



LEGAL NOTICE:

© Copyright 2007 - 2020 NVM Express, Inc. ALL RIGHTS RESERVED.

This NVM Express revision 1.4 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 NVM Express revision 1.4 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: "© 2007 - 2020 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

Technical input submitted to the NVM Express™ Workgroup is subject to the terms of the NVM Express™ Participant's agreement. Copyright © 2014-2019 NVMe™ Corporation.

NVM Express Technical Proposal for New Feature

Technical Proposal ID	TP 4046
Change Date	09/08/2020
Builds on Specification	NVM Express 1.4 NVM Express Management Interface Revision 1.1

Technical Proposal Author(s)

Name	Company
Kevin Marks, Simon Kan, Austin Bolen	Dell
Craig Lucero, Jeff Wiles	HP
David Derosa, Curtis Ballard	HPE
Paul Suhler	KIOXIA

This technical proposal adds a Command and Feature Lockdown feature that allows hosts and management entities to prohibit the execution of some commands and Set Feature command feature identifiers. The proposal includes a new admin command along with a log page that indicates which of these commands and Set Feature command feature identifiers may be prohibited from execution.

Commands and Set Feature command feature identifiers that are commanded to be prohibited are prohibited from execution until a power cycle or un-prohibited by the Command and Feature Lockdown command. The prohibition or allowance of execution may be specified independently for the in-band vs out-of-band paths.

Once a command or feature identifier is prohibited, the device behavior for the path that the command is prohibited on is to return an error defined in the proposal depending on command set.

Revision History

Revision Date	Change Description
5/08/19	Initial Creation
6/25/19	Increased field sizes for types of command sets to block
7/29/19	Update to add prohibited behavior of not supporting command or feature id.
8/15/19	Extended LSP field instead of creating a second LSP2 field and updated Logs that use this field.
8/21/19	Updated based on Comments from Mike A (Intel)
11/13/19	Change to only apply to optional features/commands for Admin command set. For MI and PCIe command set, may block mandatories with MI error code.
11/21/19	Added table headers for readability and added a couple of references.
12/02/19	Updated based on comments from Mike Allison and MI WG call.
1/13/20	Changes based on Fred Knight comments from Phase 2 exit and requests from Intel to allow a persistent mode of operation, in which only the Lockdown command becomes unblocked after a power cycle. Also added the ability to return an error to a blocked command instead of showing as not supported. This is enabled via a host behavior support set features (Not Posted)
1/29/20	After 1/23/2020 WG meeting, have gone back to original proposed error when locked for Admin commands. This is potentially not backwards compatible with an older driver. Maintained the persistence text until Mike get back with if Intel still wants.
2/06/20	Removed Persistent mode and some editorial changes reported by Jim Hatfield.
4/15/2020	Integrated into NVMe 1.4 Base Specification and NVMe-MI revision 1.1a specification
4/20/2020	For NVMe-MI changed "executing" to "processing".
7/21/2020	Added in sections 7.1 and 7.2 to list that the Lockdown command is optional. Reformatted heading to show in navigation.
7/24/2020	Accepted all changes to get ready for member review.
9/8/2020	Integrated into the NVMe Management Interface Specification, Revision 1.1 and the NVM Express Base Specification.

Editorial convention:

<<..Editorial text ..>>

~~Deleted Text~~

New Text

Value may change when approved

Description of Specification Changes to NVM Express 1.4

4.6.1.2.1 Generic Command Status Definition

<<..>>

Figure 126: Status Code – Generic Command Status Values

<<.Add Status value row to Figure 126.>>

Value	Description
22h	Transient Transport Error: A transient transport error was detected. If the command is retried on the same controller, the command is likely to succeed. A command that fails with a transient transport error four or more times should be treated as a persistent transport error that will not succeed if retried on the same controller.
23h	Command Prohibited by Command and Feature Lockdown: The command was aborted due to command execution being prohibited by the Command and Feature Lockdown (refer to section 8.2).
24h to 7Fh	Reserved

4.6.1.2.2 Command Specific Status Definition

<<..>>

Figure 128: Status Code – Command Specific Status Values

<<.Add Status value row to Figure 128.>>

Value	Description	Commands Affected
25h	ANA Attach Failed	Namespace Attachment
28h	Prohibition of Command Execution Not Supported	Lockdown
29h to 6Fh	Reserved	

