



LEGAL NOTICE:

© **Copyright 2008 to 2024 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 2024 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

Technical input submitted to the NVM Express® Workgroup is subject to the terms of the NVM Express® Participant’s agreement. Copyright © 2008 to 2024 NVM Express, Inc.

NVM Express® Technical Proposal (TP)

Technical Proposal ID	TP8024a mDNS Discovery update
Revision Date	2024.01.03
Builds on Specification(s)	TP8009 Automated Discovery of NVMe-oF Discovery Controllers for IP Networks
References	TP8010a NVMe-oF Centralized Discovery Controller

Technical Proposal Author(s)

Name	Company
Erik Smith	Dell

Technical Proposal Overview

This Technical Proposal obsoletes the DNS-SD subtype “_ddcpull._sub._nvme-disc”.

The “_ddcpull._sub._nvme-disc” subtype enabled Direct Discovery Controllers to announce this subtype and use it to signal Centralized Discovery Controllers that a pull registration should be performed.

DDCs may request a pull registration by using the Kickstart Discovery Request (refer to section 8.NEW.2.2.2 in TP8010).

Revision History

Revision Date	Change Description
2022.10.04	Initial draft
2022.10.05	Added TP8010 specific sections
2022.10.19	Incorporated changes from FMDS meeting
2022.12.12	Updated document name for phase 3 and member review.
2023.01.11	Incorporated comments from 30 day member review.
2023.01.19	Incorporated additional comments from 30 day member review.
2023.02.26	Integrated
2023.03.04	Incorporated comments from review
2023.12.20	<ul style="list-style-type: none">• Removing text that references the section removed by this TP.• Update for 2024 release.
2023.01.02	<ul style="list-style-type: none">• Used consistent colors.
2023.01.03	<ul style="list-style-type: none">• Spelled out acronyms
2024.06.11	<ul style="list-style-type: none">• Made ready for ratification

Description for Changes Document for TP8009 Automated Discovery of NVMe-oF Discovery Controllers for IP Networks

New Features/Feature Enhancements/Required Changes:

- Obsolescence of the DNS-SD subtype “_ddcpull._sub”.
 - The description of the DNS-SD subtype “_ddcpull._sub” has been updated to indicate it is obsolete.
 - All instances of the subtype “_ddcpull._sub” have been removed, along with any associated text describing how to use this subtype.
 - References
 - TP8009, TP8010a, TP8024

Description for Changes Document for TP8010a NVMe-oF Centralized Discovery Controller

New Features/Feature Enhancements/Required Changes:

- Delete section 8.NEW.2.2.1 mDNS Pull Registration Requests
 - References
 - TP8009, TP8010a, TP8024a

Markup Conventions:

Black:	Unchanged (however, hot links are removed)
Green:	Text from TP8009
Red Strikethrough:	Deleted
Blue:	New
Blue Highlighted:	TBD values, anchors, and links to be inserted in new text.
<Orange Bracketed>:	Notes to editor

Description of Specification Changes for TP8009 Automated Discovery of NVMe-oF Discovery Controllers for IP Networks

8.NEW.A.1 Discovery of NVMe-oF Discovery Controllers

8.NEW.A.1.1 Query

...

Figure NEW.ES5: mDNS Subtype

Subtype	Usage
"_cdc"	Used by CDC and DDC instances to detect the presence of a CDC service.
"_ddcpull"	Obsolete. May be used by CDC instances to detect the presence of DDC instances that are requesting a pull registration.

...

8.NEW.A.3 DDC Operation

8.NEW.A.3.1 DDC mDNS Initialization

During initialization (e.g., following a link transition or power cycle), before the DDC's mDNS responder function is enabled, the DDC shall probe to ensure the unique resource records the DDC are responsible for are unique on the local link (refer to section 8.1 in RFC 6762).

Upon successful completion of the probe, the DDC shall Announce (refer to section 8.2 in RFC 6762) its newly registered resource records. ~~If a DDC is configured for pull registration, the service name of "_ddcpull._sub._nvme-disc._<protocol>.local" is one of these resource records.~~

Upon announcing its resource records, if a DDC:

- has not been configured to perform push registration (refer to section 8.NEW.2.1) <Note to editor: 8.NEW.2.1 is located in TP-8010>, or has not been configured to request a pull registration (refer to section 8.NEW.2.2) <Note to editor: 8.NEW.2.2 is located in TP-8010> from a CDC, it may respond to mDNS queries for the service name of "_nvme-disc.<protocol>.local"; ~~or~~
- ~~b. has not been configured to request a pull registration (refer to section 8.NEW.2.2) <Note to editor: 8.NEW.2.2 is located in TP-8010>, it may respond to queries for the service name of "_ddcpull._sub._nvme-disc._<protocol>.local" as described in section 8.NEW.A.3.5; or~~
- b. has been configured to perform push registration (refer to section 8.NEW.2.1) <Note to editor: 8.NEW.2.1 is located in TP-8010> with a CDC, it should not respond to mDNS queries for the service name of "_nvme-disc.<protocol>.local", unless it has been administratively configured to do so or until it has performed a query and determined a CDC is not present as defined in section 8.NEW.A.3.3.

...

8.NEW.A.3.5 DDC response to mDNS queries

A DDC may respond to mDNS queries for the service names of ~~either:~~

"_nvme-disc.<protocol>.local"; ~~or~~
"~~_ddcpull._sub._nvme-disc.<protocol>.local~~".

mDNS responses to queries for ~~either of~~ these service names shall contain the information described in section 8.[NEW.A.1.2](#).

