



LEGAL NOTICE:

© **Copyright 2008 to 2023 NVM Express®**, Inc. **ALL RIGHTS RESERVED.**

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

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

NOTICE TO NON-MEMBERS OF NVM EXPRESS, INC.: If you are not a Member of NVM Express, Inc. and you have obtained a copy of this document, you only have a right to review this document or make reference to or cite this document. Any such references or citations to this document must acknowledge NVM Express, Inc. copyright ownership of this document. The proper copyright citation or reference is as follows: “© **2008 to 2023 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 4156a Reachability Architecture
Change Date	2023-09-07
Builds on Specification	NVM Express Base Specification 2.0c NVMe Management Interface Specification 1.2.c
References Specification	TP4130 Cross Namespace Copy TP4091 Computational Programs TP4131 Subsystem Local Memory

Technical Proposal Author(s)

Name	Company
Bill Martin	Samsung
David Black	Dell EMC
Kim Malone	Intel
Stephen Bates	Eideticom

This proposal intends to:

- Define a reachability architecture that applies to Cross Namespace Copy, Computational Programs, and Subsystem Local Memory

Revision History

Revision Date	Change Description
2022-10-23	<ul style="list-style-type: none">• Initial rough version
2023-01-27	<ul style="list-style-type: none">• Revision based on decisions of TWG
2023-02-08	<ul style="list-style-type: none">• Added Asynchronous Events
2023-02-10	<ul style="list-style-type: none">• Removed previous markup that were discussed in TWG meeting• changes from TWG meeting<ul style="list-style-type: none">○ removed all Reachability Association Characteristics except 00h
2023-03-01	<ul style="list-style-type: none">• Changes agreed to by FMDS TG related to:<ul style="list-style-type: none">○ adding bit 1 in CRCAP○ modified figure 147 in support of keeping RGRPID in the Identify – I/O Command Set Independent Identify Namespace Data Structure○ added discussion in 8.TBD.1 about changing RGRPID while a namespace is attached to any controller
2023-05-17	<ul style="list-style-type: none">• Accepted all previous changes and removed comments that are resolved• Changed Reachability Characteristics to point to fast copy operation in the NVM Command Set Specification• Added Namespace Not Reachable status code for I/O Command Set specific status codes• Added necessary changes for NVMe MI specification
2023-06-27	<ul style="list-style-type: none">• Member Review comment resolution
2023-08-03	<ul style="list-style-type: none">• Integrated
2023-08-04	<ul style="list-style-type: none">• Bolded the acronym NRID and moved it before the colon.• Fixed the Reachability Group Identifier NRID-1 upper byte offset.
2023-09-01	<ul style="list-style-type: none">• Updated to match the template and include NVMe-MI as an affected specification.

New Feature:

- Reachability Architecture (optional)
 - Description of change.
 - Create a Reachability Groups log page
 - Create a Reachability Associations log page
 - Defined the Reachability Architecture
 - New Requirement
 - Reachability Groups log page and Create a Reachability Associations log page are mandatory id Reachability Reporting is supported.
 - References
 - TP4130: Cross-Namespace Copy
 - TP4091 Computational Programs
 - TP4131 Subsystem Local Memory
 - TP4156 Reachability Architecture

Markup Conventions:

- Black: Unchanged (however, hot links are removed)
- ~~Red Strikethrough~~: Deleted
- Blue: New
- Blue Highlighted: TBD values, anchors, and links to be inserted in new text
- Purple: Text moved without change.

<Green Bracketed>: Notes to editor

Modify portions of NVM Express Base Specification 2.0c as follows:

Modify section 3 as follows:

3 NVM Express Architecture

3.1 NVM Controller Architecture

...

3.1.2 Controller Types

...

3.1.2.1 I/O Controller

...

3.1.2.1.2 Log Page Support

...

Figure 24: I/O Controller – Log Page Support

Log Page Name	Log Page Support Requirements ¹
...	
Rotational Media Information	O
Reachability Groups	M TBD
Reachability Associations	M TBD
...	...
Notes: 1. O/M/P definition: O = Optional, M = Mandatory, P = Prohibited 2. Mandatory for controllers that support Fixed Capacity Management (refer to section 8.3.2). 3. Optional for NVM Express revision 1.4 and earlier. TBD. Optional for controllers that do not support Reachability Reporting (refer to section 8.TBD).	

...

3.1.2.2 Administrative Controller

...

3.1.2.2.2 Log Page Support

...

Figure 32: Administrative Controller – Log Page Support

Log Page Name	Command Support Requirements ¹
...	
Rotational Media Information	P
Reachability Groups	P
Reachability Associations	P

Figure 32: Administrative Controller – Log Page Support

Log Page Name	Command Support Requirements ¹
...	
Notes:	
<ol style="list-style-type: none"> O/M/P definition: O = Optional, M = Mandatory, P = Prohibited Optional if Set Features command is not supported (refer to Figure 31). Optional if NVMe-MI Send command and NVMe-MI Receive command is not supported (refer to Figure 31). Optional for NVM Express revision 1.4 and earlier. Prohibited for an Exported NVM Subsystem (refer to section 8.12). If this log page is not described for a specific physical interface (refer to the applicable NVM Express transport specification), then this log page is prohibited for that transport. 	

...

3.1.2.3 Discovery Controller

...

3.1.2.3.3 Log Page Support

...

Figure 37: Discovery Controller – Log Page Support

Log Page Name	Command Support Requirements ¹
...	
Rotational Media Information	P
Reachability Groups	P
Reachability Associations	P
...	
Notes:	
<ol style="list-style-type: none"> O/M/P definition: O = Optional, M = Mandatory, P = Prohibited Optional if Set Features command is not supported (refer to Figure 36). Optional for versions 1.1 and earlier of the NVMe over Fabrics specification. Mandatory for CDCs and prohibited for Discovery controllers that are not a CDC. 	