5 Admin Command Set

<<..>>

<<.. Add Opcode row into table 139 ..>>

Figure 139: Opcodes for Admin Commands

Opcode by Field			Combined Opcode ¹	Namespace Identifier Used ²	Command
(07)	(06:02)	(01:00)			
Generic Command	Function	Data Transfer ³			
<<..>>	<<..>>	<<..>>	<<..>>	<<..>>	<<..>>
0b	010 01b	00b	24h	No	Lockdown
<<..>>	<<..>>	<<..>>	<<..>>	<<..>>	<<..>>

<<..>>

5.x Lockdown command

The Lockdown command is used to control the Command and Feature Lockdown feature (refer to section 8.z) which configures the prohibition or allowance of execution of the specified command or Set Features command targeting a specific Feature Identifier.

After a successful completion of a Lockdown command prohibiting a command or Feature Identifier, all controllers, if applicable, and all management endpoints, if applicable, in the NVM subsystem behave as described in 8.z.

The Lockdown command uses Command Dword 10 (refer to figure FTBD1) and Command Dword 14 (refer to figure FTBD2). All other command specific fields are reserved.

Figure FTBD1: Lockdown – Command Dword 10

Bits	Description														
31:16	Reserved														
15:8	Opcode or Feature Identifier (OFI): This field specifies the command opcode or Set Features Feature Identifier identified by the Scope field.														
07:06	Reserved														
06:05	Interface (IFC): This field identifies the interfaces affected by this command. The actions of this command apply if a command is received on the specified interfaces.. <table><tr><th>Value</th><th>Affected Interfaces</th></tr><tr><td>00b</td><td>Admin Submission Queue</td></tr><tr><td>01b</td><td>Admin Submission Queue and out-of-band on a Management Endpoint</td></tr><tr><td>10b</td><td>Out-of-band on a Management Endpoint</td></tr><tr><td>11b</td><td>Reserved</td></tr></table>	Value	Affected Interfaces	00b	Admin Submission Queue	01b	Admin Submission Queue and out-of-band on a Management Endpoint	10b	Out-of-band on a Management Endpoint	11b	Reserved				
Value	Affected Interfaces														
00b	Admin Submission Queue														
01b	Admin Submission Queue and out-of-band on a Management Endpoint														
10b	Out-of-band on a Management Endpoint														
11b	Reserved														
04	Prohibit (PRHBT): This bit specifies whether to prohibit or allow the command opcode or Set Features Feature Identifier specified by this command. If set to '1', then this command prohibits the execution of the command based on other fields specified in Dword 10. If cleared to '0', then this command allows the execution of the command based on other fields specified in Dword 10.														
03:00	Scope (SCP): This field specifies the contents of the Opcode or Feature Identifier field.. <table><tr><th>Value</th><th>Opcode or Feature Identifier Definition</th></tr><tr><td>0h</td><td>Admin command opcode</td></tr><tr><td>1h</td><td>Reserved</td></tr><tr><td>2h</td><td>A Set Features Feature Identifier</td></tr><tr><td>3h</td><td>Management Interface Command Set opcode (refer to the NVM Express Management Interface Specification)</td></tr><tr><td>4h</td><td>PCIe Command Set opcode (refer to the NVM Express Management Interface Specification)</td></tr><tr><td>5h-Fh</td><td>Reserved</td></tr></table>	Value	Opcode or Feature Identifier Definition	0h	Admin command opcode	1h	Reserved	2h	A Set Features Feature Identifier	3h	Management Interface Command Set opcode (refer to the NVM Express Management Interface Specification)	4h	PCIe Command Set opcode (refer to the NVM Express Management Interface Specification)	5h-Fh	Reserved
Value	Opcode or Feature Identifier Definition														
0h	Admin command opcode														
1h	Reserved														
2h	A Set Features Feature Identifier														
3h	Management Interface Command Set opcode (refer to the NVM Express Management Interface Specification)														
4h	PCIe Command Set opcode (refer to the NVM Express Management Interface Specification)														
5h-Fh	Reserved														

If the controller supports selection of a UUID:

- by the Lockdown command; and
- by the Set Features command (refer to Figure 271 and section 8.24) and for the vendor specific Feature Identifier specified by the Opcode or Feature Identifier field, if the Scope field is set to 2h, then Command Dword 14 (refer to Figure FTBD2) is used to specify a UUID Index value.

