



LEGAL NOTICE:

© **Copyright 2008 to 2022 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 2022 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.
Redfish® is a registered trademark of DMTF.

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 4094 – Management Network Addresses
Revision Date	2022.04.14
Builds on Specification(s)	NVM Express® Base Specification 2.0b NVM Express® Management Interface Specification 1.2b
Ratified Technical Proposals Referenced	TP4074a – Defining Scope for Features

Technical Proposal Author(s)

Name	Company
Mark Carlson, Paul Suhler	Kioxia
David Black	Dell EMC
Michael Allison, Judy Brock	Samsung

Technical Proposal Overview

Simplify device management by providing a method to automatically obtain the management URIs for a device.

This technical proposal adds new mechanisms to the NVMe Base specification to enable the drive to indicate management agents. This method will enable a host to detect whether a network management interface is available for the device or the system that contains the device and to obtain the network addresses used to manage the device and/or the system that contains the device.

Revision History

Revision Date	Change Description
2021.11.29	Initial draft
2021.11.30	Recorded suggested changes from 2021-11-30 meeting of the FMDS TG.
2021.12.06	Changes from 2021-11-30 meeting of the FMDS TG and the 2021-12-06 meeting of the NVMe-MI TG: <ul style="list-style-type: none">• Moved two addresses into a new feature ID.• Moved two addresses into the command set independent Identify Controller data structure (CNS 01h).• Added TP4074 as a ratified referenced TP.
2021.12.07	Added notes from 2021-12-07 meeting of the FMDS TG.
2021.12.08	Changes from 2021-12-07 meeting of the FMDS TG:

Revision Date	Change Description
	<ul style="list-style-type: none"> • Moved the Management Address NVM Subsystem (MASUB) field and the Management Address Fabric Interface (MAFI) field from the Identify Controller data structure to descriptors in a new Management Address List Identify data structure). • Split the Management Address feature into the Management Controller Agent Address feature and the Host Management Agent Address feature. • Changed length of all URI fields to 252 bytes. • Updated the Technical Proposal Overview.
2021-12-14	<p>Changes from 2021-12-14 meeting of the FMDS TG:</p> <ul style="list-style-type: none"> • Added Management Address List to Figure 273 (Identify – CNS Values). • Reduced Management Address Type field from four bytes to one byte. • Increase URI field lengths from 252 to 508. • Renamed Management Controller Agent Address to System Management Agent Address. • Simplified lockdown of a feature to a single Lockdown command, specifying prohibited. • Rewrote various parts of the feature descriptions.
2022-01-06	<p>Changes from 2021-12-16 and 2022-01-06 meetings of the Technical WG and the 2022-01-04 meeting of the FMDS TG:</p> <ul style="list-style-type: none"> • Added definitions of host and embedded management controllers. • Renamed System Management Agent to Embedded Management Controller. • Moved non-changeable URIs from an Identify data structure to a log page. • Added Model section 8.TBD. • Added changes to NVMe-MI lists of log page support and feature support.
2022-01-10	<p>Changes from the 2022-01-10 meetings of the NVMe-MI TG and the 2022-01-11 FMDS TG:</p> <ul style="list-style-type: none"> • Base specification figures 25, 30, and 34: Added footnote prohibiting Set Features command to specify the Embedded Management Controller Address feature if the command was not received over the OOB Management Endpoint. • NVMe-MI: Removed Host Management Agent Address feature from the Features Supported figures for the I/O and Administrative controllers.

