group Descriptors shall be returned in ascending ANA Group Identifier order.

If the RGO bit is cleared to '0' in Command Dword 10, then the LPOL field in Command Dword 12 and the LPOU field in Command Dword 13 of the Get Log Page command should be cleared to 0h.

If the host performs multiple Get Log Page commands to read the ANA log page (e.g., using the LPOL field or the LPOU field), the host should re-read the header of the log page and ensure that the Change Count field in the Asymmetric Namespace Access Log matches the original value read. If it does not match, then the data captured is not consistent and needs to be re-read.

Figure 50a: Command Dword 10 – Log Specific Field

Bit	Description
11:09	Reserved
08	Return Groups Only (RGO): If set to '1', then the controller shall return ANA Group Descriptors with the Number of NSID Values field in each ANA Group Descriptor cleared to 0h (i.e., no Namespace Identifiers are returned). If cleared to '0', then the controller shall return ANA Group Descriptors that contain the Namespace Identifiers of attached namespaces that are members of the ANA Group described by that ANA Group Descriptor and the Number of NSID Values field set to the number of Namespace Identifier values in that ANA Group Descriptor.

Figure 51a: Get Log Page – Asymmetric Namespace Access Log

Bytes	Description
07:00	Change Count: This field contains a 64-bit incrementing Asymmetric Namespace Access log change count, indicating an identifier for this set of asymmetric namespace access information. The count starts at 0h following a controller reset, is incremented each time the contents of the log page change (e.g., not only if an Asymmetric Namespace Access Change AEN is generated), and rolls over to zero when the maximum count is reached.
09:08	Number of ANA Group Descriptors: This field indicates the number of ANA Group Descriptors available in the log page. The log page shall contain one ANA Group Descriptor for each ANA Group that contains Namespaces that are attached to the controller. If, for an ANA Group, there are no namespaces attached to the controller processing the command, then no ANA Group Descriptor is returned for that ANA Group (i.e., an ANA Group Descriptor is returned only if that ANA Group contains namespaces that are attached to the controller processing the command).
15:10	Reserved
n:16	ANA Group Descriptor 0
m:n+1	ANA Group Descriptor 1
...	...
x:y	ANA Group Descriptor n

The format of the ANA Group Descriptor is defined in Figure 52. Namespace Identifiers shall be listed in ascending NSID order.

Figure 52a: ANA Group Descriptor format

Bytes	Description				
03:00	ANA Group ID: The ANA Group Identifier associated with all namespaces in the ANA Group (refer to section 8.18.2) described by this ANA Group Descriptor.				
07:04	Number of NSID Values: This field indicates the number of Namespace Identifier values in this ANA Group Descriptor. If the RGO bit is set to '1', then this field is cleared to 0h.				
15:08	Change Count: This field contains a 64-bit incrementing count, indicating an identifier for the information contained in this ANA Group Descriptor. A value of 0h indicates that the controller does not report a Change Count for this ANA Group Descriptor. If a Change Count is reported, then the count starts at 1h following a controller reset, is incremented each time the data in this ANA Group Descriptor change and rolls over to one when the maximum count is reached. If this field contains 0h, the host should examine this ANA Group Descriptor for any changes and not use this field as an indicator that a change has occurred.				
16	<table> <tr> <th>Bits</th><th>Description</th></tr> <tr> <td>07:04</td><td>Reserved</td></tr> </table>	Bits	Description	07:04	Reserved
Bits	Description				
07:04	Reserved				

Bytes	Description																							
	03:00	Asymmetric Namespace Access State: This field indicates the Asymmetric Namespace Access state for all namespaces in this ANA Group when accessed through this controller. <table><tr><th>Value</th><th>Description</th><th>Reference</th></tr><tr><td>01h</td><td>ANA Optimized state</td><td>8.18.3.1</td></tr><tr><td>02h</td><td>ANA Non-Optimized state</td><td>8.18.3.2</td></tr><tr><td>03h</td><td>ANA Inaccessible state</td><td>8.18.3.3</td></tr><tr><td>04h</td><td>ANA Persistent Loss state</td><td>8.18.3.4</td></tr><tr><td>0Fh</td><td>ANA Change state</td><td>8.18.3.5</td></tr><tr><td>All others</td><td>Reserved</td><td></td></tr></table>		Value	Description	Reference	01h	ANA Optimized state	8.18.3.1	02h	ANA Non-Optimized state	8.18.3.2	03h	ANA Inaccessible state	8.18.3.3	04h	ANA Persistent Loss state	8.18.3.4	0Fh	ANA Change state	8.18.3.5	All others	Reserved	
Value	Description	Reference																						
01h	ANA Optimized state	8.18.3.1																						
02h	ANA Non-Optimized state	8.18.3.2																						
03h	ANA Inaccessible state	8.18.3.3																						
04h	ANA Persistent Loss state	8.18.3.4																						
0Fh	ANA Change state	8.18.3.5																						
All others	Reserved																							
31:17	Reserved																							
35:32	Namespace Identifier 0: The Namespace Identifier of the first namespace that is a member of this ANA Group.																							
39:36	Namespace Identifier 1: The Namespace Identifier of the second namespace, that is a member of this ANA Group.																							
...	...																							
((n*4) + 35): ((n*4) + 32)	Namespace Identifier n: The Namespace Identifier of the n+1 namespace that is a member of this ANA Group.																							

Add table rows and text to the sections as shown below:

5.15 Identify command

The Identify command returns a data buffer that describes information about the NVM subsystem, the controller or the namespace(s). The data structure is 4096 bytes in size.

...