If the controller does not support selection of a UUID:

- by the Lockdown command;
- by the Set Features command; or

- c) for the vendor specific feature identifier specified by the Opcode or Feature Identifier field, if the Scope field is set to 2h.

then Command Dword 14 does not specify a UUID Index value. If the Scope field is not set to 2h, then UUID Index field is ignored.

Figure FTBD2: Lockdown – Command Dword 14

Bits	Description
<u>31:07</u>	<u>Reserved</u>
<u>06:00</u>	<u>UUID Index: Refer to Figure 498.</u>

If a controller processes this command specifying a command opcode or Feature Identifier that is not supported as being prohibitable, then the command shall be aborted with a status code of Prohibition of Command Execution Not Supported.

If a controller processes this command with the Interface field set to 01b or 10b and the NVM subsystem does not contain a Management Endpoint, then the command shall be aborted with a status code of Invalid Field in Command.

If a controller processes this command with the Interface field set to 00b or 01b and the Scope field is set to 4h, then the command shall be aborted with a status code of Invalid Field in Command.

It is not an error to attempt to prohibit a command or Feature Identifier that is already prohibited from execution or allow a command or Feature Identifier that is already allowed to be executed.

5.x.1 Command Completion

Upon completion of the Lockdown command, the controller posts a completion queue entry to the Admin Completion Queue.

Lockdown command specific status values are defined in Figure FTBD3.

Figure FTBD3: Lockdown – Command Specific Status Values

Value	Description
<u>28h</u>	<u>Prohibition of Command Execution Not Supported: The command was aborted due to the specified opcode or Feature Identifier not supporting being prohibited from execution by the command.</u>

5.14 Get Log Page command

<<..>>

Figure 186: Get Log Page – Command Dword 10

Bits	Description
<u><<..>></u>	<u><<..>></u>

Bits	Description
15	Retain Asynchronous Event (RAE): This bit specifies when to retain or clear an Asynchronous Event. If this bit is cleared to '0', the corresponding Asynchronous Event is cleared after the command completes successfully. If this bit is set to '1' the corresponding Asynchronous Event is retained (i.e., not cleared) after the command completes successfully. Host software should clear this bit to '0' for log pages that are not used with Asynchronous Events. Refer to section 5.2
14:12	Reserved
14:08	Log Specific Field (LSP): If not defined for the log specified by the Log Page Identifier field, this field is reserved.

5.14.1 Log Specific Information

<<..>>

Figure 191: Get Log Page – Log Page Identifiers

<<.. Add log row into Figure 191 ..>>

Log Identifier	Scope	Log Page Name	Reference Section
<<..>>	<<..>>	<<..>>	<<..>>
0Fh	NVM subsystem	Endurance Group Event Aggregate	5.14.1.15
14h	NVM subsystem	Command and Feature Lockdown ⁵	5.14.1y
10h to 6Fh	Reserved		

<<..>>

5.14.1.7 Telemetry Host-Initiated (Log Identifier 07h)

<<..>>

Figure 201: Command Dword 10 – Log Specific Field

Bits	Description
14:09	Reserved
08	<p>Create Telemetry Host-Initiated Data: If set to '1', then the controller shall capture the Telemetry Host-Initiated Data representing the internal state of the controller at the time the associated Get Log Page command is processed. If cleared to '0', then the controller shall not update the Telemetry Host-Initiated Data. The Host-Initiated Data shall not change until the controller processes:</p> <ul style="list-style-type: none"> a) a subsequent Telemetry Host-Initiated Log with this bit set to '1'; b) a Firmware Commit command; or c) a power on reset.