...

3.3 NVM Queue Models

...

3.3.3 Queueing Data Structures

...

3.3.3.2 Common Completion Queue Entry

...

3.3.3.2.1 Status Field Definition

...

3.3.3.2.1.1 Generic Command Status Definition

...

Figure 96: Status Code – Command Specific Status Values, I/O Commands

Value	Description
80h	Conflicting Attributes
81h	Invalid Protection Information
82h	Attempted Write to Read Only Range
83h	Command Size Limit Exceeded
84h	Invalid Command ID
88h	Namespace Not Reachable
89h 85h to B7h	Reserved
...	
Notes:	
1. A = All I/O Command Sets, C = Command Set Specific.	

...

Modify section 5 as follows:

5 Admin Command Set

...

5.2 Asynchronous Event Request command

...

Figure 147: Asynchronous Event Information – Notice

Value	Description
00h	<p>Namespace Attribute Changed: Indicates a change to one or more of the following:</p> <ul style="list-style-type: none"> the Identify Namespace data structure (refer to the applicable NVMe I/O Command Set specification) for one or more namespaces; the I/O Command Set Independent Identify Namespace data structure; the Namespace List returned when the Identify command is issued with the CNS field set to 02h; or other data structures as specified in applicable NVMe I/O Command Set specifications. <p>To clear this event, host software issues a Get Log Page command for the Changed Namespace List log page (refer to section 5.16.1.5) with the Retain Asynchronous Event bit cleared to '0'.</p> <p>A controller shall not send this event if:</p> <ol style="list-style-type: none"> Namespace Status (refer to Figure 280) has changed and shutdown processing is either occurring (i.e., CSTS.SHST is set to 01b) or complete (i.e., CSTS.SHST is set to 10b); the ANAGRPID field (refer to Figure 280) has changed; or the RGRPID field (refer to figure 280) has changed; or an I/O Command Set specific change occurs (refer to the applicable I/O Command Set specification). <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 0h, or from 0h 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>

Figure 147: Asynchronous Event Information – Notice

Value	Description
02h	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 command with Retain Asynchronous Event bit cleared to '0' for the Telemetry Controller-Initiated log.
03h	Asymmetric Namespace Access Change: The Asymmetric Namespace Access information (refer to section 5.16.1.13) 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 5.16.1.13). To clear this event, the host issues a Get Log Page command (refer to section 5.16) with the Retain Asynchronous Event bit cleared to '0' for the Asymmetric Namespace Access log. A controller shall not send this event if a Namespace Attribute Changed notice is sent for the same event, such as a change due to: <ul style="list-style-type: none"> a) the attachment of a namespace (refer to section 5.22); b) the deletion of a namespace (refer to section 5.23); or c) the detachment of a namespace (refer to section 5.22).
04h	Predictable Latency Event Aggregate Log Change: Indicates that event pending entries for one or more NVM Sets (refer to section 5.16.1.12) have been added to the Predictable Latency Event Aggregate log.
05h ¹	LBA Status Information Alert: I/O Command Set specific definition.
06h	Endurance Group Event Aggregate Log Page Change: Indicates that event entries for one or more Endurance Groups (refer to section 5.16.1.10) have been added to the Endurance Group Event Aggregate log. To clear this event, the host issues a Get Log Page command with the Retain Asynchronous Event bit cleared to '0' for the Endurance Group Event Aggregate log.
07h	Reachability Group Change: The Reachability Group information (refer to section 5.16.1.TBD-1) related to a Reachability Group that contains namespaces attached to this controller has changed (e.g., a member was added to or deleted from a Reachability Group). The current Reachability Group information for attached namespaces is indicated in the Reachability Groups log page. To clear this event, the host issues a Get Log Page command (refer to section 5.16) with the Retain Asynchronous Event bit cleared to '0' for the Reachability Groups log page. A controller shall not send this event if a Namespace Attribute Changed notice is sent for the same event, such as a change due to: <ul style="list-style-type: none"> a) the attachment of a namespace (refer to section 5.22); b) the deletion of a namespace (refer to section 5.23); or c) the detachment of a namespace (refer to section 5.22).
08h	Reachability Association Change: The Reachability Association information (refer to section 5.16.1.TBD-2) related to a Reachability Association that contains namespaces attached to this controller has changed (e.g., a member was added to or deleted from a Reachability Association). The current Reachability Association information for attached namespaces is indicated in the Reachability Associations log page. To clear this event, the host issues a Get Log Page command (refer to section 5.16) with the Retain Asynchronous Event bit cleared to '0' for the Reachability Associations log page. A controller shall not send this event if a Namespace Attribute Changed notice is sent for the same event, such as a change due to: <ul style="list-style-type: none"> a) the attachment of a namespace (refer to section 5.22); b) the deletion of a namespace (refer to section 5.23); or c) the detachment of a namespace (refer to section 5.22).
09h 07h to EEh	Reserved
EFh ²	Zone Descriptor Changed: I/O Command Set specific definition.
F0h	Discovery Log Page Change: A change has occurred to one or more of the Discovery Log Pages. The host should submit a Get Log Page command to receive updated Discovery Log Pages.
F1h to FFh	Reserved for future NVMe over Fabrics Asynchronous Event Notifications

Figure 147: Asynchronous Event Information – Notice

Value	Description
NOTE:	
1. Refer to the NVM Command Set Specification.	
2. Refer to the Zoned Namespace Command Set Specification.	

...

5.16 Get Log Page command

...

5.16.1 Log Specific Information

...