Figure 109: Identify – Identify Controller Data Structure

Bytes	O/M	Description
Controller Capabilities and Features		
75:73	M	IEEE OUI Identifier (IEEE): Contains the Organization Unique Identifier (OUI) for the controller vendor. The OUI shall be a valid IEEE/RAC assigned identifier that may be registered at http://standards.ieee.org/develop/regauth/oui/public.html .
76	O	Controller Multi-Path I/O and Namespace Sharing Capabilities (CMIC): This field specifies multi-path I/O and namespace sharing capabilities of the controller and NVM subsystem. Bits 7: 34 are reserved. Bit 3 if set to '1', then the NVM subsystem supports Asymmetric Namespace Access Reporting (refer to section 8.18). If cleared to '0', then the NVM subsystem does not support Asymmetric Namespace Access Reporting. Bit 2: If set to '1' then the controller is associated with an SR-IOV Virtual Function. If cleared to '0' then the controller is associated with a PCI Function. Bit 1: If set to '1' then the NVM subsystem may be connected to more than one host (e.g., it may contain two or more controllers). If cleared to '0' then the NVM subsystem may only be connected to a single host (e.g., it contains only a single controller). Bit 0: If set to '1' then the NVM subsystem may contain more than one NVM subsystem port. If cleared to '0' then the NVM subsystem contains only a single NVM subsystem port.
...		

Bytes	O/M	Description
95:92	M	<p>Optional Asynchronous Events Supported (OAES): This field indicates the optional asynchronous events supported by the controller. A controller shall not send optional asynchronous events before they are enabled by host software.</p> <p>Bits 31:4012 are reserved.</p> <p>Bit 11 is set to '1' if the controller supports sending Asymmetric Namespace Access Change Notices. If cleared to '0', then the controller does not support the Asymmetric Namespace Access Change Notices event.</p> <p>Bit 10 is reserved.</p> <p>Bit 9 is set to '1' if the controller supports sending Firmware Activation Notices. If cleared to '0', then the controller does not support the Firmware Activation Notices event.</p> <p>Bit 8 is set to '1' if the controller supports sending Namespace Attribute Notices and the associated Changed Namespace List log page. If cleared to '0', then the controller does not support the Namespace Attribute Notices event nor the associated Changed Namespace List log page.</p> <p>Bits 7:0 are reserved.</p>
...		
544 341:340		Reserved
342	O	<p>ANA Transition Time (ANATT): This field indicates the maximum amount of time, in seconds, for a transition between ANA states or the maximum amount of time, in seconds, that the controller reports the ANA change state. If the controller supports Asymmetric Namespace Access Reporting (refer to the CMIC field), then this field shall be set to a non-zero value. If the controller does not support Asymmetric Namespace Access Reporting, then this field shall be cleared to 0h. Refer to section 8.19.4.</p>

Bytes	O/M	Description
343	O	<p>Asymmetric Namespace Access Capabilities (ANACAP): This field indicates the capabilities associated with Asymmetric Namespace Access Reporting (refer to section 8.18).</p> <p>Bit 7 if set to '1', then the controller supports a non-zero value in the ANAGRPID field of the Namespace Management command. If cleared to '0', then the controller does not support a non-zero value in the ANAGRPID field of the Namespace Management command. If the Namespace Management command is not supported, then this bit shall be cleared to '0'.</p> <p>Bit 6 if set to '1' then the ANAGRPID field in the Identify Namespace data structure (refer to Figure 114) does not change while the namespace is attached to any controller. If cleared to '0', then the ANAGRPID field may change while the namespace is attached to any controller. Refer to section 8.18.2.</p> <p>Bit 5 is reserved.</p> <p>Bit 4 if set to '1', then the controller is able to report ANA Change state (refer to section 8.18.3.5). If cleared to '0', then the controller does not report ANA Change state.</p> <p>Bit 3 if set to '1', then the controller is able to report ANA Persistent Loss state (refer to section 8.18.3.4). If cleared to '0', then the controller does not report ANA Persistent Loss state.</p> <p>Bit 2 if set to '1', then the controller is able to report ANA Inaccessible state (refer to section 8.18.3.3). If cleared to '0', then the controller does not report ANA Inaccessible state.</p> <p>Bit 1 if set to '1', then the controller is able to report ANA Non-Optimized state (refer to section 8.18.3.2). If cleared to '0', then the controller does not report ANA Non-Optimized state.</p> <p>Bit 0 if set to '1', then the controller is able to report ANA Optimized state (refer to section 8.18.3.1). If the controller supports Asymmetric Namespace Access Reporting, then this bit is set to '1'.</p>
347:344	O	<p>ANA Group Identifier Maximum (ANAGRPMAX): This field indicates the maximum value of a valid ANA Group Identifier for any controller in the NVM subsystem. If the controller supports Asymmetric Namespace Access Reporting (refer to the CMIC field), then this field shall be set to a non-zero value. If the controller does not support Asymmetric Namespace Access Reporting, then this field shall be cleared to 0h.</p>
351:348	O	<p>Number of ANA Group Identifiers (NANAGRPID): This field indicates the number of ANA groups supported by this controller. If the controller supports Asymmetric Namespace Access Reporting (refer to the CMIC field), then this field shall be set to a non-zero value that is less than or equal to the ANAGRPMAX value. If the controller does not support Asymmetric Namespace Access Reporting, then this field shall be cleared to 0h.</p>
511:340352		Reserved
...		
519:516	M	<p>Number of Namespaces (NN): This field defines the maximum number of namespaces supported by the controller. This field also represents indicates the maximum value of a valid NSID for the controller. If the MNAN field is cleared to 0h, then this field also indicates the maximum number of namespaces supported by the controller.</p>
...		
543:540	O	<p>Maximum Number of Allocated Namespaces (MNAN): This field indicates the maximum number of allocated namespaces. If this field is cleared to 0h, then the maximum number of allocated namespaces is less than or equal to the value in the NN field. If the controller supports Asymmetric Namespace Access Reporting, then this field shall be set to a non-zero value that is less than or equal to the NN value.</p>

Bytes	O/M	Description
767:5440		Reserved
...		