<<..>>

5.14.1.12 Asymmetric Namespace Access (Log Identifier 0Ch)

<<...>>

Figure 207: Command Dword 10 – Log Specific Field

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

<<...>>

5.14.1.13 Persistent Event Log (Log Identifier 0Dh)

<<...>>

Figure 210: Command Dword 10 – Log Specific Field

Bits	Description
14:10	Reserved
09:08	Action: This field specifies the action the controller shall take during processing this Get Log Page command. <<...>>

<<...>>

5.14.1.14 Command and Feature Lockdown (Log Identifier 14h)

This log page is used to indicate which commands and Set Features Feature Identifiers are supported to be prohibited from execution using the Command and Feature Lockdown (refer to section 8.2) and which commands are currently prohibited if received on an NVM Express controller Admin Submission Queue or received out-of-band on a Management Endpoint (refer to the NVM Express Management Interface Specification). This log page uses the Log Specific Field field (refer to figure FTBD4) and may use the UUID Index field in the Get Log Page command to specify the scope and content of the list returned in the Command and Feature Identifier List field of this log page. The UUID Index field may be used if the Scope field is set to 2h, allowing returning of vendor specific Set Features Feature Identifier lockdown information.

Figure FTBD4: Command Dword 10 – Log Specific Field

Bits	Description
14	Reserved

Figure FTBD4: Command Dword 10 – Log Specific Field

Bits	Description														
13:12	<p>Contents (CNTTS): This field in combination with the Scope field specifies the contents of the Command and Feature Identifier List field in the log page.</p> <table> <tr> <th>Value</th><th>Command and Feature Identifier List Definition</th></tr> <tr> <td>00b</td><td>List of command opcodes or Feature Identifiers based on the Scope field that are supported to be prohibited</td></tr> <tr> <td>01b</td><td>List of command opcodes or Feature Identifiers based on the Scope field that are currently prohibited if received on an NVM Express controller submission queue</td></tr> <tr> <td>10b</td><td>List of command opcodes or Feature Identifiers based on the Scope field that are currently prohibited if received out-of-band on a Management Endpoint</td></tr> <tr> <td>11b</td><td>Reserved</td></tr> </table>	Value	Command and Feature Identifier List Definition	00b	List of command opcodes or Feature Identifiers based on the Scope field that are supported to be prohibited	01b	List of command opcodes or Feature Identifiers based on the Scope field that are currently prohibited if received on an NVM Express controller submission queue	10b	List of command opcodes or Feature Identifiers based on the Scope field that are currently prohibited if received out-of-band on a Management Endpoint	11b	Reserved				
Value	Command and Feature Identifier List Definition														
00b	List of command opcodes or Feature Identifiers based on the Scope field that are supported to be prohibited														
01b	List of command opcodes or Feature Identifiers based on the Scope field that are currently prohibited if received on an NVM Express controller submission queue														
10b	List of command opcodes or Feature Identifiers based on the Scope field that are currently prohibited if received out-of-band on a Management Endpoint														
11b	Reserved														
11:08	<p>Scope (SCP): This field in combination with the Contents field specifies the contents of the Command and Feature Identifier List field in the log page.</p> <table> <tr> <th>Value</th><th>Command and Feature Identifier List Contents</th></tr> <tr> <td>0h</td><td>List of Admin Command Set opcodes</td></tr> <tr> <td>1h</td><td>Reserved</td></tr> <tr> <td>2h</td><td>List of Feature Identifiers</td></tr> <tr> <td>3h</td><td>List of a Management Interface Command Set opcodes (refer to the NVM Express Management Interface Specification)</td></tr> <tr> <td>4h</td><td>List of a PCIe Command Set opcodes (refer to the NVM Express Management Interface Specification)</td></tr> <tr> <td>5h to Fh</td><td>Reserved</td></tr> </table>	Value	Command and Feature Identifier List Contents	0h	List of Admin Command Set opcodes	1h	Reserved	2h	List of Feature Identifiers	3h	List of a Management Interface Command Set opcodes (refer to the NVM Express Management Interface Specification)	4h	List of a PCIe Command Set opcodes (refer to the NVM Express Management Interface Specification)	5h to Fh	Reserved
Value	Command and Feature Identifier List Contents														
0h	List of Admin Command Set opcodes														
1h	Reserved														
2h	List of Feature Identifiers														
3h	List of a Management Interface Command Set opcodes (refer to the NVM Express Management Interface Specification)														
4h	List of a PCIe Command Set opcodes (refer to the NVM Express Management Interface Specification)														
5h to Fh	Reserved														

If a UUID Index is specified in the Get Log Page command (refer to section 5.14) with the Scope field is set to 2h, then the controller should return vendor specific Set Features lockdown information defined by the vendor identified by the specified UUID index field. If the Scope field is not set to 2h, then the UUID index field is ignored.