Figure 202: Get Log Page – Log Page Identifiers

Log Identifier	Scope	Log Page Name	Reference Section
00h	Controller	Supported Log Pages	5.16.1.1
01h	Controller	Error Information	5.16.1.2
02h	Controller ¹	SMART / Health Information	5.16.1.3
	Namespace ²		
03h	Domain / NVM subsystem ⁶	Firmware Slot Information	5.16.1.4
04h	Controller	Changed Namespace List	5.16.1.5
05h	Controller	Commands Supported and Effects	5.16.1.6
06h	Controller ³	Device Self-test ⁵	5.16.1.7
	Domain / NVM subsystem ^{4, 6}		
07h	Vendor Specific	Telemetry Host-Initiated ⁵	5.16.1.8
08h	Vendor Specific	Telemetry Controller-Initiated ⁵	5.16.1.9
09h	Domain / NVM subsystem ⁶	Endurance Group Information	5.16.1.10
0Ah	Domain / NVM subsystem ⁶	Predictable Latency Per NVM Set	5.16.1.11
0Bh	Domain / NVM subsystem ⁶	Predictable Latency Event Aggregate	5.16.1.12
0Ch	Controller	Asymmetric Namespace Access	5.16.1.13
0Dh	NVM subsystem	Persistent Event Log ⁵	5.16.1.14
0Eh	Refer to the NVM Command Set		
0Fh	Domain / NVM subsystem ⁶	Endurance Group Event Aggregate	5.16.1.15
10h	Domain / NVM subsystem ^{5, 6}	Media Unit Status	5.16.1.16
11h	Domain / NVM subsystem ⁶	Supported Capacity Configuration List	5.16.1.17
12h	Controller	Feature Identifiers Supported and Effects	5.16.1.18
13h	Controller	NVMe-MI Commands Supported and Effects	5.16.1.19
14h	NVM subsystem	Command and Feature Lockdown ⁵	5.16.1.20
15h	Controller	Boot Partition	5.16.1.21
16h	Endurance Group	Rotational Media Information	5.16.1.22
1Ah	Controller	Reachability Groups	5.16.1.TBD -1
1Bh	Controller	Reachability Associations	5.16.1.TBD -2

Figure 202: Get Log Page – Log Page Identifiers

Log Identifier	Scope	Log Page Name	Reference Section
1Ch to 6Fh	Reserved		
70h		Discovery	5.16.1.23
71h to 7Fh	Reserved		
80h	Controller	Reservation Notification	5.16.1.24
81h	NVM subsystem	Sanitize Status	5.16.1.25
82h to BEh	I/O Command Set Specific		
BFh	Refer to the Zoned Namespace Command Set		
C0h to FFh	Vendor specific ⁵		
<p>Key:</p> <p>Namespace = The log page contains information about a specific namespace. Endurance Group = The log page contains information about a specific Endurance Group. Controller = The log page contains information about the controller that is processing the command. Domain = The log page contains information about the Domain. NVM subsystem = The log page contains information about the NVM subsystem. Vendor Specific = The log page contains information that is vendor specific.</p> <p>Notes:</p> <ol style="list-style-type: none"> For namespace identifiers of 0h or FFFFFFFFh. For namespace identifiers other than 0h or FFFFFFFFh. Bit 0 is cleared to '0' in the DSTO field in the Identify Controller data structure (refer to Figure 275). Bit 0 is set to '1' in the DSTO field in the Identify Controller data structure. Selection of a UUID may be supported. Refer to section 8.25. For NVM subsystems that support multiple domains (refer to the MDS bit in the Identify Controller data structure, Figure 275), Domain scope information is returned. 			

...

5.16.1.TBD-1 Reachability Groups (Log Identifier 1Ah)

This log page consists of a header describing the log page and a list of descriptors containing the Reachability Groups (refer to section 8.TBD) that contain namespaces attached to the controller processing this command that all have the same reachability attributes.

If the Index Offset Supported bit is cleared to '0' in the LID Support and Effects data structure for this log page (refer to Figure 204), then:

- 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 Index Offset Supported bit is set to '1' in the LID Supported and Effects data structure for this log page (refer to Figure 204), then:

- for indexes greater than 0, the entry data structure that is indexed is a Reachability Group descriptor (e.g., specifying an index offset of 2 returns this log page starting at the second descriptor (i.e., Reachability Group Descriptor 1)); and
- for index 0, the data structure starting from the beginning of the log page is returned.

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

The Log Specific Parameter field in Command Dword 10 (refer to Figure 197) for this log page is defined in Figure FIGTBD1.

Figure FIGTBD1: Reachability Groups Log Specific Parameter Field

Bits	Description
14:09	Reserved
08	Return Groups Only (RGO): If set to '1', then the controller shall return Reachability Group Descriptors with the Number of NSID Values field in each Reachability Group Descriptor cleared to 0h (i.e., no Namespace Identifiers are returned). If cleared to '0', then the controller shall return Reachability Group Descriptors that contain the Namespace Identifiers of attached namespaces that are members of the Reachability Group described by that Reachability Group Descriptor and the Number of NSID Values field set to the number of Namespace Identifier values in that Reachability Group Descriptor.

Figure FIGTBD2: Reachability Groups Log Page

Bytes	Description
07:00	Change Count: This field contains a 64-bit incrementing Reachability Groups log page change count, indicating an identifier for this set of reachability groups information. The count starts at 0h following a Controller Level Reset and is incremented each time the contents of the log page change. If the value of this field is FFFFFFFF_FFFFFFFFh, then the field shall be cleared to 0h when incremented (i.e., rolls over to 0h).
09:08	Number of Reachability Group Descriptors (NRGD): This field indicates the number of Reachability Group Descriptors available in the log page. The log page shall contain one Reachability Group Descriptor for each Reachability Group that contains namespaces that are attached to the controller. If, for a Reachability Group, there are no namespaces attached to the controller processing the command, then no Reachability Group Descriptor is returned for that Reachability Group (i.e., a Reachability Group Descriptor is returned only if that Reachability Group contains namespaces that are attached to the controller processing the command). If no namespaces are attached to the controller, then the log page does not contain any Reachability Group Descriptors and this field is cleared to 0h.
15:10	Reserved
n:16	Reachability Group Descriptor 0 , if any (refer to Figure FIGTBD3)
m:n+1	Reachability Group Descriptor 1 , if any (refer to Figure FIGTBD3)
...	...
x:y	Reachability Group Descriptor NRGD-1 , if any (refer to Figure FIGTBD3)

The format of the Reachability Group Descriptor is defined in Figure FIGTBD3. Namespace Identifiers shall be listed in ascending NSID order.