Figure 114 shows the Identify Namespace data structure for the NVM Command Set.

Figure 114: Identify – Identify Namespace Data Structure, NVM Command Set Specific

Bytes	O/M	Description
7:0	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. This field is undefined prior to the namespace being formatted.
15:8	M	<p>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. This field is undefined prior to the namespace being formatted. This field is used in the case of thin provisioning and reports a value that is smaller than or equal to the Namespace Size. Spare LBAs are not reported as part of this field.</p> <p>A logical block is allocated when it is written with a Write or Write Uncorrectable command. A logical block may be deallocated using the Dataset Management, Sanitize, or Write Zeroes command.</p>
23:16	M	<p>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.</p> <p>When using the NVM command set: A logical block is allocated when it is written with a Write or Write Uncorrectable command. A logical block may be deallocated using the Dataset Management, Sanitize, or Write Zeroes command.</p> <p>A controller may report NUSE equal to NCAP at all times if the product is not targeted for thin provisioning environments.</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 section 8.18.3.3) or the ANA Persistent Loss state (refer to section 8.18.3.4), then this field shall be cleared to 0h.</p>
...		
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 Namespace Management and Namespace Attachment commands are 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 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 section 8.18.3.3) or the ANA Persistent Loss state (refer to section 8.18.3.4), then this field shall be cleared to 0h.</p>
40391:64		Reserved

Bytes	O/M	Description
95:92	O	ANA Group Identifier (ANAGRPID): This field indicates the ANA Group Identifier of the ANA group (refer to section 8.18.2) 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. 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.
103:96		Reserved
...		

5.19 Namespace Attachment command

...

5.19.1 Command Completion

When the command is completed, the controller posts a completion queue entry to the Admin Completion Queue indicating the status for the command.

Command specific status values associated with the Namespace Attachment command are defined in Figure 125. For failures, the byte offset of the first failing entry is reported in the Command Specific Information field of the Error Information Log Entry. The controller does not process further entries in the Controller List after an error is encountered.

Figure 125: Namespace Attachment – Command Specific Status Values

Value	Description
18h	Namespace Already Attached: The controller is already attached to the namespace specified.
19h	Namespace Is Private: The controller is not attached to the namespace. The request to attach the controller could not be completed because the namespace is private and is already attached to one controller.
1Ah	Namespace Not Attached: The controller is not attached to the namespace. The request to detach the controller could not be completed.
1Ch	Controller List Invalid: The controller list provided is invalid.
25h	ANA Attach Failed: The controller is not attached to the namespace as a result of an ANA condition (e.g., attaching the controller would result in an ANA Persistent Loss state (refer to section 8.18.3.4)).

5.20 Namespace Management command

The Namespace Management command is used to manage namespaces, including create and delete operations. Note: The controller continues to execute commands submitted to I/O Submission Queues while this operation is in progress.

Host software uses the Namespace Attachment command to attach or detach a namespace to or from a controller. The create operation does not attach the namespace to a controller. As a side effect of the delete operation, the namespace is detached from any controller as it is no longer present in the system. It is recommended that host software detach all controllers from a namespace prior to deleting the namespace. If the namespace is attached to another controller when a delete operation is requested, then as part of the delete operation a Namespace Attribute Notice is issued by that controller to indicate a namespace change (if Namespace Attribute Notices are supported).

The data structure used for the create operation is defined in Figure 129 and has the same format as the Identify Namespace data structure defined in Figure 114. After successful completion of a Namespace Management command with the create operation, the namespace is formatted with the specified attributes. The fields that host software may specify in the create operation are defined in Figure 126. Fields that are reserved shall be cleared to 0h by host software. There is no data structure transferred for the delete operation.

Figure 126: Namespace Attachment – Host Software Specified Fields

Bytes	Description	Host Specified
7:0	Namespace Size (NSZE)	Yes
15:8	Namespace Capacity (NCAP)	Yes
25:16	Reserved	
26	Formatted LBA Size (FLBAS)	Yes
28:27	Reserved	
29	End-to-end Data Protection Type Settings (DPS)	Yes
30	Namespace Multi-path I/O and Namespace Sharing Capabilities (NMIC)	Yes
91:31	Reserved	
95:92	ANA Group Identifier (ANAGRPID) ¹	Yes
383:3196	Reserved	
Notes:		
1. A value of zero specifies that the controller determines the value to use (refer to section 8.12).		

The Namespace Management command uses the Data Pointer and Dword 10 fields. All other command specific fields are reserved.

...

5.20.1 Command Completion

When the command is completed, the controller posts a completion queue entry to the Admin Completion Queue indicating the status for the command.

Namespace Management command specific status values are defined in Figure 130.

Figure 130: Namespace Management – Command Specific Status Values

Value	Description
0Ah	Invalid Format: The LBA Format specified is not supported. This may be due to various conditions, including: 1) specifying an invalid LBA Format number, or 2) enabling protection information when there is not sufficient metadata per LBA, or 3) the specified format is not available in the current configuration, or 4) invalid security state (refer to TCG SIIS), etc.
15h	Namespace Insufficient Capacity: Creating the namespace requires more free space than is currently available. The Command Specific Information field of the Error Information Log specifies the total amount of NVM capacity required to create the namespace in bytes.
16h	Namespace Identifier Unavailable: The number of namespaces supported has been exceeded.
1Bh	Thin Provisioning Not Supported: Thin provisioning is not supported by the controller.
24h	ANA Group Identifier Invalid: The specified ANA Group Identifier (ANAGRPID) is not supported in the submitted command. This may be due to various conditions, including: a) specifying an ANAGRPID that does not exist; b) the controller does not allow an ANAGRPID to be specified (i.e., bit 7 in the ANACAP field is cleared to '0'); or c) the specified ANAGRPID is not supported by the controller processing the command (e.g., the specified value exceeds ANAGRPMAX (refer to Figure 109)). If the host specified a non-zero ANAGRPID, retrying the command with the ANAGRPID field cleared to 0h may succeed.