If a controller processes this command with the Contents field set to 10b and the NVM subsystem does not contain a Management Endpoint, then the command shall be aborted with a status code of Invalid Field in Command.

The log page returned is defined in Figure FTBD5.

Figure FTBD5: Get Log Page – Command and Feature Lockdown Log

Bytes	Description																																
0	<table> <tr> <th>Bits</th><th>Description</th></tr> <tr> <td>7:6</td><td>Reserved</td></tr> <tr> <td>5:4</td><td> <p>Contents Selected (CS): This field in combination with the Scope Selected field indicates the contents of the Command and Feature Identifier List field in the log page. The Content Selected field is specified by the contents of the Contents field in the Log Specific Field field of the Get Log Page command.</p> <table> <tr> <th>Value</th><th>Description</th></tr> <tr> <td>00b</td><td>List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are supported to be prohibited</td></tr> <tr> <td>01b</td><td>List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are currently prohibited if received on an NVM Express controller submission queue</td></tr> <tr> <td>10b</td><td>List contains command opcodes or Set Features Feature Identifiers based on the Scope field that are currently prohibited if received out-of-band on a Management Endpoint</td></tr> <tr> <td>11b</td><td>Reserved</td></tr> </table> </td></tr> <tr> <td>3:0</td><td> <p>Scope Selected (SS): This field in combination with the Contents Selected field indicates what the Command and Feature Identifier List field contains in the log page. The Scope Selected field is specified by the contents of the Scope field in the Log Specific field of the Get Log Page command.</p> <table> <tr> <th>Value</th><th>Description</th></tr> <tr> <td>0h</td><td>List contains Admin Command Set opcodes</td></tr> <tr> <td>1h</td><td>Reserved</td></tr> <tr> <td>2h</td><td>List contains Set Features Feature Identifiers</td></tr> <tr> <td>3h</td><td>List contains Management Interface Command Set opcodes</td></tr> <tr> <td>4h</td><td>List contains PCIe Command Set opcodes</td></tr> <tr> <td>5h to Fh</td><td>Reserved</td></tr> </table> </td></tr> </table>	Bits	Description	7:6	Reserved	5:4	<p>Contents Selected (CS): This field in combination with the Scope Selected field indicates the contents of the Command and Feature Identifier List field in the log page. The Content Selected field is specified by the contents of the Contents field in the Log Specific Field field of the Get Log Page command.</p> <table> <tr> <th>Value</th><th>Description</th></tr> <tr> <td>00b</td><td>List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are supported to be prohibited</td></tr> <tr> <td>01b</td><td>List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are currently prohibited if received on an NVM Express controller submission queue</td></tr> <tr> <td>10b</td><td>List contains command opcodes or Set Features Feature Identifiers based on the Scope field that are currently prohibited if received out-of-band on a Management Endpoint</td></tr> <tr> <td>11b</td><td>Reserved</td></tr> </table>	Value	Description	00b	List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are supported to be prohibited	01b	List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are currently prohibited if received on an NVM Express controller submission queue	10b	List contains command opcodes or Set Features Feature Identifiers based on the Scope field that are currently prohibited if received out-of-band on a Management Endpoint	11b	Reserved	3:0	<p>Scope Selected (SS): This field in combination with the Contents Selected field indicates what the Command and Feature Identifier List field contains in the log page. The Scope Selected field is specified by the contents of the Scope field in the Log Specific field of the Get Log Page command.</p> <table> <tr> <th>Value</th><th>Description</th></tr> <tr> <td>0h</td><td>List contains Admin Command Set opcodes</td></tr> <tr> <td>1h</td><td>Reserved</td></tr> <tr> <td>2h</td><td>List contains Set Features Feature Identifiers</td></tr> <tr> <td>3h</td><td>List contains Management Interface Command Set opcodes</td></tr> <tr> <td>4h</td><td>List contains PCIe Command Set opcodes</td></tr> <tr> <td>5h to Fh</td><td>Reserved</td></tr> </table>	Value	Description	0h	List contains Admin Command Set opcodes	1h	Reserved	2h	List contains Set Features Feature Identifiers	3h	List contains Management Interface Command Set opcodes	4h	List contains PCIe Command Set opcodes	5h to Fh	Reserved