Figure FIGTBD3: Reachability Group Descriptor format

Bytes	Description
03:00	Reachability Group ID (RGID): The Reachability Group Identifier associated with namespaces in the Reachability Group (refer to section 8.TBD) described by this Reachability Group Descriptor.
07:04	Number of NSID Values (NNID): This field indicates the number of Namespace Identifier values in this Reachability Group Descriptor. If the RGO bit is set to '1', then this field shall be cleared to 0h.

Figure FIGTBD3: Reachability Group Descriptor format

Bytes	Description
15:08	<p>Change Count: This field contains a 64-bit incrementing count, indicating an identifier for the information contained in this Reachability Group Descriptor. A value of 0h indicates that the controller does not report a Change Count for this Reachability Group Descriptor. If a Change Count is reported, then the count starts at 1h following a Controller Level Reset and is incremented each time the data in this Reachability Group Descriptor changes. If the value of this field is FFFFFFFF_FFFFFFFFh, then the field shall be set to 1h when incremented (i.e., rolls over to 1h).</p> <p>If this field contains 0h, the host should examine this Reachability Group Descriptor for any changes and not use this field as an indicator that a change has occurred.</p>
31:16	Reserved
35:32	Namespace Identifier 0: The Namespace Identifier of the first attached namespace, if any, that is a member of this Reachability Group.
39:36	Namespace Identifier 1: The Namespace Identifier of the second attached namespace, if any, that is a member of this Reachability Group.
...	...
(((NNID-1)*4) + 35): (((NNID-1)*4) + 32)	Namespace Identifier NNID-1: The Namespace Identifier of the NNID attached namespace, if any, that is a member of this Reachability Group.

5.16.1.TBD-2 Reachability Associations (Log Identifier TBD-2h)

This log page consists of a header describing the log page and a list of descriptors containing the Reachability Associations (refer to section 8.TBD), for the controller processing the Get Log Page command requesting the Reachability Associations log page, that contain Reachability Groups (refer 5.16.1.TBD-1) that all have the same reachability attributes.

If the Index Offset Supported bit is cleared to '0' in the LID Support and Effects data structure for this log page (refer to Figure 204), then:

- if the RAO 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 Index Offset Supported bit is set to '1' in the LID Supported and Effects data structure for this log page (refer to Figure 204), then:

- for indexes greater than 0, the entry data structure that is indexed is a Reachability Association descriptor (e.g., specifying an index offset of 2 returns this log page starting at the second descriptor (i.e., Reachability Association Descriptor 1)).

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

The Log Specific Parameter field in Command Dword 10 (refer to Figure 197) for this log page is defined in Figure FIGTBD4.

Figure FIGTBD4: Reachability Associations Log Specific Parameter Field

Bits	Description
14:09	Reserved

Figure FIGTBD4: Reachability Associations Log Specific Parameter Field

Bits	Description
08	Return Associations Only (RAO): If set to '1', then the controller shall return Reachability Association Descriptors with the Number of RGID Values field in each Reachability Association Descriptor cleared to 0h (i.e., no Reachability Group Identifiers are returned). If cleared to '0', then the controller shall return Reachability Association Descriptors that contain the RGIDs of reachability groups that are members of the Reachability Association described by that Reachability Association Descriptor and the Number of RGID Values field set to the number of Reachability Group Identifier values in that Reachability Association Descriptor.

Figure FIGTBD5: Reachability Associations Log Page

Bytes	Description
07:00	Change Count: This field contains a 64-bit incrementing Reachability Associations log change count, indicating an identifier for this set of reachability associations information. The count starts at 0h following a Controller Level Reset and is incremented each time the contents of the log page change. If the value of this field is FFFFFFFF_FFFFFFFFh, then the field shall be cleared to 0h when incremented (i.e., rolls over to 0h).
09:08	Number of Reachability Association Descriptors (NRAD): This field indicates the number of Reachability Association Descriptors available in the log page. The log page shall contain one Reachability Association Descriptor for each Reachability Association that contains a Reachability Group that contains namespaces that are attached to the controller. If, for a Reachability Association, there are no Reachability Groups that contains namespaces that are attached to the controller processing the command, then no Reachability Association Descriptor shall be returned for that Reachability Association (i.e., a Reachability Association Descriptor is returned only if that Reachability Association contains Reachability Groups associated with the controller processing the command). If no namespaces are attached to the controller, then the log page does not contain any Reachability Association Descriptors and this field shall be cleared to 0h.
15:10	Reserved
n:16	Reachability Association Descriptor 0 , if any (refer to Figure FIGTBD6)
m:n+1	Reachability Association Descriptor 1 , if any (refer to Figure FIGTBD6)
...	...
x:y	Reachability Association Descriptor NRAD-1 , if any (refer to Figure FIGTBD6)

The format of the Reachability Association Descriptor is defined in Figure FIGTBD6. Reachability Group Identifiers shall be listed in ascending order.

Figure FIGTBD6: Reachability Association Descriptor format

Bytes	Description
03:00	Reachability Association ID (RASID): The Reachability Association Identifier associated with Reachability Groups in the Reachability Association (refer to section 8.TBD) described by this Reachability Association Descriptor.
07:04	Number of RGID Values (NRID): This field indicates the number of Reachability Group Identifier values in this Reachability Association Descriptor. If the RAO bit is set to '1', then this field shall be cleared to 0h.