...

5.21 Set Features command

...

5.21.1.11 Asynchronous Event Configuration (Feature Identifier 0Bh)

This Feature controls the events that trigger an asynchronous event notification to the host. ...

Figure 148: Asynchronous Event Configuration – Command Dword 11

Bit	Description
31:412	Reserved
11	Asymmetric Namespace Access Change Notices: This bit determines whether an asynchronous event notification is sent to the host when an asymmetric namespace access change occurs (i.e., the contents of the Asymmetric Namespace Access log page (refer to section 5.14.1.12) change). If this bit is set to '1', then the Asymmetric Namespace Access Change Notices event is sent to the host when this condition occurs. If this bit is cleared to '0', then the controller shall not send the Asymmetric Namespace Access Change Notices event to the host.
10	Telemetry Log Notices: This field determines whether an asynchronous event notification is sent to the host when the Telemetry Controller-Initiated Data Available field transitions from '0' to '1' in the Telemetry Controller-Initiated log page. If this bit is set to '1', then the Telemetry Log Changed event is sent to the host when this condition occurs. If this bit is cleared to '0', then the controller shall not send the Telemetry Log Changed event to the host.
...	

...

Make changes to the following sections as shown below (changes to section 6.1.6 are based on changes made in TP4018):

6.1 Namespaces

A namespace is a collection of logical blocks that range from 0 to the capacity of the namespace – 1. A namespace ID (NSID) is an identifier used by a controller to provide access to a namespace.

...

6.1.6 NSID and Namespace Usage

If Namespace Management (refer to section 8.12), ANA Reporting (refer to section 8.18), or NVM Sets (refer to section 4.9) are supported, (refer to the OACS field in Figure 110) then NSIDs shall be unique within the NVM subsystem (e.g., NSID of 3 shall refer to the same physical namespace regardless of the accessing controller). If Namespace Management, ANA Reporting, and NVM Sets are not supported, then NSIDs:

- for shared namespaces shall be unique; and
- for private namespaces are not required to be unique.

The Identify command (refer to section 5.15) may be used to determine the active NSIDs for a controller and the allocated NSIDs in the NVM subsystem.

If the MNAN field (refer to Figure 109) is cleared to 0h, then the maximum number of allocated NSIDs is the same as the value reported in the NN field (refer to Figure 109). If the MNAN field is non-zero, then the maximum number of allocated NSIDs may be less than the number of namespaces (e.g., an NVM subsystem may support a maximum valid NSID value (i.e., the NN field) set to 1,000,000 but support a maximum of 10 allocated NSID values).

To determine the active NSIDs for a particular controller, the host may follow either of the following methods:

...

The Namespace Size (**NSZE**) field in the Identify Namespace data structure defines the total size of the namespace in logical blocks (LBA 0 through $n-1$). The Namespace Utilization (**NUSE**) field in the Identify Namespace data structure defines the number of logical blocks currently allocated in the namespace. The Namespace Capacity (**NCAP**) field in the Identify data structure defines the maximum number of logical blocks that may be allocated at one time as part of the namespace in a thin provisioning usage model. The following relationship holds: Namespace Size \geq Namespace Capacity \geq Namespace Utilization.

If the controller supports Asymmetric Namespace Access Reporting (i.e., bit 3 set to '1' in the CMIC field in the Identify Controller data structure (refer to Figure 109)), then the NUSE field (refer to Figure 114) and the NVMCAP field (refer to Figure 114) are cleared to 0h if the relationship between the controller and the namespace is in the ANA Inaccessible state (refer to section 8.18.3.3) or the ANA Persistent Loss state (refer to section 8.18.3.4). The Namespace Attribute Changed event is not generated for changes to these fields that result from ANA state changes as described in Figure 49. The host uses the Asymmetric Namespace Access Change Notices as an indication of these changes.

A namespace may or may not have a relationship to a Submission Queue ...

Make changes to section 8.12 (Namespace Management) as shown below:

8.12 Namespace Management (Optional)

The Namespace Management command is used to create a namespace or delete a namespace. The Namespace Attachment command is used to attach and detach controllers from a namespace. Namespace management is intended for use during manufacturing or by a system administrator.

...

The total and unallocated NVM capacity for the NVM subsystem is reported in the Identify Controller data structure (refer to Figure 109). For each namespace, the NVM capacity used for that namespace is reported in the Identify Namespace data structure (refer to Figure 114). The controller may allocate NVM capacity in units such that the requested size for a namespace may be rounded up to the next unit boundary. For example, if host software requests a namespace of 32 logical blocks with a logical block size of 4KB for a total size of 128KB and the allocation unit for the implementation is 1MB then the NVM capacity consumed may be rounded up to 1MB. The NVM capacity fields may not correspond to the logical block size multiplied by the total number of logical blocks.

The method of allocating ANA Group identifiers is outside the scope of this specification. If the ANA Group Identifier (refer to Figure 126) is cleared to 0h, then the controller shall determine the ANAGRPID that is assigned to that namespace.

To create a namespace, host software performs the following actions:

...

8.15 Sanitize Operations (Optional)

...

8.15.1 Command Restrictions

...

Figure 131: Sanitize Operations – Admin Commands Allowed

Admin Command	Additional Restrictions	
...		
Get Log Page	The log pages allowed are listed below.	
	Log Pages	Additional Restrictions
	Error Information	Return zeros in the LBA field.
	SMART / Health Information	
	Changed Namespace List	
	Reservation Notification	
	Sanitize Status	
	Asymmetric Namespace Access	
...		