Revision Date	Change Description
2022-01-24	<ul style="list-style-type: none"> Fig. 316: Corrected "Uses Memory Buffer for Attributes" from No to Yes for both features. NVMe-MI: Added Host Management Agent Address feature to figures 124 and 125, per NVMe-MI TG meeting.
2022-01-25	<p>Changes from the 2022-01-25 meeting of the FMDS TG:</p> <ul style="list-style-type: none"> Deleted table footnotes prohibiting host from setting the Embedded Management Controller Address feature. Clarified wording of recommendations for the Command and Feature Lockdown log page.
2022-01-26	<ul style="list-style-type: none"> Specification list: Updated Base 1.0a to Base 1.0b. Specification list: Added NVMe-MI 1.2a. Added trademark information for Redfish to the first page of this TP and to the Base spec. Added Management Address Type FFh to indicate end of list. (Suggested by Mike Allison.)
2022-01-27	<p>Changes from the 2022-01-27 meeting of the Technical WG:</p> <ul style="list-style-type: none"> Aligned to NVMe-MI 1.2b. Figure FigMAL: Corrected byte offsets. 8.TBD: Removed "(Optional)" from heading. NVMe-MI Figure 126: Modified to align with 1.2b.
Phase 3	
2022-02-07	<ul style="list-style-type: none"> Changed filename. Updated copyright dates and symbols and trademark symbols. NVMe-MI Fig. 122: Modified heading row to align with NVMe-MI 1.2b.
Member Review	
2022-02-10	<ul style="list-style-type: none"> Accepted all changes. Changed filename.
2022-02-15	<ul style="list-style-type: none"> Resolution of editorial comments from Edward Hsieh (Silicon Motion).
2022-02-16	<ul style="list-style-type: none"> Proposed resolution of comments from Mike Allison and Judy Brock (Samsung).
2022-02-22	<ul style="list-style-type: none"> Proposed resolution of additional comments from Mike Allison, David Black, and Fred Knight. Added reference for RFC 4088. Revised description of MAT field value 2h. Revised description of MADRS field. Added description of scanning the Management Address List.
2022-03-14	<ul style="list-style-type: none"> Added ® symbol to "NVM Express" in various places.
Integration	
2022-03-17	<ul style="list-style-type: none"> Deleted all comments. Accepted all changes.
2022-04-11	<ul style="list-style-type: none"> Integrated

Revision Date	Change Description
2022-04-13	<ul style="list-style-type: none"><li data-bbox="418 241 1351 304">• In 1.5.TBD1 (definition of embedded management controller, change the last word in the definition from “controller” to “Management Endpoint”.<li data-bbox="418 312 1263 344">• Added change to Base Spec 2.0b Figure 202 (Log Page Identifiers).<li data-bbox="418 352 1351 384">• Added hexadecimal values for log and feature identifiers to three headings.<li data-bbox="418 392 964 424">• Replied to editor’s questions in comments.<li data-bbox="418 432 1224 447">• Updated new log page title to include “- Log Page”, per ECN110.

Description for Changes Document for NVM Express Base Specification 2.0b

New Features:

- Management Network Addresses (optional)
 - Description of change.
 - Added the Management Address List log page.
 - Added the Embedded Management Controller Address feature.
 - Added the Host Management Agent Address feature.
 - Added the Management Addresses section to the Extended Capabilities section.
 - References
 - Technical Proposal 4094

Description for Changes Document for NVM Express Management Interface Specification 1.2b

New Feature:

- Management Addresses (optional)
 - Added the following to the appropriate support tables:
 - Management Address List log page.
 - Embedded Management Controller Address feature.
 - Host Management Agent Address feature.
 - References:
 - Technical Proposal 4094

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.
<Green Bracketed>:	Notes to editor

Modify portions of NVMe Base Specification 2.0b as shown below:

Specification disclaimer page:

...

InfiniBand™ is a trademark and servicemark of the InfiniBand Trade Association.
Redfish® is a registered trademark of DMTF.

...

1 Introduction

...

1.5 Definitions

...

1.5.TBD1 embedded management controller

An embedded management controller is a Management Controller (refer to the NVM Express Management Interface Specification) that provides an external management interface (e.g., Redfish®), typically implemented via commands to the Management Endpoint.