Figure FIGTBD6: Reachability Association Descriptor format

Bytes	Description															
15:08	<p>Change Count: This field contains a 64-bit incrementing count, indicating an identifier for the information contained in this Reachability Association Descriptor. A value of 0h indicates that the controller does not report a Change Count for this Reachability Association Descriptor. If a Change Count is reported, then the count starts at 1h following a Controller Level Reset and is incremented each time the data in this Reachability Association Descriptor changes. If the value of this field is FFFFFFFF_FFFFFFFFh, then the field shall be set to 1h when incremented (i.e., rolls over to 1h).</p> <p>If this field contains 0h, the host should examine this Reachability Group Descriptor for any changes and not use this field as an indicator that a change has occurred.</p>															
16	<p>Reachability Association Characteristics: This field indicates the Access characteristics for all Reachability Groups in this Reachability Association Descriptor for operations between the namespaces represented by the Reachability Groups.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> <th>Reference</th> </tr> </thead> <tbody> <tr> <td>01h</td> <td>Reachable with no performance characteristics specified</td> <td>8.TBD</td> </tr> <tr> <td>02h</td> <td>Fast copy operations are supported (refer to the Fast copy operations section in the NVM Command Set Specification)</td> <td>NVM Command Set</td> </tr> <tr> <td>03h</td> <td>Fast copy operations are not supported (refer to the Fast copy operations section in the NVM Command Set Specification)</td> <td>NVM Command Set</td> </tr> <tr> <td>All other values</td> <td>Reserved</td> <td></td> </tr> </tbody> </table>	Value	Description	Reference	01h	Reachable with no performance characteristics specified	8.TBD	02h	Fast copy operations are supported (refer to the Fast copy operations section in the NVM Command Set Specification)	NVM Command Set	03h	Fast copy operations are not supported (refer to the Fast copy operations section in the NVM Command Set Specification)	NVM Command Set	All other values	Reserved	
Value	Description	Reference														
01h	Reachable with no performance characteristics specified	8.TBD														
02h	Fast copy operations are supported (refer to the Fast copy operations section in the NVM Command Set Specification)	NVM Command Set														
03h	Fast copy operations are not supported (refer to the Fast copy operations section in the NVM Command Set Specification)	NVM Command Set														
All other values	Reserved															
31:17	Reserved															
35:32	Reachability Group Identifier 0: The Reachability Group Identifier of the first Reachability Group, if any, that is a member of this Reachability Association.															
39:36	Reachability Group Identifier 1: The Reachability Group Identifier of the second Reachability Group, if any, that is a member of this Reachability Association.															
...	...															
(((NRID-1)*4)+35): (((NRID-1)*4)+32)	Reachability Group Identifier NRID-1: The Reachability Group Identifier of the last Reachability Group, if any, that is a member of this Reachability Association.															

...

5.17 Identify command

...

5.17.2 Identify Data Structures

5.17.2.1 Identify Controller Data Structure (CNS 01h)

...

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

Bytes	I/O ¹	Admin ¹	Disc ¹	Description
Controller Capabilities and Features				
...

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

Bytes	I/O ¹	Admin ¹	Disc ¹	Description
95:92	M	M	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>Bit 31 is set to '1' if the controller supports sending Discovery Log Page Change Notifications. If cleared to '0', then the controller does not support the Discovery Log Page Change Notification events.</p> <p>Bits 30:28 are reserved.</p> <p>Bit 27 is set to '1' if the controller supports the Zone Descriptor Changed Notices event and the associated Changed Zone List log page (refer to the Zoned Namespace Command Set Specification). If cleared to '0', then the controller does not support the Zone Descriptor Changed Notices event nor the associated Changed Zone List log page.</p> <p>Bits 26:4618 are reserved.</p> <p>Bit 17 (Reachability Groups Change Notices Support) is set to '1' if the controller supports the Reachability Groups Change Notices event, and the Reachability Association Change Notices event. If cleared to '0', then the controller does not support the Reachability Groups Change Notices event, nor the Reachability Association Change Notices event.</p> <p>Bit 15 is set to '1' if the controller supports the Normal NVM Subsystem Shutdown event. If cleared to '0', then the controller does not support the Normal NVM Subsystem Shutdown event.</p> <p>Bit 14 is set to '1' if the controller supports the Endurance Group Event Aggregate Log Page Change Notices event. If cleared to '0', then the controller does not support the Endurance Group Event Aggregate Log Page Change Notices event.</p> <p>Bit 13 is set to '1' if the controller supports the LBA Status Information Alert Notices event (refer to the NVM Command Set Specification). If cleared to '0', then the controller does not support the LBA Status Information Alert Notices event.</p> <p>Bit 12 is set to '1' if the controller supports the Predictable Latency Event Aggregate Log Change Notices event. If cleared to '0', then the controller does not support the Predictable Latency Event Aggregate Log Change Notices event.</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 the Firmware Activation Notices event. 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 the Namespace Attribute Notices event 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>
...