Bits	Description																																
7:6	Reserved																																
5:4	<p>Contents Selected (CS): This field in combination with the Scope Selected field indicates the contents of the Command and Feature Identifier List field in the log page. The Content Selected field is specified by the contents of the Contents field in the Log Specific Field field of the Get Log Page command.</p> <table> <tr> <th>Value</th><th>Description</th></tr> <tr> <td>00b</td><td>List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are supported to be prohibited</td></tr> <tr> <td>01b</td><td>List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are currently prohibited if received on an NVM Express controller submission queue</td></tr> <tr> <td>10b</td><td>List contains command opcodes or Set Features Feature Identifiers based on the Scope field that are currently prohibited if received out-of-band on a Management Endpoint</td></tr> <tr> <td>11b</td><td>Reserved</td></tr> </table>	Value	Description	00b	List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are supported to be prohibited	01b	List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are currently prohibited if received on an NVM Express controller submission queue	10b	List contains command opcodes or Set Features Feature Identifiers based on the Scope field that are currently prohibited if received out-of-band on a Management Endpoint	11b	Reserved																						
Value	Description																																
00b	List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are supported to be prohibited																																
01b	List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are currently prohibited if received on an NVM Express controller submission queue																																
10b	List contains command opcodes or Set Features Feature Identifiers based on the Scope field that are currently prohibited if received out-of-band on a Management Endpoint																																
11b	Reserved																																
3:0	<p>Scope Selected (SS): This field in combination with the Contents Selected field indicates what the Command and Feature Identifier List field contains in the log page. The Scope Selected field is specified by the contents of the Scope field in the Log Specific field of the Get Log Page command.</p> <table> <tr> <th>Value</th><th>Description</th></tr> <tr> <td>0h</td><td>List contains Admin Command Set opcodes</td></tr> <tr> <td>1h</td><td>Reserved</td></tr> <tr> <td>2h</td><td>List contains Set Features Feature Identifiers</td></tr> <tr> <td>3h</td><td>List contains Management Interface Command Set opcodes</td></tr> <tr> <td>4h</td><td>List contains PCIe Command Set opcodes</td></tr> <tr> <td>5h to Fh</td><td>Reserved</td></tr> </table>	Value	Description	0h	List contains Admin Command Set opcodes	1h	Reserved	2h	List contains Set Features Feature Identifiers	3h	List contains Management Interface Command Set opcodes	4h	List contains PCIe Command Set opcodes	5h to Fh	Reserved																		
Value	Description																																
0h	List contains Admin Command Set opcodes																																
1h	Reserved																																
2h	List contains Set Features Feature Identifiers																																
3h	List contains Management Interface Command Set opcodes																																
4h	List contains PCIe Command Set opcodes																																
5h to Fh	Reserved																																
2:1	Reserved																																
3	Length (LNGTH): This field indicates the length in bytes (n) of the Command and Feature Identifier List field that follow in the log page. If the Command and Feature Identifier List field contains no coded values, then this field shall be cleared to 0h.																																
n+3:4	Command and Feature Identifier List (CFIL): The contents of this field are dependent on the setting of the Contents Selected field and Scope Selected field. This field contains a list of coded values identified by the Scope Selected field and the Content Selected field. The list shall be in order from lowest numerical value to highest numerical value.																																
511:n+4	Reserved																																