~~DDCs should only respond to mDNS queries for the service name of “_ddcpull._sub._nvme-disc.<protocol>.local” if the DDC is requesting a pull registration (refer to section [8.NEW.2.2](#)) <Note to editor: 8.NEW.2.2 is located in TP-8010> be performed.~~

DDCs should ignore mDNS queries for the service name of “_ddcpull._sub._nvme-disc.<protocol>.local” (refer to [Figure NEW.ES5](#)), as the subtype “_ddcpull._sub” is obsolete.

...

8.NEW.A.4.3 CDC Query

A CDC may query for ~~both~~ other CDC ~~and DDC~~ instances. When performed the mDNS or DNS query shall include a DNS PTR record (refer to RFC 6763) with the name in the form of:

“_cdc._sub._nvme-disc.<protocol>.<domain>”; ~~or~~
“~~_ddcpull._sub._nvme-disc.<protocol>.<domain>~~”.

The protocol field shall be set as shown in [Figure NEW.ES4](#).

The domain field shall be set as shown in [Figure NEW.ES6](#).

8.NEW.A.4.4 CDC Processing of DNS-SD records

Upon reception of an mDNS or DNS response that contains a DNS SRV record with the service name set to “_cdc._sub._nvme-disc.<protocol>.local” the CDC may provide an alert to the administrator to indicate the presence of more than one CDC in a broadcast domain.

~~Upon reception of an mDNS or DNS response that contains a DNS SRV record with the service name set to “_ddcpull._sub._nvme-disc.<protocol>.local” the CDC may choose to perform a pull registration (refer to section [8.NEW.2.2](#)) <Note to editor: 8.NEW.2.2 is located in TP-8010> for the responding DDC. If the CDC performs a pull registration, it shall use the IP address in the A or AAAA record as the destination IP address for a subsequent connect command.~~

A CDC should ignore an mDNS or DNS response that contains a DNS SRV record with the service name set to “_ddcpull._sub._nvme-disc.<protocol>.local”.

...

Description of Specification Changes for TP8010a Automated Discovery of NVMe-oF Discovery Controllers for IP Networks

8.NEW.2.2 Pull Registrations and Pull De-Registrations

~~A pull registration may be requested using either:~~

- ~~• an mDNS request and response sequence (refer to section [8.NEW.2.2.1](#)); or~~
- ~~• a kickstart discovery request and response sequence (refer to section [8.NEW.2.2.2](#)).~~

A pull registration may be requested using a kickstart discovery request and response sequence (refer to section [8.NEW.2.2.2](#)).

~~If the Centralized Discovery controller (CDC) and the Direct Discovery controller (DDC) are both contained in the same broadcast domain, then either an mDNS request and response sequence or a kickstart discovery request and response sequence may be used to request a pull registration.~~

~~If the CDC and the DDC are contained in separate broadcast domains, then a kickstart discovery request and response sequence may be used to request a pull registration.~~

Upon completion of acknowledging a pull registration request (e.g., after a KDRsp NVMe/TCP PDU (refer to the Kickstart Discovery Response PDU section in the NVMe TCP Transport Specification) has been sent by the [Centralized Discovery controller \(CDC\)](#)), the CDC performs the pull registration by:

1. sending a Connect command to the [Direct Discovery controller \(DDC\)](#) to establish an NVMe connection with that DDC; and
2. sending a Get Log Page command to the DDC with the Log Page Identifier (LID) field set to 70h to retrieve NVM subsystem discovery information contained in the associated Discovery log page from that DDC. The CDC may:
 - a. set the Port Local Entries Only (PLEO) bit to '1' in the Log Specific Parameter (LSP) field in that Get Log Page command to request that NVM subsystem discovery information entries for only NVM subsystem ports that are presented through the same NVM subsystem port on the DDC that received the Get Log Page command be returned; and
 - b. set the All NVM Subsystem Entries (ALLSUBE) bit to '1' in the LSP field in that Get Log Page command to request that records for all NVM subsystem ports be returned.

~~If kickstart discovery was used to request the pull registration, then the~~ [The](#) DDC should match the NQN contained in the CDC NVMe Qualified Name (CDCNQN) field of the KDRsp NVMe/TCP PDU with the NQN contained in the Host NVMe Qualified Name (HOSTNQN) field in the data portion of the Connect command to know that a CDC is performing a pull registration.

...

~~8.NEW.2.2.1 mDNS Pull Registration Requests~~

~~A Centralized Discovery controller (CDC) queries the local broadcast domain for any Direct Discovery controller (DDC) that requires a pull registration by sending an mDNS query for the service name of “_ddcpull._sub._nvme-disc._<protocol>.local” as described in section 8.NEW.A.4.3 <Note to editor: 8.NEW.A.4.3 is located in TP8009>. If the mDNS query from the CDC is completed by an mDNS response from a DDC containing a DNS SRV record with the service name set to “_ddcpull._sub._nvme-disc._<protocol>.local” (refer to section 8.NEW.A.4.4 <Note to editor: 8.NEW.A.4.4 is located in TP8009>), then the CDC performs a pull registration with the responding DDC as described in section 8.NEW.2.2.~~

~~A DDC may request a pull registration during initialization by sending an mDNS announcement containing a DNS SRV record with the service name set to “_ddcpull._sub._nvme-disc._<protocol>.local” as described in section 8.NEW.A.3.1 <Note to editor: 8.NEW.A.3.1 is located in TP8009>. If a CDC receives an mDNS announcement from a DDC containing a DNS SRV record with the service name set to “_ddcpull._sub._nvme-disc._<protocol>.local” (refer to section 8.NEW.A.4.4 <Note to editor: 8.NEW.A.4.4 is located in TP8009>), then the CDC performs a pull registration with the requesting DDC as described in section 8.NEW.2.2.~~