



LEGAL NOTICE:

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

This erratum to the NVM Express revision 1.4 specification 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 erratum to the NVM Express revision 1.4 specification subject, however, to the Member's continued compliance with the Company's Intellectual Property Policy and Bylaws and the Member's Participation Agreement.

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

LEGAL DISCLAIMER:

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

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

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

NVM Express™ Technical Proposal for New Feature

Technical Proposal ID	TP 8014
Change Date	07/12/2021
Builds on Specification	NVMe Base Specification 2.0
References Ratified TPs	

Technical Proposal Author(s)

Name	Company
Fred Knight	NetApp
John Meneghini	NetApp

Technical Proposal Overview

Discovery controllers are only able to describe things they point to (referral entries for other Discovery subsystems, or to NVM subsystems (with namespaces)). There is no way for a discovery controller to describe itself.

As such, a discovery subsystem with multiple ports cannot describe all the ports/paths that can be used to connect to that discovery subsystem.

This TP makes it possible for a discovery subsystem to describe itself (and all the available ports that can be used to connect to that subsystem). In addition, this TP adds the ability to detect which ports report duplicate information.

Revision History

Revision Date	Change Description
2/5/2021	Initial creation.
3/16/2021	Update “duplicate” bit descriptive text.
3/29/2021	Update to use NVMe 2.0 as the base, remove Entry offset description, name DUPRETINFO bit and add host usage “may” text.
4/8/2021	Use discovery service rather than discovery subsystem – a future ECN will be needed to resolve the uses of “discovery service” and “discovery subsystem” – which appear to be equivalent terms.
4/14/2021	Update descriptions for SUBTYPE fields to add clarity that > 1 entry may exist if multiple ports/protocols are being described for each subsystem that is being described.
4/27/2021	Add the EPCSD bit (as per FMDS group).
5/6/2021	Copy for 30 day review.
5/11/2021	Incorporate changes from 30 day review (change EFLAGS from bit descriptions to a table format); fix case of Discovery Log Page and Discovery Service; make sure the black/red text is correct.
07/12/2021	Integrated into the NVMe Base Specification, revision 2.0.

Editor’s Note:

BLACK text indicates unchanged text; **RED** text indicates new or deleted text; **BLUE** text indicates copied or moved text; **GREEN** text indicates editor notes.

Incompatible Changes:

None

Description of Specification Changes

5.16.1.23 Discovery Log Page (Log Identifier 70h)

The Discovery Log Page shall only be supported by Discovery controllers. The Discovery Log Page shall not be supported by controllers that expose namespaces for NVMe over PCIe or NVMe over Fabrics. The Discovery Log Page provides an inventory of NVM subsystems with which a host may attempt to form an association. The Discovery Log may be specific to the host requesting the log. The Discovery Log page is persistent across power cycles.

The Log Page Offset may be used to retrieve specific records. The number of records is returned in the header of the log page. The format for a Discovery Log Page Entry is defined in Figure 264. The format for the Discovery Log is defined in Figure 265.

A single Get Log Page command used to read the Discovery Log Page shall be atomic. If the host reads the Discovery Log Page using multiple Get Log Page commands the host should ensure that there has not been a change in the contents of the data. The host should read the Discovery Log Page contents in order (i.e., with increasing Log Page Offset values) and then re-read the Generation Counter after the entire log page is transferred. If the Generation Counter does not match the original value read, the host should discard the log page read as the entries may be inconsistent. If the log page contents change during this command sequence, the controller may return a status of Discover Restart.

Every record indicates via the SUBTYPE field if that record is **describing the current Discovery Service**, referring to another Discovery Service, or if the record indicates an NVM subsystem composed of controllers that may expose namespaces. A referral to another Discovery Service (i.e., SUBTYPE field set to 01h) is a mechanism to find additional Discovery subsystems. An NVM subsystem entry (i.e., SUBTYPE field set to 02h) is a mechanism to find NVM subsystems that contain controllers that may expose namespaces. **An entry that describes the current Discovery Service (i.e., SUBTYPE field set to 03h) is a mechanism to find additional access information (e.g., other NVM subsystem ports) for the current Discovery service.** Referrals to another Discovery Service shall not be deeper than eight levels.