1.5.TBD2 host management agent

A host management agent is a part of the host that provides an external management interface (e.g., Redfish) to external managers, typically via Admin commands to the controller (refer to the NVM Express Management Interface Specification).

...

1.8 References

...

RFC 4088, D. Black, K. McCloghrie, J. Schoenwaelder, "Uniform Resource Identifier (URI) Scheme for the Simple Network Management Protocol (SNMP)", June 2005. Available from <https://www.ietf.org/rfc.html>.

...

3 NVM Express Architecture

...

3.1 NVM Controller Architecture

...

3.1.2 Controller Types

...

3.1.2.1 I/O Controller

...

3.1.2.1.3 Features Support

...

Figure 25: I/O Controller – Feature Support

Feature Name	Feature Support Requirements ¹	Logged in Persistent Event Log ¹
...		
Embedded Management Controller Address	○	○
Host Management Agent Address	○	○
...		

...

3.1.2.2 Administrative Controller

...

3.1.2.2.3 Features Support

...

Figure 30: Administrative Controller – Feature Support

Feature Name	Feature Support Requirements ¹	Logged in Persistent Event Log ¹
...		
Embedded Management Controller Address	○	○
Host Management Agent Address	○	○
...		

...

3.1.2.3 Discovery Controller

...

3.1.2.3.4 Features Support

...

Figure 34: Discovery Controller – Feature Support

Feature Name	Feature Support Requirements ¹	Logged in Persistent Event Log ¹
...		
Embedded Management Controller Address	○	○
Host Management Agent Address	○	○
...		

...

5 Admin Command Set

...

5.15 Get Features command

...

Figure 194: Get Features – Feature Identifiers

Description	Section Defining Format of Attributes Returned
...	
Embedded Management Controller Address	5.27.1.TBD1
Host Management Agent Address	5.27.1.TBD2
...	

...

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
...			
18h	NVM subsystem	Management Address List	5.16.1.TBD
...			

...

5.16.1.TBD Management Address List (Log Identifier 18h)

This log consists of a Management Address List (refer to [Figure FigMAL](#)) of up to eight Management Address Descriptors (refer to [Figure FigMAD](#)). Management addresses are described in section [8.TBD](#).

Figure FigMAL: Management Address List Log Page

Bytes	Description
511:00	Management Address Descriptor 0: This field contains the first Management Address Descriptor (refer to Figure FigMAD), if any. If the Management Address Type field is set to FFh, then subsequent Management Address Descriptors are reserved.
1023:512	Management Address Descriptor 1: This field contains the second Management Address Descriptor (refer to Figure FigMAD), if any. If the Management Address Type field is set to FFh, then subsequent Management Address Descriptors are reserved.
...	...
4095:3584	Management Address Descriptor 7: This field contains the last Management Address Descriptor (refer to Figure FigMAD), if any.

Figure FigMAD describes the Management Address Descriptor.

Figure FigMAD: Management Address Descriptor

Bytes	Description												
00	Management Address Type (MAT): This field describes the contents of the Management Address field:												
	<table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>The Management Address field is reserved.</td> </tr> <tr> <td>1h</td> <td>The Management Address field contains the address of a management agent in the NVM subsystem (e.g., in an Ethernet-attached SSD).</td> </tr> <tr> <td>2h</td> <td>The Management Address field contains the address of a fabric interface manager (e.g., a Service URI for the Simple Network Management Protocol (SNMPv3) agent as defined in RFC 4088).</td> </tr> <tr> <td>FFh</td> <td>The Management Address field is reserved and all subsequent Management Address Descriptors, if any, are reserved.</td> </tr> <tr> <td>Other values</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Description	0h	The Management Address field is reserved.	1h	The Management Address field contains the address of a management agent in the NVM subsystem (e.g., in an Ethernet-attached SSD).	2h	The Management Address field contains the address of a fabric interface manager (e.g., a Service URI for the Simple Network Management Protocol (SNMPv3) agent as defined in RFC 4088).	FFh	The Management Address field is reserved and all subsequent Management Address Descriptors, if any, are reserved.	Other values	Reserved
	Value	Description											
	0h	The Management Address field is reserved.											
	1h	The Management Address field contains the address of a management agent in the NVM subsystem (e.g., in an Ethernet-attached SSD).											
	2h	The Management Address field contains the address of a fabric interface manager (e.g., a Service URI for the Simple Network Management Protocol (SNMPv3) agent as defined in RFC 4088).											
FFh	The Management Address field is reserved and all subsequent Management Address Descriptors, if any, are reserved.												
Other values	Reserved												
03:01	Reserved												
511:04	Management Address (MADRS): A uniform resource indicator (URI; refer to RFC 3986) as a UTF-8 null-terminated string, containing the address of a management entity of the type indicated by the MAT field.												