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

Bytes	I/O ¹	Admin ¹	Disc ¹	Description								
134	O	R	R	Controller Reachability Capabilities (CRCAP): This field specifies reachability capabilities of the controller and NVM subsystem.								
				<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>Reachability Group ID Changeable (RGIDC): If set to '1', then the RGRPID field in the I/O Command Set Independent Identify Namespace data structure (refer to Figure 280) does not change while the namespace is attached to any controller. If cleared to '0', then the RGRPID field may change while the namespace is attached to any controller. Refer to section 8.TBD.</td> </tr> <tr> <td>0</td> <td>Reachability Reporting Supported (RRSUP): If set to '1', then the NVM subsystem supports Reachability Reporting (refer to section 8.TBD). If cleared to '0', then the NVM subsystem does not support Reachability Reporting.</td> </tr> </tbody> </table>	Bits	Description	7:2	Reserved	1	Reachability Group ID Changeable (RGIDC): If set to '1', then the RGRPID field in the I/O Command Set Independent Identify Namespace data structure (refer to Figure 280) does not change while the namespace is attached to any controller. If cleared to '0', then the RGRPID field may change while the namespace is attached to any controller. Refer to section 8.TBD .	0	Reachability Reporting Supported (RRSUP): If set to '1', then the NVM subsystem supports Reachability Reporting (refer to section 8.TBD). If cleared to '0', then the NVM subsystem does not support Reachability Reporting.
				Bits	Description							
7:2	Reserved											
1	Reachability Group ID Changeable (RGIDC): If set to '1', then the RGRPID field in the I/O Command Set Independent Identify Namespace data structure (refer to Figure 280) does not change while the namespace is attached to any controller. If cleared to '0', then the RGRPID field may change while the namespace is attached to any controller. Refer to section 8.TBD .											
0	Reachability Reporting Supported (RRSUP): If set to '1', then the NVM subsystem supports Reachability Reporting (refer to section 8.TBD). If cleared to '0', then the NVM subsystem does not support Reachability Reporting.											
239:134 5				Reserved								
252:240				Reserved for the NVMe Management Interface.								
...												
Notes:												
1. O/M/R definition: O = Optional, M = Mandatory, R = Reserved.												
2. Mandatory for I/O controllers using a message-based transport. Reserved for controllers using a memory-based transport.												

...

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

...

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

Bytes	O/M ¹	Description
...
19:18		Reserved
23:20	O	Reachability Group Identifier (RGRPID): For NSID other than FFFFFFFFh, this field indicates the Reachability Group Identifier of the reachability group (refer to section 8.TBD.1) of which the namespace is a member. Each namespace that is attached to a controller that supports Reachability Group Reporting (refer to the RRSUP bit in the CRCAP field in Figure 275)) shall report a valid RGRPID. If the controller does not support Reachability Group Reporting, then this field shall be cleared to 0h. If the value in this field changes and Reachability Group Change Notices are supported and enabled, then the controller shall issue a Reachability Group Change Notice.
4095:1 5 24		Reserved
Notes:		
1. O/M definition: O = Optional, M = Mandatory.		

...

5.27 Set Features command

...

5.27.1 Feature Specific Information

...

5.27.1.8 Asynchronous Event Configuration (Feature Identifier 0Bh)

...

Figure 326: Asynchronous Event Configuration – Command Dword 11

Bits	Description
31	Discovery Log Page Change Notification: This bit indicates that the Discovery controller reports Discovery Log Page Change Notifications. If set to '1', the Discovery controller shall send a notification if Discovery Log Page changes occur.
30:28	Reserved
27	Zone Descriptor Changed Notices²: I/O Command Set specific definition.
26:16 9	Reserved
18	Reachability Group: This bit determines whether an asynchronous event notification is sent to the host when a Reachability Group change occurs (i.e., the contents of the Reachability Groups log page (refer to section 5.16.TBD-1) changed). If this bit is set to '1', then the Reachability Group 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 Reachability Group Change Notices event to the host.
17	Reachability Association: This bit determines whether an asynchronous event notification is sent to the host when a Reachability Association change occurs (i.e., the contents of the Reachability Associations log page (refer to section 5.16.TBD-1) changed). If this bit is set to '1', then the Reachability Association 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 Reachability Association Change Notices event to the host.
15	Normal NVM Subsystem Shutdown: This bit determines whether an asynchronous event notification is sent to the host when the NVM subsystem has started performing a normal shutdown due to an NVM Subsystem Shutdown (refer to Figure 149). If this bit is set to '1', then the Normal NVM Subsystem Shutdown event is sent to the host if an outstanding Asynchronous Event Request command exists at the time this condition occurs. If this bit is cleared to '0', then the controller shall not send the Normal NVM Subsystem Shutdown event to the host.
14	Endurance Group Event Aggregate Log Change Notices: This bit determines whether an asynchronous event notification is sent to the host when an event entry for an Endurance Group (refer to section 3.2.3) has been added to the Endurance Group Event Aggregate log (refer to section 5.16.1.15). If this bit is set to '1', then the Endurance Group Event Aggregate Log Change event is sent to the host when this condition occurs. If this bit is cleared to '0', then the controller shall not send the Endurance Group Event Aggregate Log Change event to the host. If Endurance Groups are not supported and this bit is set to '1', then the Set Features command shall be aborted with a status of Invalid Field in Command.
13	LBA Status Information Alert Notices¹: I/O Command Set specific definition.
12	Predictable Latency Event Aggregate Log Change Notices: This bit determines whether an asynchronous event notification is sent to the host when an event pending entry for an NVM Set (refer to section 5.16.1.12) has been added to the Predictable Latency Event Aggregate Log. If this bit is set to '1', then the Predictable Latency Event Aggregate Log Change event is sent to the host when this condition occurs. If this bit is cleared to '0', then the controller shall not send the Predictable Latency Event Aggregate Log Change event to the host.
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 5.16.1.13) 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 bit determines whether an asynchronous event notification is sent to the host when the Telemetry Controller-Initiated Data Available field transitions from 0h to 1h 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.