For entries that describe the current Discovery Service, the DUPRETINFO bit set to '1' in the Entry Flags field indicates those ports that return duplicate information. A host need not retrieve the Discovery Log Page from all NVM subsystem ports described by a Discovery Log Page Entry that contains the DUPRETINFO bit set to '1' (refer to Figure 263), as each of those ports return the same information. A host may retrieve the Discovery Log Page from multiple ports with the DUPRETINFO bit set to '1' to act as redundant access to the same information. A host may choose to enable the Discovery Log Page Change AEN on one of the connections associated with an entry with the DUPRETINFO bit set to '1'. If the Discovery Log Page Change AEN is enabled on multiple ports represented by entries with the DUPRETINFO bit set to '1', then the host may process a Discovery Log Page Change AEN for one connection as applying to all ports represented by entries that also have the DUPRETINFO bit set to '1'.

If an NVM subsystem supports the dynamic controller model, then all entries for that NVM subsystem shall have the Controller ID field set to FFFFh. For a particular NVM subsystem port and NVMe Transport address in an NVM subsystem, there shall be no more than one entry with the Controller ID field set to:

- FFFFh if that NVM subsystem supports the dynamic controller model; or
- FFFEh if that NVM subsystem supports the static controller model.

Figure 263: Get Log Page – Discovery Log Page Entry

Bytes	Description												
...													
02	<p>Subsystem Type (SUBTYPE): Specifies the type of the NVM subsystem that is indicated in this entry.</p> <table> <tr> <th>Value</th><th>Definition</th></tr> <tr> <td>00</td><td>Reserved.</td></tr> <tr> <td>01</td><td>Referral: The entry describes a referral to another Discovery Service composed of Discovery controllers for that provide additional discovery records. Multiple Referral entries may be reported for each Discovery Service (e.g., if that Discovery Service has multiple NVM subsystem ports or supports multiple protocols).</td></tr> <tr> <td>02</td><td>NVM Subsystem: The entry describes an NVM subsystem that is not associated with Discovery controllers and whose controllers may have attached namespaces (i.e., an NVM subsystem that is not composed of Discovery controllers). Multiple NVM Subsystem entries may be reported for each NVM subsystem if that NVM subsystem has multiple NVM subsystem ports.</td></tr> <tr> <td>03</td><td>Current Discovery Subsystem: The entry describes this Discovery subsystem (i.e., the Discovery Service that contains the controller processing the Get Log Page command). Multiple Current Discovery Subsystem entries may be reported for this Discovery subsystem if the current Discovery subsystem has multiple NVM subsystem ports.</td></tr> <tr> <td>03 04 to 255</td><td>Reserved</td></tr> </table>	Value	Definition	00	Reserved.	01	Referral: The entry describes a referral to another Discovery Service composed of Discovery controllers for that provide additional discovery records. Multiple Referral entries may be reported for each Discovery Service (e.g., if that Discovery Service has multiple NVM subsystem ports or supports multiple protocols).	02	NVM Subsystem: The entry describes an NVM subsystem that is not associated with Discovery controllers and whose controllers may have attached namespaces (i.e., an NVM subsystem that is not composed of Discovery controllers). Multiple NVM Subsystem entries may be reported for each NVM subsystem if that NVM subsystem has multiple NVM subsystem ports.	03	Current Discovery Subsystem: The entry describes this Discovery subsystem (i.e., the Discovery Service that contains the controller processing the Get Log Page command). Multiple Current Discovery Subsystem entries may be reported for this Discovery subsystem if the current Discovery subsystem has multiple NVM subsystem ports.	03 04 to 255	Reserved
Value	Definition												
00	Reserved.												
01	Referral: The entry describes a referral to another Discovery Service composed of Discovery controllers for that provide additional discovery records. Multiple Referral entries may be reported for each Discovery Service (e.g., if that Discovery Service has multiple NVM subsystem ports or supports multiple protocols).												