Because the list is able to contain from zero to eight Management Address Descriptors with a MAT field less than FFh, a host scans the list starting with Management Address Descriptor 0, and ending with either Management Address Descriptor 7 or a Management Address Descriptor having a MAT field set to FFh.

...

5.27 Set Features command

...

5.27.1 Feature Specific Information

...

Figure 316: Set Features – Feature Identifiers

Feature Identifier	Current Setting Persists Across Power Cycle and Reset ²	Uses Memory Buffer for Attributes	Feature Name	Scope ⁶
...				
78h	Yes	Yes	Embedded Management Controller Address	NVM subsystem
79h	Yes	Yes	Host Management Agent Address	NVM subsystem
...				

< **Note to Editor:** Please assign FeatEMCAh and FeatHMAAh from the range “Reserved for Management Features.” >

5.27.1.TBD1 Embedded Management Controller Address (Feature Identifier 78h)

This Feature configures a URI (refer to RFC 3986) containing the address of a management agent provided by the system (e.g., a Management Controller (refer to the NVMe Management Interface Specification) or an enclosure manager) for management of the NVM subsystem.

If a Set Features command is issued for this Feature, the data structure specified in Figure FigEMCA is transferred in the data buffer for that command.

If a Get Features command is issued for this Feature, the data structure specified in Figure FigEMCA is returned in the data buffer for that command.

Figure FigEMCA: Embedded Management Controller Address

Bytes	Definition
03:00	Reserved
511:04	Embedded Management Controller Address (EMCA): A URI (refer to RFC 3986), containing the address of a management agent provided by the system.

The Command and Feature Lockdown log page (refer to section 5.16.1.20), if supported, should indicate that this Feature is supported to be prohibited from execution.

5.27.1.TBD2 Host Management Agent Address (Feature Identifier 79h)

This Feature configures a URI (refer to RFC 3986) containing the address of a management agent provided by host software for management of the NVM subsystem.

If a Set Features command is issued for this Feature, the data structure specified in Figure FigHMAA is transferred in the data buffer for that command.

If a Get Features command is issued for this Feature, the data structure specified in Figure FigHMAA is returned in the data buffer for that command.

Figure FigHMAA: Host Management Agent Address

Bytes	Definition
03:00	Reserved
511:04	Host Management Agent Address (HMAA): A URI (refer to RFC 3986), containing the address of a management agent residing in host software.

The Command and Feature Lockdown log page (refer to section 5.16.1.20), if supported, should indicate that this Feature is supported to be prohibited from execution.

...

8 Extended Capabilities

...

8.TBD Management Addresses

Various entities on a network are able to request a management agent to perform management operations on NVM subsystems. A controller in an NVM subsystem use the Management Addresses capability to indicate the network addresses of those management agents. When an NVM subsystem is provisioned in a storage system, the management addresses are established in the controller.

Management agents are able to be located in various networked entities, including:

- NVM subsystems;
- Fabric interface managers;
- Embedded management controllers; and
- Host software.