Add the following new section as shown below (this becomes the last subsection of section 8):

8.18 Asymmetric Namespace Access Reporting (Optional)

8.18.1 Asymmetric Namespace Access Reporting Overview

Asymmetric namespace access (ANA) occurs in environments where namespace access characteristics (e.g., performance or ability to access the media) may vary based on the controller used to access the namespace (e.g., Fabrics) and the internal configuration of the NVM subsystem. Asymmetric Namespace Access Reporting is used to indicate to the host information about those access characteristics.

Shared namespaces may be accessed through controllers via multiple PCIe ports (refer to section 1.4.1) or fabric ports (refer to the NVMe over Fabrics specification). The controllers that provide access to a shared namespace may provide identical access characteristics through all controllers (i.e., symmetric access), or may provide different access characteristics through some controllers (i.e., asymmetric access).

Private namespaces are accessed by only one controller at a time. The access characteristics of the namespace through that controller may be impacted as a result of changes to the internal configuration of the NVM Subsystem. If the access characteristics of the namespace through that controller are impacted by the internal configuration of the NVM subsystem, then asymmetric access occurs.

Symmetric access to a namespace occurs when:

- the access characteristics using one controller are identical to the access characteristics when using a different controller; and
- changes to the internal configuration of the NVM subsystem do not impact the access characteristics.

Asymmetric access to a namespace occurs when:

- the access characteristics using one controller may differ from the access characteristics when using a different controller; or
- changes to the internal configuration of the NVM subsystem may impact the access characteristics.

While commands may be sent to a shared namespace through any attached controller with asymmetric access, the characteristics (e.g., performance or ability to access the media) may differ based on which controller is used; as a result, the host should consider those characteristics when selecting which controller to use for each command that accesses the namespace. The NVM subsystem may perform autonomous internal reconfiguration that results in a change to the access characteristics.

If an NVM subsystem supports Asymmetric Namespace Access Reporting, then all controllers in that NVM subsystem shall:

- set bit 3 to '1' in the Controller Multi-path I/O and Namespace Sharing Capabilities (CMIC) field in the Identify Controller data structure (refer to Figure 109) to indicate support for Asymmetric Namespace Access Reporting;
- set bit 0 to '1' in the Asymmetric Namespace Access Capabilities (ANACAP) field in the Identify Controller data structure to indicate that the ANA Optimized state is able to be reported;
- set bit 1 to '1' in the ANACAP field in the Identify Controller data structure if ANA Non-Optimized state is able to be reported;
- set bit 2 to '1' in the ANACAP field in the Identify Controller data structure if ANA Inaccessible state is able to be reported;
- set bit 3 to '1' in the ANACAP field in the Identify Controller data structure if ANA Persistent Loss state is able to be reported;

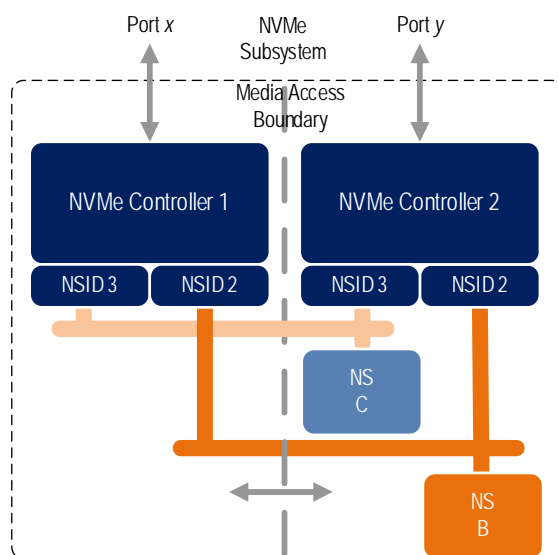
- set bit 4 to '1' in the ANACAP field in the Identify Controller data structure if ANA Change state is able to be reported;
- support Asymmetric Namespace Access Change Notices (refer to section 5.21.1.11); and
- support the Asymmetric Namespace Access log page (refer to section 5.14.1.12).

Namespaces attached to a controller that supports Asymmetric Namespace Access Reporting shall:

- be members of an ANA Group; and
- supply a valid ANA Group Identifier in the ANA Group Identifier (ANAGRPID) field in Identify Namespace data structure (refer to Figure 114).

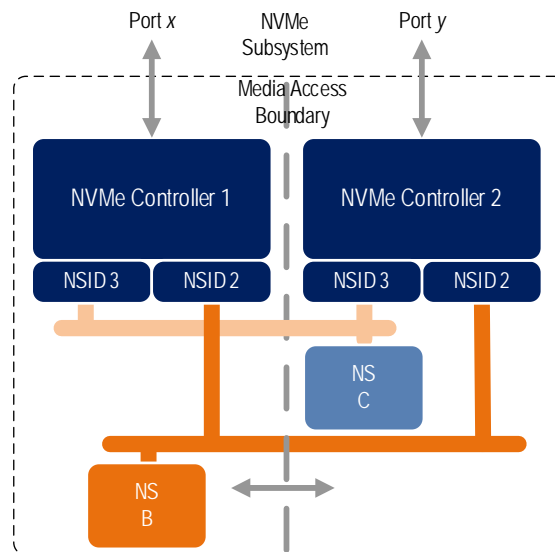
Figure 1000 shows an example of an NVM subsystem where access characteristics vary as a result of the presence of two independent domains. In this example, the non-volatile media for namespace B and for namespace C are contained within the same domain that contains controller 2. As a result, controller 2 provides optimized access to namespace B and to namespace C while controller 1 does not provide optimized access to namespace B or to namespace C.

Figure 1000: Namespace B and C optimized through Controller 2



To provide optimized access to namespace B through controller 1, the NVM subsystem may be administratively reconfigured, or may perform autonomous internal reconfiguration actions that change the access characteristics of namespace B when accessed through controller 1 and controller 2 as shown in Figure 1001. Controller 2 provides optimized access to namespace C while controller 1 provides optimized access to namespace B.