Figure 326: Asynchronous Event Configuration – Command Dword 11

Bits	Description
09	Firmware Activation Notices: This bit determines whether an asynchronous event notification is sent to the host for a Firmware Activation Starting event (refer to Figure 147). If this bit is set to '1', then the Firmware Activation Starting event is sent to the host when this condition occurs. If this bit is cleared to '0', then the controller shall not send the Firmware Activation Starting event to the host.
08	Namespace Attribute Notices: This bit determines whether an asynchronous event notification is sent to the host for a Namespace Attribute change (refer to Figure 147). If this bit is set to '1', then the Namespace Attribute 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 Namespace Attribute Changed event to the host.
07:00	SMART / Health Critical Warnings: This field determines whether an asynchronous event notification is sent to the host for the corresponding Critical Warning specified in the SMART / Health Information log (refer to Figure 207). If a bit is set to '1', then an asynchronous event notification is sent when the corresponding critical warning bit is set to '1' in the SMART / Health Information log. If a bit is cleared to '0', then an asynchronous event notification is not sent when the corresponding critical warning bit is set to '1' in the SMART / Health Information log.
NOTE: 1. Refer to the NVM Command Set Specification. 2. Refer to the Zoned Namespace Command Set Specification.	

...

8.TBD Reachability Reporting architecture

8.TBD.1 Reachability Reporting overview

For some operations that specify multiple namespaces it is necessary to indicate what namespaces are able to be used in those operations (e.g., Copy command that specifies multiple namespaces, Execute command in the Computational Programs command set). This ability is called reachability. Reachability is defined at the controller level (i.e., only namespaces attached to the same controller are reachable).

A Reachability Group becomes available to a controller if:

- a namespace becomes attached to that controller and that namespace is associated with a Reachability Group that did not currently have any namespaces attached to that controller; or
- a configuration change in an NVM subsystem changes the membership in a Reachability Group (e.g., a new Reachability Group is present in the Reachability Groups log page (refer to section 5.16.1.TBD-1)).

A Reachability Group becomes unavailable to a controller if a namespace is detached that is the only namespace attached to that controller for that Reachability Group (i.e., a Reachability Group that was present in the Reachability Groups log page is no longer present).

A Reachability Association change occurs if a Reachability Group that is associated with that Reachability Association becomes available or unavailable to that controller. (e.g., a new Reachability Association is present in the Reachability Associations log page (refer to section 5.16.1.TBD-2) or a Reachability Association that was present in the Reachability Associations log page is no longer present).

If an NVM subsystem supports Reachability Reporting, then all controllers in that NVM subsystem shall:

- set the RRSUP bit to '1' in the Controller Reachability Capabilities (CRCAP) field in the Identify Controller data structure (refer to Figure 275) to indicate support for Reachability Reporting;
- support Reachability Groups Change Notices (refer to section 5.27.1.8);
- support Reachability Association Change Notices (refer to section 5.27.1.8);
- support the Reachability Groups log page; and
- support the Reachability Associations log page.

Namespaces attached to a controller in a NVM subsystem that supports Reachability Reporting shall:

- be members of a Reachability Group; and
- supply a valid Reachability Group Identifier in the Reachability Group Identifier (RGRPID) field in the I/O Command Set Independent Identify Namespace Data Structure (refer to [Figure 280](#)).

Reachability and characteristics of reachability are indicated in the combination of the Reachability Groups log page (refer to section [5.16.1.TBD-1](#)) and the Reachability Associations log page (refer to section [5.16.1.TBD-2](#)). A Reachability Group defines a group of namespaces that are all associated with namespaces in other Reachability Groups defined by a Reachability Association. The method of assigning Reachability Group identifiers and Reachability Association identifiers is outside the scope of this specification. If members of a Reachability Group are able to reach members of that Reachability Group, then a Reachability Association that contains only that Reachability Group specifies the characteristics of that reachability. All Reachability Groups in a Reachability Association have identical access characteristics. All namespaces in a Reachability Group have identical access characteristics. If a Reachability Group is not in any Reachability Association then the members of that Reachability Group are not reachable by any other namespaces, including the members of that Reachability Group. A namespace is always able to reach itself (i.e., a command that has two or more NSIDs as part of the command where those NSIDs are all the same is not constrained by reachability).

If the Reachability Association Characteristics field (refer to [Figure FIGTBD6](#)) is set to 01h, then all namespaces in the associated Reachability Groups are able to reach other namespaces in the other Reachability Groups in that Reachability Association without any indication of a performance characteristic (e.g., for the Execute command in the Computational Programs Command Set). If the Reachability Association Characteristics field is set to 02h, then all namespaces in the associated Reachability Groups are reachable and support fast copy operations as defined by the applicable I/O Command Set specification (e.g., the Fast copy operations section in the NVM Command Set Specification). If the Reachability Association Characteristics field is set to 03h, then all namespaces in the associated Reachability Groups are reachable but do not support fast copy operations.

A namespace is only allowed to be in one Reachability Group. A Reachability Group is in zero or more Reachability Associations with different values of the Reachability Association Characteristics field for each Reachability Association.

