



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
c/o VTM, Inc.
3855 SW 153rd Drive
Beaverton, OR 97003
USA
info@nvmexpress.org

Technical input submitted to NVM Express® is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2023 NVMe Express, Inc.

NVM Express® Technical Proposal (TP)

Technical Proposal ID	TP8017 Discovery Subsystem Authentication Recommendations
Revision Date	2023.01.11
Builds on Specification(s)	TP8013 Unique Discovery Controller ID
References	NVM Express Base Specification 2.0c

Technical Proposal Author(s)

Name	Company
Erik Smith	Dell
Claudio DeSanti	Dell
Curtis Ballard	HPE
Matthew Goepfert	HPE

Technical Proposal Overview

This proposal updates the wording of section 8.13.2 “NVMe In-band Authentication” to clarify the behavior of a host when connecting to a Discovery subsystem that requires authentication.

The proposal also updates all instances of “unique NQN” to “unique Discovery Service NQN” to make it more consistent with the term “well-known Discovery Service NQN”

Revision History

Revision Date	Change Description
2022.02.15	Initial draft
2022.03.28	Updated concept based on Feedback from FMDS meeting on 3/22.
2022.03.29	Updated based on feedback from co-authors and during the FMDS meeting
2022.04.07	Updated based on feedback from co-authors.
2022.04.28	Updated sequence diagram and clarified the meaning of unidirectional authentication.
2022.05.11	Updated sequence diagram based on comments from FMDS meeting on 2022.05.10
2022.05.19	Updated sequence diagram based on comments from TWG meeting on 2022.05.19
2022.08.23	Updated sequence diagram based on comments during FMDS meeting on 2022.08.16 and 2022.08.23
2022.09.22	Updated revision date and adjusted sequence diagram based on feedback during the TWG meeting on 2022.09.22
2022.09.26	Removed residual comment from the review process.
2022.12.07	Swapped the “Builds on Specification(s)” and “References” fields. Added two new color codes to highlight text from TP8013
2022.12.13	Incorporated editorial changes suggested during FMDS meeting.
2023.01.04	Integrated
2023.01.10	Updated based on feedback from Erik Smith and Mike Allison
2023.01.11	Fixed the figure numbers to match NVM Base Specification 2.0c.

Description for Changes Document for NVMe Express Base Specification 2.0b

New Features/Feature Enhancements/Required Changes:

- Updated wording of section 3.1.2.3 Discovery Controller to convert instances of “unique NQN” to “unique Discovery Service NQN”
- Updated wording of section 8.13.2 NVMe In-band Authentication to provide host guidance when connecting to a Discovery subsystem that requires authentication
- **New requirement / incompatible change:**
 - None
- References:
 - TP8017

Markup Conventions:

Black:	Unchanged (however, hot links are removed)
Orange:	Text from TP8013
Orange Strikethrough:	Text added in TP8013 and deleted by this TP
Red Strikethrough:	Deleted
Blue:	New
Blue Highlighted:	TBD values, anchors, and links to be inserted in new text.
<Green Bracketed>:	Notes to editor

Description of Specification Changes for NVM Express Base Specification 2.0c

3 NVM Express Architecture

...

3.1.2.3 Discovery Controller

A Discovery controller only implements features related to Discovery Log Pages and does not implement I/O Queues, I/O commands, or expose namespaces. The functionality supported by the Discovery controller is defined in section 3.1.2.3.4.

If the Discovery subsystem provides a unique **Discovery Service NQN** (i.e., the NVM Subsystem NVMe Qualified Name (SUBNQN) field in that Discovery subsystem's Identify Controller data structure contains a unique **Discovery Service NQN** value), then that Discovery subsystem shall support both the unique **Discovery Service NQN** and the well-known Discovery Service NQN (i.e., nqn.2014-08.org.nvmexpress.discovery) being specified in the Connect command (refer to section 6.3) from the host.

If the Discovery subsystem does not provide a unique **Discovery Service NQN** (i.e., the SUBNQN field in that Discovery subsystem's Identify Controller data structure contains the well-known Discovery Service NQN), then that Discovery subsystem shall support the well-known Discovery Service NQN being specified in the Connect command from the host.

In the Connect command to a Discovery subsystem that provides a unique **Discovery Service NQN**, the host may use either of the following:

- the well-known Discovery Service NQN (nqn.2014-08.org.nvmexpress.discovery); or
- the unique **Discovery Service NQN** of that Discovery subsystem.

In the Connect command to a Discovery subsystem that does not provide a unique **Discovery Service NQN**, the host uses the well-known Discovery Service NQN.

The method that a host uses to obtain the NVMe Transport information necessary to connect to the well-known Discovery Service is implementation specific.

The Discovery Log Page provided by a Discovery controller contains one or more entries. Each entry specifies information necessary for the host to connect to an NVM subsystem. An entry may be associated with an NVM subsystem that exposes namespaces or a referral to another Discovery Service. There are no ordering requirements for log page entries within the Discovery Log Page.

...

5 Admin Command Set

...

5.16.1.23 Discovery Log Page (Log Identifier 70h)

...