Figure 1001: Namespace B optimized through Controller 1



8.18.2 ANA Groups

Namespaces that are members of the same ANA Group perform identical asymmetric namespace access state transitions. The ANA Group maintains the same asymmetric namespace access state for all namespaces that are members of that ANA Group (i.e., a change in the asymmetric namespace access state of one namespace only occurs as part of a change in the asymmetric namespace access state of all namespaces that are members of that ANA Group). The method for assigning namespaces to ANA Groups is outside the scope of the specification.

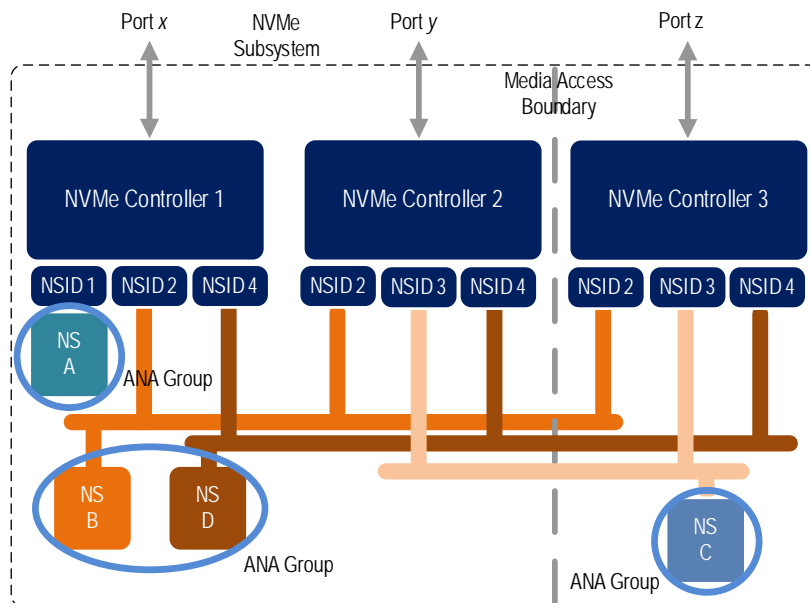
A valid ANA Group Identifier is a non-zero value that is less than or equal to ANAGRPMAX (refer to Figure 109).

The ANA Group Identifier (ANAGRPID) for each ANA Group shall be unique within the NVM Subsystem. If bit 6 in the ANACAP field in the Identify Controller data structure is set to '1', then the ANA Group Identifier shall not change while the namespace is attached to any controller in the NVM subsystem. If bit 6 in the ANACAP field is cleared to '0', then the ANA Group Identifier may change while the namespace is attached to any controller in the NVM subsystem. If the ANA Group Identifier changes, the controller shall issue the Asymmetric Namespace Access Change Notice as described in 8.18.3.6.

Figure 1002 shows the following four namespaces:

- the private namespace A in a first ANA Group;
- namespace B and namespace D, that are in the same second ANA Group; and
- namespace C that is in a third ANA Group.

Figure 1002: Multiple Namespace groups



8.18.3 Asymmetric Namespace Access states

The Asymmetric Namespace Access State indicates information about the characteristics of the relationship between a controller and an ANA Group. The following asymmetric namespace access states are defined:

- ANA Optimized (refer to section 8.18.3.1);
- ANA Non-Optimized (refer to section 8.18.3.2);
- ANA Inaccessible (refer to section 8.18.3.3);
- ANA Persistent Loss (refer to section 8.18.3.4); and
- ANA Change (refer to section 8.18.3.5).

8.18.3.1 ANA Optimized state

While the relationship between the controller and an ANA group is in this state, the characteristic of that relationship to each namespace in that group is optimized. Commands processed by a controller that reports this state for an ANA Group provide optimized access characteristics to any namespace in that ANA Group. A controller that supports ANA Reporting shall support reporting this state.

While in this state, all commands, functions, and operations supported by the namespace shall perform as described in this specification.

8.18.3.2 ANA Non-Optimized state

While the relationship between the controller and an ANA group is in this state, the characteristic of that relationship to each namespace in that group is non-optimized. Commands processed by a controller that reports this state for an ANA Group provide non-optimized access characteristics (e.g., the processing of some commands, especially those involving data transfer, may operate with lower performance or may use NVM subsystem resources less effectively than if a controller is used that reports the optimized state) to any namespace in that ANA Group. Support for reporting this state is optional.

While in this state, all commands, functions, and operations supported by the namespace shall perform as described in this specification.

8.18.3.3 ANA Inaccessible state

While the relationship between the controller and an ANA group is in this state, the characteristic of that relationship to each namespace in that group is inaccessible. Commands processed by a controller that reports this state for an ANA Group are not able to access user data of namespaces in that ANA Group. The namespaces may become accessible through the controller reporting this state at a future time (i.e., a subsequent ANA state transition may occur). Support for reporting this state is optional.

While in this state, accurate namespace related capacity information may not be available. As a result, some namespace capacity information returned in the Identify namespace data (i.e., the NUSE field and the NVMeCAP field), are cleared to 0h. For that namespace capacity information, hosts should use the Identify namespace data returned from a controller that reports the relationship between the controller and the namespace to be in the ANA Optimized state or in the ANA Non-Optimized state.

A controller shall fail commands, other than those described in 8.18.4, with a status code of Asymmetric Namespace Access Inaccessible if those commands are submitted while the relationship between the namespace specified by the command and the controller processing the command is in this state.

While ANA Inaccessible state is reported by a controller for the namespace, the host should retry the command on a different controller that is reporting ANA Optimized state or ANA Non-Optimized state. If no controllers are reporting ANA Optimized state or ANA Non-Optimized state, then a transition may be occurring such that a controller reporting the Inaccessible state may become accessible and the host should retry the command on the controller reporting Inaccessible state for at least ANATT seconds (refer to Figure 109). Refer to section 8.19.2.