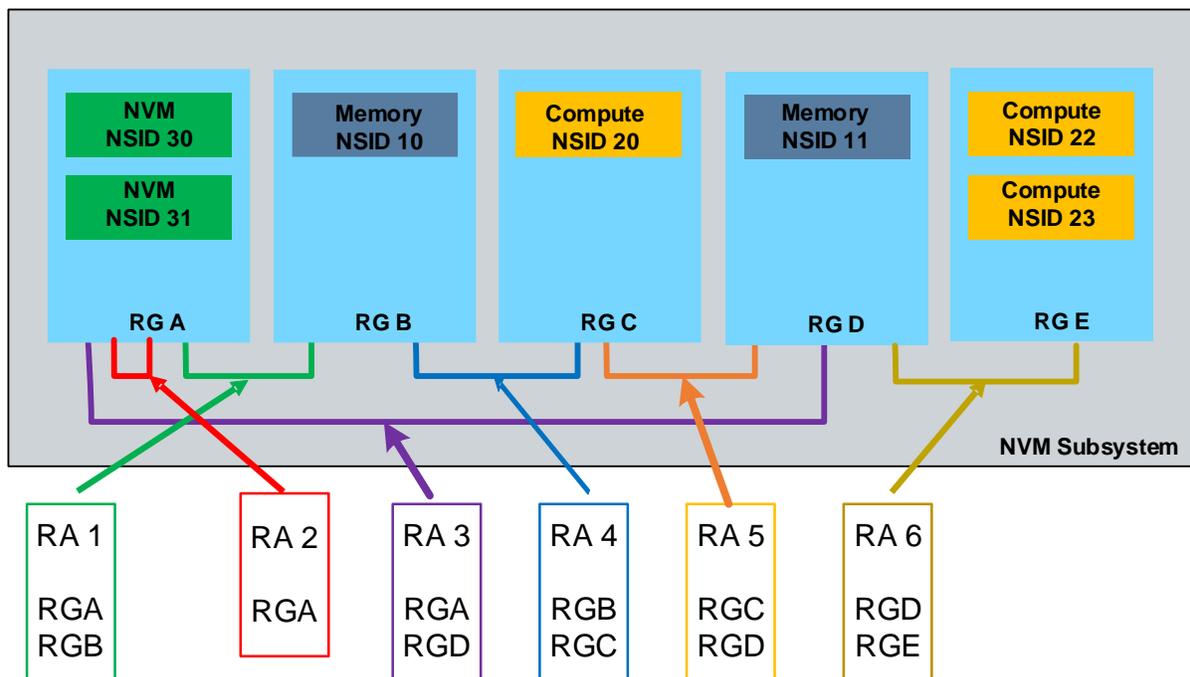
The Reachability Group Identifier (RGRPID) for each Reachability Group shall be unique within the NVM subsystem. If RGIDC bit in the CRCAP field in the Identify Controller data structure is set to '1', then the Reachability Group Identifier shall not change while the namespace is attached to any controller in the NVM subsystem. If RGIDC bit in the CRCAP field is cleared to '0', then the Reachability Group Identifier may change while the namespace is attached to any controller in the NVM subsystem. If the Reachability Group Identifier changes, the controller shall issue the Reachability Groups Change Notice as described in this section.

[Figure FIGTBD7](#) is an example configuration that shows that:

- Reachability Association 1 (i.e., RA 1) :
 - specifies that:
 - NSID 10 and NSID 30 are able to reach each other; and
 - NSID 10 and NSID 31 are able to reach each other;
 - and
 - does not specify any reachability between NSID 30 and NSID 31;
- Reachability Association 2 (i.e., RA 2) is a self-reference for Reachability Group A (i.e., RG A) that specifies:
 - that NSID 30 and NSID 31 are able to reach each other;
- Reachability Association 3 (i.e., RA 3):

- specifies that:
 - NSID 11 and NSID 30 are able to reach each other; and
 - NSID 11 and NSID 31 are able to reach each other;
- and
- does not specify any reachability between NSID 30 and NSID 31;
- The lack of a Reachability Association that contains Reachability Group B (i.e., RG B) and Reachability Group D (i.e., RG D) specifies:
 - that NSID 10 and NSID 11 are not able to reach each other;
- and
- The lack of a self-referencing Reachability Association for Reachability Group E (i.e., RG E) specifies:
 - that NSID 22 and NSID 23 are not able to reach each other.

Figure FIG TBD7: Reachability Associations Example



8.TBD.2 Reachability event notifications

Reachability may generate a Reachability Group Change notice (refer to section 5.2.1) or a Reachability Association Change notice (refer to section 5.2.1).

Receipt of a Reachability Group Change Notice from a controller may indicate:

- a) an NSID has been added to one or more of the Reachability Group Descriptors;
- b) an NSID has been removed from one or more of the Reachability Group Descriptors;
- c) a Reachability Group no longer has any NSIDs attached to this controller as members of that Reachability Group and therefore is no longer reported in the Reachability Groups log page for this controller; or
- d) the NSID of a namespace has moved from one Reachability Group Descriptor to a different Reachability Group Descriptor (i.e., the RGRPID field in the I/O Command Set Independent Identify Namespace Data Structure for that namespace has changed), if RGIDC bit in the CRCAP field is cleared to '0' in the Identify Controller data structure (refer to Figure 275).

As a result of receiving a Reachability Group Change notice, the host should read the Reachability Groups log page (refer to section 5.16.1.TBD-1) to check for each of those possible changes.

Receipt of a Reachability Association Change Notice from a controller may indicate:

- a) an RGID has been added to one or more of the Reachability Association Descriptors;
- b) an RGID has been removed from one or more of the Reachability Association Descriptors; and/or
- c) a Reachability Association no longer has any RGIDs associated with this controller as members of that Reachability Association and therefore is no longer reported in the Reachability Associations log page for this controller.

As a result of receiving a Reachability Association Change notice, the host should read the Reachability Associations log page (refer to section 5.16.1.TBD-2) to check for each of those possible changes.

...

Modify portions of NVMe Management Interface Specification 1.2.c as follows:

Modify section 6 as follows:

6 NVM Express Admin Command Set

...

6.3 Get Log Page

...

Figure 123: Management Endpoint - Log Page Support

Log Page Name ³	Log Identifier	Support Requirements ¹	
		NVMe Storage Device	NVMe Enclosure
...			
Rotational Media Information	16h	O	O
Reachability Groups	1Ah	O	O
Reachability Associations	1Bh	O	O
Discovery	70h	O	O
...			
Notes:			
1. O = Optional, M = Mandatory, P = Prohibited.			
2. Optional for versions 1.1 and earlier of this specification.			
3. Refer to the NVM Express Base Specification unless another footnote specifies otherwise.			
4. Refer to the NVM Command Set Specification.			
5. Refer to the Zoned Namespace Command Set Specification.			

...