Figure 264: Discovery Log Page Entry Data Structure

Bytes	Description
...	...
255:64	Reserved
511:256	NVM Subsystem Qualified Name (SUBNQN): NVMe Qualified Name (NQN) that uniquely identifies the NVM subsystem. Refer to section 4.4. For a Discovery subsystem, if that Discovery subsystem has a unique Discovery Service NQN (i.e., the NVM Subsystem NVMe Qualified Name (SUBNQN) field in that Discovery subsystem's Identify Controller data structure contains a unique Discovery Service NQN value), then the value returned shall be that unique Discovery Service NQN . If the Discovery subsystem does not have a unique Discovery Service NQN , then the value returned shall be the well-known Discovery Service NQN (i.e., nqn.2014-08.org.nvmexpress.discovery).
...	...

...

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
...
767:564				Reserved
1023:768	M	M	R	NVM Subsystem NVMe Qualified Name (SUBNQN): This field specifies the NVM Subsystem NVMe Qualified Name as a UTF-8 null-terminated string. Refer to section 4.5 for the definition of NVMe Qualified Name. Support for this field is mandatory if the controller supports revision 1.2.1 or later as indicated in the Version property (refer to section 3.1.3.2). For a Discovery controller, if the Discovery subsystem containing the Discovery controller has a unique Discovery Service NQN, then this field shall be set to that unique Discovery Service NQN. If the Discovery subsystem containing the Discovery controller does not have a unique Discovery Service NQN, then this field shall be set to the well-known Discovery Service NQN (i.e., nqn.2014-08.org.nvmexpress.discovery).
1791:1024				Reserved
...
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 I/O controllers using a memory-based transport.				

8 Extended Capabilities

...

8.13.2 NVMe In-band Authentication

The Authentication and Security Requirements (AUTHREQ) field in the Connect response capsule (refer to Figure 382) indicates whether NVMe in-band authentication is required.

If one or more of the bits in the AUTHREQ field are set to '1', then the controller requires that the host authenticate on that queue in order to proceed with Fabrics, Admin, and I/O commands. Authentication success is defined by the specific security protocol that is used for authentication. If any command other than Connect, Authentication Send, or Authentication Receive is received prior to authentication success, then the controller shall abort the command with Authentication Required status.

If all bits in the AUTHREQ field are cleared to '0', then the controller does not require the host to authenticate, and the NVM subsystem shall not abort any command with a status value of Authentication Required.

~~If a Discovery subsystem provides a unique NQN (refer to section 3.1.2.3), then NVMe in-band authentication should be performed using that unique NQN (i.e., the well-known Discovery Service NQN should not be used). The host may connect to that Discovery subsystem using the well-known Discovery Service NQN (i.e., nqn.2014-08.org.nvmexpress.discovery) to determine if that Discovery subsystem provides a unique NQN. The connection using the well-known Discovery Service NQN may or may not require authentication. If the Discovery subsystem does provide a unique NQN, then the host may disconnect the connection to the well-known Discovery Service NQN and reconnect to that Discovery subsystem using the unique NQN. The connection using the unique NQN may or may not require authentication.~~

New Header 8.13.2.<new> ...Special considerations for Discovery subsystems

Hosts that have been configured to authenticate Discovery subsystems should behave as follows:

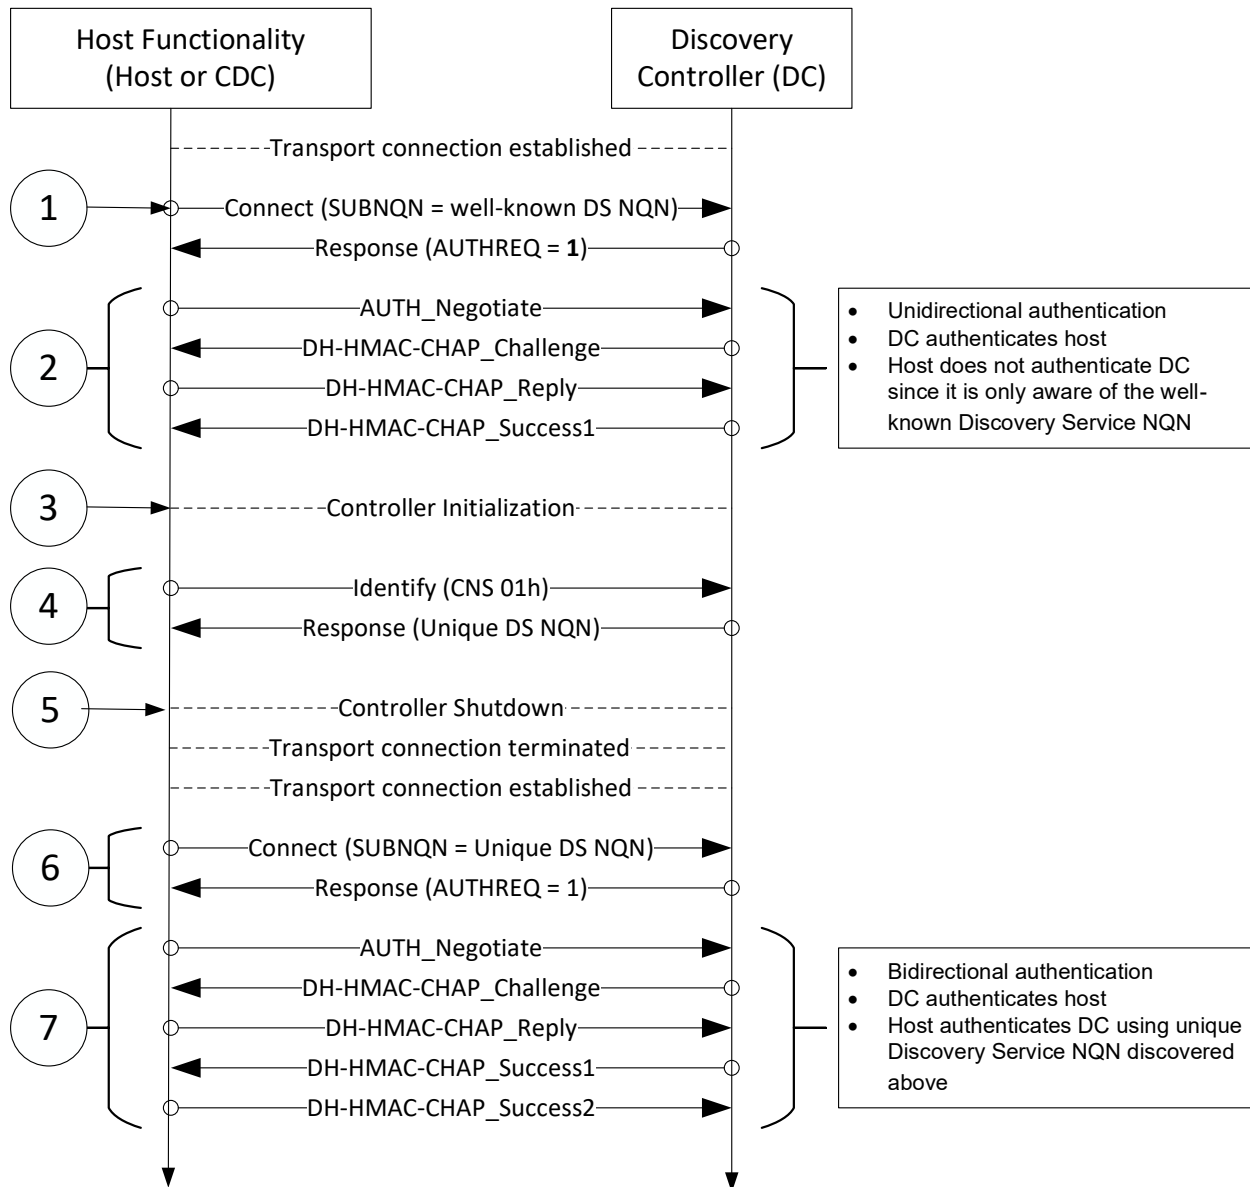
- If the host connected to a Discovery subsystem using the well-known Discovery Service NQN (i.e., nqn.2014-08.org.nvmexpress.discovery) and the Discovery subsystem did not request authentication, then the host should not perform an authentication transaction;
- If the host connected to a Discovery subsystem using the well-known Discovery Service NQN (i.e., nqn.2014-08.org.nvmexpress.discovery) and the Discovery subsystem requested authentication, then the host should perform only unidirectional authentication (i.e., the Discovery subsystem may authenticate the host, but the host should not authenticate the well-known Discovery Service NQN); or
- If the host connected to a Discovery subsystem using the unique Discovery Service NQN for that Discovery subsystem (refer to section 3.1.2.3), regardless of whether the Discovery subsystem requested authentication, then the host may perform unidirectional authentication or bidirectional authentication (i.e., the host may authenticate the unique Discovery Service NQN for that Discovery subsystem).

Figure NEW.1 illustrates the process a host may use to retrieve the unique Discovery Service NQN for a Discovery subsystem when that host has been configured to authenticate Discovery subsystems and the Discovery subsystem that host connects to also requires authentication (i.e., AUTHREQ is non-zero). This process includes:

1. connect to the Discovery subsystem using the well-known Discovery Service NQN (i.e., nqn.2014-08.org.nvmexpress.discovery);
2. perform unidirectional authentication with the Discovery subsystem;
3. perform controller initialization (refer to section 3.5);
4. retrieve the unique Discovery Service NQN (refer to section 3.1.2.3);
5. perform controller shutdown (refer to section 3.6);

6. reconnect to the Discovery subsystem using the unique Discovery Service NQN for that Discovery subsystem; and
7. perform bidirectional authentication with the Discovery subsystem using the unique Discovery Service NQN for that Discovery subsystem.

Figure NEW.1: Unique Discovery Service NQN retrieval for bidirectional authentication



If NVMe in-band authentication succeeds, then any supported commands for the associated queue type may be processed.

...