8.18.3.4 ANA Persistent Loss state

While the relationship between the controller and an ANA group is in this state, the characteristic of that relationship to each namespace in that group is persistently inaccessible. Commands processed by a controller that reports this state for an ANA Group are persistently not able to access user data of namespaces in that ANA Group. The relationship between a controller and an ANA Group in this state shall not transition to any other ANA state. Support for reporting this state is optional.

While in this state, accurate namespace related capacity information may not be available. As a result, some namespace capacity information returned in the Identify namespace data (i.e., the NUSE field and the NVMeCAP field), are cleared to 0h. For that namespace capacity information, hosts should use the Identify namespace data returned from a controller that reports the relationship between the controller and the namespace to be in the ANA Optimized state or in the ANA Non-Optimized state.

A controller shall fail commands, other than those described in 8.18.4, with a status code of Asymmetric Namespace Access Persistent Loss if those commands are submitted while the relationship between the namespace specified by the command and the controller processing the command is in this state.

While ANA Persistent Loss state is reported by a controller for the namespace, the host should retry the command on a different controller that is reporting ANA Optimized state or ANA Non-Optimized state. If no controllers are reporting ANA Optimized state or ANA Non-Optimized state, then a transition may be occurring such that a controller reporting the Inaccessible state may become accessible and the host should retry the command on the controller reporting Inaccessible state for at least ANATT seconds (refer to Figure 109).

8.18.3.5 ANA Change state

The change from one asymmetric namespace access state to another asymmetric namespace access state is called a transition. Transitions may occur in such a way that the ANA Change state is not visible to the host (i.e., the ANA Change state may or may not be reported in the Asymmetric Namespace Access State field in the Asymmetric Namespace Access log (refer to section 5.14.1.12)). Support for reporting this state is optional.

A controller shall fail commands, other than those described in 8.18.4, with a status code of Asymmetric Namespace Access Transition if those commands are submitted while the relationship between the namespace specified by the command and the controller processing the command is in this state.

While ANA Change state is reported by a controller for the namespace, the host should:

- a) after a short delay, retry the command on the same controller for at least ANATT (refer to Figure 109) seconds (e.g., if ANATT is 30, perform 3 retries at 10 s intervals, or 10 retries at 3 s intervals); or
- b) retry the command on a different controller that is reporting ANA Optimized state or ANA Non-Optimized state.

8.18.3.6 Asymmetric Namespace Access Change Notifications

If Asymmetric Namespace Access Change Notices are enabled (refer to section 5.21.1.11) on a controller, then an Asymmetric Namespace Access Change Notice shall be sent by the controllers where the change occurred:

- a) if an ANA Group Identifier (refer to Figure 109) changes;
- b) if an asymmetric namespace access state transition fails (e.g., a transition begins, but does not complete and the controller returns to the state that existed before the transition began); or
- c) upon entry to the following ANA states:
 - ANA Optimized State;
 - ANA Non-Optimized State;
 - ANA Inaccessible State; and
 - ANA Persistent Loss State.

8.18.4 Asymmetric Namespace Access States Command Processing Effects

Processing of Admin commands that:

- are not NVM Command Set Specific commands; and
- do not use the Namespace Identifier (i.e., Figure 41 – “Namespace Identifier Used” column indicates “No”),

are not affected by ANA states, except as specified in Figure 1003.

Figure 1003 describes Asymmetric Namespace Access effects on command processing.

Figure 1003: ANA effects on Command Processing

Command	ANA State	Effects on command processing
Get Features	ANA Inaccessible, ANA Persistent Loss, or ANA Change	<p>The following feature identifiers are not available¹:</p> <ol style="list-style-type: none"> a) LBA Range Type (i.e., 03h); b) Error Recovery (i.e., 05h); c) Write Atomicity Normal (i.e., 0Ah); d) Reservation Notification Mask (i.e., 82h); and e) Reservation Persistence (i.e., 83h).
Get Log Page	ANA Inaccessible, ANA Persistent Loss, or ANA Change	<p>The following log pages are affected:</p> <ol style="list-style-type: none"> a) Error Information (i.e., 01h): The log page may contain entries only for namespaces whose relationship to the controller processing the command is in the ANA Optimized state (refer to section 8.18.3.1) or the ANA Non-optimized state (refer to section 8.18.3.2).

Command	ANA State	Effects on command processing
Identify	ANA Inaccessible or ANA Persistent Loss	Capacity fields in the Identify namespace data structure (refer to Figure 114) information is cleared to zero.
Set Features	ANA Inaccessible	<p>The saving of features shall not be supported and the following feature identifiers are not available¹:</p> <ul style="list-style-type: none"> a) LBA Range Type (i.e., 03h); b) Error Recovery (i.e., 05h); c) Write Atomicity Normal (i.e., 0Ah) d) Reservation Notification Mask (i.e., 82h); and e) Reservation Persistence (i.e., 83h). <p>If the NSID is set to FFFFFFFFh, then the command shall fail with a status code of Asymmetric Namespace Access Inaccessible (refer to section 8.18.3.3).</p>
	ANA Change	<p>The saving of features shall not be supported and the following feature identifiers are not available¹:</p> <ul style="list-style-type: none"> a) LBA Range Type (i.e., 03h); b) Error Recovery (i.e., 05h); c) Write Atomicity Normal (i.e., 0Ah); d) Reservation Notification Mask (i.e., 82h); and e) Reservation Persistence (i.e., 83h). <p>If the NSID is set to FFFFFFFFh, then the command shall fail with a status code of Asymmetric Namespace Access Transition (refer to section 8.18.3.5).</p>
	ANA Persistent Loss	This command shall fail with a status code of Asymmetric Namespace Access Persistent Loss (refer to section 8.18.3.4).
<p>Notes:</p> <ol style="list-style-type: none"> 1. If the ANA state is ANA Inaccessible State, then commands that use feature identifiers that are not available shall fail with a status code of Asymmetric Namespace Access Inaccessible. If the ANA state is ANA Persistent Loss State, then commands that use feature identifiers that are not available shall fail with a status code of Asymmetric Namespace Access Persistent Loss. If the ANA state is ANA Change State, then commands that use feature identifiers that are not available shall fail with a status code of Asymmetric Namespace Access Transition. 		