02	NVM Subsystem: The entry describes an NVM subsystem that is not associated with Discovery controllers and whose controllers may have attached namespaces (i.e., an NVM subsystem that is not composed of Discovery controllers). Multiple NVM Subsystem entries may be reported for each NVM subsystem if that NVM subsystem has multiple NVM subsystem ports.												
03	Current Discovery Subsystem: The entry describes this Discovery subsystem (i.e., the Discovery Service that contains the controller processing the Get Log Page command). Multiple Current Discovery Subsystem entries may be reported for this Discovery subsystem if the current Discovery subsystem has multiple NVM subsystem ports.												
03 04 to 255	Reserved												
...													
09:08	<p>Admin Max SQ Size (ASQSZ): Specifies the maximum size of an Admin Submission Queue. This applies to all controllers in the NVM subsystem. The value shall be a minimum of 32 entries.</p>												
11:10	<p>Entry Flags (EFLAGS): This field Indicates additional information related to the current entry.</p> <table> <tr> <th>Bit</th><th>Definition</th></tr> <tr> <td>15:2</td><td>Reserved</td></tr> <tr> <td>1</td><td> <p>Explicit Persistent Connection Support for Discovery (EPCSD): If set to '1', then Explicit Persistent Connections are supported for the Discovery controller described by this entry. If cleared to '0', then support for Explicit Persistent Connections is not reported.</p> <p>For entries with the SUBTYPE field set to 2h, this bit shall be cleared to '0'.</p> </td></tr> <tr> <td>0</td><td> <p>Duplicate Returned Information (DUPRETINFO): If set to '1', then using the content of this entry to access this Discovery Service returns the same information that is returned by using the content of other entries in this log page that also have this bit set to '1'. If cleared to '0', then using the content of this entry to access this Discovery Service may or may not return different information than is returned by using the content of any other entry in this log page.</p> <p>For entries with the SUBTYPE field set to a value other than 3h, this bit shall be cleared to '0'.</p> </td></tr> </table>	Bit	Definition	15:2	Reserved	1	<p>Explicit Persistent Connection Support for Discovery (EPCSD): If set to '1', then Explicit Persistent Connections are supported for the Discovery controller described by this entry. If cleared to '0', then support for Explicit Persistent Connections is not reported.</p> <p>For entries with the SUBTYPE field set to 2h, this bit shall be cleared to '0'.</p>	0	<p>Duplicate Returned Information (DUPRETINFO): If set to '1', then using the content of this entry to access this Discovery Service returns the same information that is returned by using the content of other entries in this log page that also have this bit set to '1'. If cleared to '0', then using the content of this entry to access this Discovery Service may or may not return different information than is returned by using the content of any other entry in this log page.</p> <p>For entries with the SUBTYPE field set to a value other than 3h, this bit shall be cleared to '0'.</p>				
Bit	Definition												
15:2	Reserved												
1	<p>Explicit Persistent Connection Support for Discovery (EPCSD): If set to '1', then Explicit Persistent Connections are supported for the Discovery controller described by this entry. If cleared to '0', then support for Explicit Persistent Connections is not reported.</p> <p>For entries with the SUBTYPE field set to 2h, this bit shall be cleared to '0'.</p>												
0	<p>Duplicate Returned Information (DUPRETINFO): If set to '1', then using the content of this entry to access this Discovery Service returns the same information that is returned by using the content of other entries in this log page that also have this bit set to '1'. If cleared to '0', then using the content of this entry to access this Discovery Service may or may not return different information than is returned by using the content of any other entry in this log page.</p> <p>For entries with the SUBTYPE field set to a value other than 3h, this bit shall be cleared to '0'.</p>												
31: 40 12	Reserved												
63:32	<p>Transport Service Identifier (TRSVCID): Specifies the NVMe Transport service identifier as an ASCII string. The NVMe Transport service identifier is specified by the associated NVMe Transport binding specification.</p>												
255:64	Reserved												
...													