Each management address is represented as a uniform resource indicator (URI; refer to RFC 3986).

The address of a management agent contained in an NVM subsystem (e.g., an Ethernet-attached SSD) is indicated in the Management Address List log page (refer to section 5.16.1.TBD), in a Management Address Descriptor indicating a Management Address Type of 1h. The method by which the address is determined is outside the scope of this specification.

The address of a management agent contained in a fabric interface manager is indicated in the Management Address List log page (refer to section 5.16.1.TBD), in a Management Address Descriptor indicating a Management Address Type of 2h. The method by which the address is determined is outside the scope of this specification.

The address of a management agent contained in an embedded management controller (e.g., a BMC) is set and retrieved using the Embedded Management Controller Address feature (refer to section 5.27.1.TBD1). The Embedded Management Controller Address feature is intended to be set only by the storage system containing the embedded management controller. The embedded management controller is able to ensure this by issuing a Lockdown command (refer to section 5.19) with the OFI field specifying this Feature, and the PRHBT bit set to '1' and the IFC field set to Admin Submission Queue (i.e., 00b).

The address of a management agent contained in host software is set and retrieved using the Host Management Controller Address feature (refer to section 5.27.1.TBD2). The Host Management Agent Address feature is intended to be set only by host software. Host software is able to ensure this by issuing a Lockdown command (refer to section 5.19) with the OFI field specifying this Feature, and the PRHBT bit set to '1' and the IFC field set to Out-of-band on a Management Endpoint (i.e., 10b).

< end of changes to NVMe Base specification >

Modify portions of NVMe-MI 1.2b as shown below:

6 NVM Express Admin Command Set

...

6.3 Get Log Page

Figure 122 defines the log pages that are mandatory, optional, and prohibited for SMBus/I2C and PCIe VDM Management Endpoint on NVMe Storage Devices and NVMe Enclosures.

Figure 122: Management Endpoint - Log Page Support

Log Page Name ³	Log Identifier	Requirements ¹	
		NVMe Storage Device	NVMe Enclosure
...			
Management Address List	18h	○	○
...			

...

6.5 Set Features and Get Features

Figure 124 defines features that are mandatory or optional for an I/O Controller. Refer to the NVM Express Base Specification for the definition of Set Features and Get Features commands and I/O Controllers.

All Feature Identifiers supported shall be supported if received in-band on an NVMe Controller or received out-of-band on a Management Endpoint unless otherwise stated.

Figure 124: I/O Controller – Feature Support

Feature Name	Feature Support Requirements ¹	Logged in Persistent Event Log ¹
...		
Embedded Management Controller Address	○	○
Host Management Agent Address	○	○
...		

Figure 125 defines features that are mandatory or optional for an Administrative Controller. Refer to the NVM Express Base Specification for the definition of Set Features and Get Features commands and Administrative Controller.

Figure 125: Administrative Controller – Feature Support

Feature Name	Feature Support Requirements ¹	Logged in Persistent Event Log ¹
...		
Embedded Management Controller Address	○	○
Host Management Agent Address	○	○
...		

Figure 126 define the features that are mandatory, optional, and prohibited for SMBus/I2C and PCIe VDM Management Endpoints on NVMe Storage Devices and NVMe Enclosures.

Figure 126: Management Endpoint - Feature Support

Feature Name ²	Feature Identifier	Support Requirements ¹	
		NVMe Storage Device	NVMe Enclosure
...			
Embedded Management Controller Address	78h	O	O
Host Management Agent Address	79h	O	O
...			
Notes:			
<ol style="list-style-type: none"> 1. O = Optional, M = Mandatory, P = Prohibited for Set Features/Optional for Get Features. 2. Refer to the NVM Express Base Specification unless another footnote specifies otherwise. 3. Refer to the NVM Command Set Specification. 4. Refer to the Key Value Command Set Specification. 			

< end of changes to NVMe Management Interface specification >