8.19 Host Operation with Asymmetric Namespace Access Reporting (Informative)

8.19.1 Host ANA Normal Operation

The host determines if ANA is supported by examining bit 3 in the CMIC field in the Identify Controller data structure (refer to Figure 109). The NSID or Identifier (refer to section 7.10) is used to determine when multiple paths to the same namespace are available. The host examines the ANA Log page (refer to section 5.14.1.12) for each controller to determine the ANA state of each group of namespaces attached to that controller.

To send a command to a namespace, the host should select a controller that reports the ANA Optimized State (refer to section 8.18.3.1) and send the command to that controller. If more than one controller that reports the ANA Optimized state for a namespace are found, then the host may use all of those controllers to send commands.

If there are no controllers that report the ANA Optimized state for a namespace, then the host should select a controller that reports ANA Non-Optimized State (refer to section 8.18.3.2) for that namespace and send the command to that controller. If more than one controller that reports ANA Non-Optimized state for a namespace are found, then the host may use all of those controllers to send commands.

If multiple controllers are being used, the algorithm for determining which controller to use next is outside the scope of this specification (e.g., the host may select a simple round robin algorithm, a queue depth weighted algorithm, a transfer length weighted algorithm, or any other algorithm).

If there are no controllers that report the ANA Optimized state for a namespace and there are no controllers that report the ANA Non-Optimized state for that namespace, then the host should examine controllers that report the ANA Inaccessible state as described in section 8.19.2.

8.19.2 Host ANA Inaccessible Operation

If the ANA Log page reports an ANA state of ANA Inaccessible State for an ANA Group or a command returns a status of Asymmetric Namespace Access Inaccessible, then the host should:

- not use that controller to send commands to any namespace in that ANA Group; and
- select a different controller for sending commands to all namespaces in that ANA Group.

If there are no controllers that report the ANA Optimized state for a namespace and there are no controllers that report the ANA Non-Optimized state, then a transition may be occurring that also impacts controllers that are reporting the ANA Inaccessible state. As a result, the host should use the methods described for Host ANA Transition operation (refer to section 8.19.4) to determine if the controller reporting ANA Inaccessible state transitions during the ANATT time interval to an ANA state that enables commands to be processed by that controller.

8.19.3 Host ANA Persistent Loss Operation

If the ANA Log page reports an ANA state of ANA Persistent Loss State for an ANA Group or a command returns a status of Asymmetric Namespace Access Persistent Loss, then the host should not use that controller to send commands to any namespace in that ANA Group, and select a different controller for sending commands to any namespace in that ANA Group. If the controller supports Namespace Management, then the namespaces in an ANA Group reporting this state should be detached.

8.19.4 Host ANA Transition Operation

If the ANA Log page reports an ANA state of ANA Change State for an ANA Group or a command returns a status of Asymmetric Namespace Access Transition, then the host should temporarily not use that controller to send commands to any namespace in that ANA Group. If only controllers reporting ANA Inaccessible State are available, then the host should follow these procedures to determine which controller to use. To use a controller, the host may:

- a) if Asymmetric Namespace Access Change Notices are enabled (refer to section 5.21.1.11) on the controller, wait for an Asymmetric Namespace Access Change Notice from that controller. Upon receipt of that notice, the host should examine the ANA Log page to determine the new ANA state and resume sending commands based on the new ANA state. Such notice should occur within the ANATT time (refer to Figure 109); or
- b) delay and retry the command during the ANATT time interval. The host should not immediately retry, but rather, divide the ANATT time into equal intervals for command retry purposes (e.g., if ANATT is 30, perform 3 retries at 10 s intervals, or 10 retries at 3 s intervals). During or upon completion of the ANATT time interval, the new ANA state of the ANA Group should be known (e.g., one of the command retries returned a different status that indicates completion of the transition to a new ANA state). If the retried command did not complete without error, the ANA Log page should be examined on each controller that provides access to the namespace and the host should resume sending commands based on the new ANA state.

If the ANATT time interval expires, then the host should use a different controller for sending commands to the namespaces in that ANA Group. The ANATT interval reported by the controller should prevent this type of timer expiration from occurring.

8.19.5 Host ANA Change Notice Operation

Receipt of an Asymmetric Namespace Access Change Notice from a controller may indicate:

- a) that the ANA state reported in one or more ANA Group Descriptors has changed;
- b) a new NSID has been added to one or more of the ANA Group Descriptors;
- c) an NSID has been removed from one or more of the ANA Group Descriptors; and/or
- d) the NSID of a namespace has moved from one ANA Group Descriptor to a different ANA Group Descriptor (i.e., the ANAGRPID field in the Identify Namespace data structure for that namespace has changed), if bit 6 in the ANACAP field is cleared to '0' in the Identify Controller data structure (refer to Figure 109).

As a result of receiving an Asymmetric Namespace Access Change Notice, the host should read the ANA Log page (refer to section 5.14.1.12) to check for each of those possible changes.

8.19.6 All Paths Down Condition

An all paths down condition occurs when there are no paths available on the host to access the namespaces in an ANA Group (i.e., the NVM media). To determine whether an all paths down condition has occurred, the host may examine the ANA Log page on each controller that provides access to the namespaces in a particular ANA Group. All paths that are not in the ANA Persistent Loss state should be checked. If no paths to the namespaces in that ANA Group become available (i.e., transition to the ANA Optimized state or the ANA Non-Optimized state) for the duration of an ANATT time interval, then an all paths down condition has occurred for the namespaces in that ANA Group.