5.15.2.2 Identify Controller data structure (CNS 01h)

<<..>>

Figure 247: Identify – Identify Controller Data Structure

<<.. Add Opcode row into Figure 247..>>

Bytes	O/M ¹	Description
Admin Command Set Attributes & Optional Controller Capabilities		
257:256	M	<p>Optional Admin Command Support (OACS): This field indicates the optional Admin commands and features supported by the controller. Refer to section 5.</p> <p>Bits 15:11 are reserved.</p> <p>Bit 10 if set to '1', then the controller supports the Command and Feature Lockdown (refer to section 8.z). If cleared to '0', then the controller does not support the Command and Feature Lockdown. This value shall be the same for all controllers in the NVM Subsystem.</p> <p>Bit 9 if set to '1', then the controller supports the Get LBA Status capability (refer to section 8.22). If cleared to '0', then the controller does not support the Get LBA Status capability.</p> <p><<...>></p>

7.1 I/O Controller

<<...>>

Figure 417: I/O Controller – Admin Command Support

Command	Command Support Requirements ¹
Delete I/O Submission Queue	M
Create I/O Submission Queue	M
Get Log Page	M
Delete I/O Completion Queue	M
Create I/O Completion Queue	M
Identify	M
Abort	M
Set Features	M
Get Features	M
Asynchronous Event Request	M
Namespace Management	O
Firmware Commit	O
Firmware Image Download	O
Device Self-test	O
Namespace Attachment	O
Keep Alive	NOTE 2
Directive Send	O
Directive Receive	O
Virtualization Management	O
NVMe-MI Send	O
NVMe-MI Receive	O
Doorbell Buffer Config	O
Lockdown	O
NVMe over Fabrics Commands	Refer to the NVMe over Fabrics specification
I/O Command Set Specific Admin Command	O
Vendor Specific	O
Notes: 1. O = Optional, M = Mandatory, P = Prohibited 2. For NVMe over PCIe implementations, the Keep Alive command is optional. For NVMe over Fabrics implementations, the associated NVMe Transport binding defines whether the Keep Alive command is optional or mandatory.	

<<...>>

Figure 420: I/O Controller – Log Page Support

Log Page Name	Command Support Requirements ¹
Error Information	M
SMART / Health Information (Controller scope)	M
SMART / Health Information (NVM subsystem scope)	O
Firmware Slot Information	M
Changed Namespace List	O
Commands Supported and Effects	O
Device Self-test	O
Telemetry Host-Initiated	O
Telemetry Controller-Initiated	O
Endurance Group Information	O
Predictable Latency Per NVM Set	O
Predictable Latency Event Aggregate	O
Asymmetric Namespace Access	O
Persistent Event	O
LBA Status Information	O
Endurance Group Event Aggregate	O
Command and Feature Lockdown	O
Notes:	
1. O = Optional, M = Mandatory, P = Prohibited	

7.2 Administrative Controller

[<<..>>](#)

Figure 426: Administrative Controller – Admin Command Support

Command	Command Support Requirements ¹
Delete I/O Submission Queue	P
Create I/O Submission Queue	P
Get Log Page	M
Delete I/O Completion Queue	P
Create I/O Completion Queue	P
Identify	M
Abort	O
Set Features	O ³
Get Features	O ³
Asynchronous Event Request	O ⁴
Namespace Management	O
Firmware Commit	O
Firmware Image Download	O
Device Self-test	O
Namespace Attachment	O
Keep Alive	NOTE 2
Directive Send	O
Directive Receive	O
Virtualization Management	O
NVMe-MI Send	O
NVMe-MI Receive	O
Doorbell Buffer Config	O

Figure 426: Administrative Controller – Admin Command Support

Command	Command Support Requirements ¹
Lockdown	O
NVMe over Fabrics Commands	Refer to the NVMe over Fabrics specification
I/O Command Set Specific Admin Command	O
Vendor Specific	O
Notes: 1. O = Optional, M = Mandatory, P = Prohibited 2. For NVMe over PCIe implementations, the Keep Alive command is optional. For NVMe over Fabrics implementations, the associated NVMe Transport binding defines whether the Keep Alive command is optional or mandatory. 3. Mandatory if any of the features in Figure 430 are implemented. 4. Mandatory if Telemetry Log, Firmware Commit, or SMART/Health Critical Warnings are supported.	

<<..>>

Figure 428: Administrative – Controller Log Page Support

Log Page Name	Command Support Requirements ¹
Error Information	M
SMART / Health Information (Controller scope)	O
SMART / Health Information (NVM subsystem scope)	O
Firmware Slot Information	O
Changed Namespace List	O
Commands Supported and Effects	M
Device Self-test	O
Telemetry Host-Initiated	O
Endurance Group Information	O
Predictable Latency Per NVM Set	O
Predictable Latency Event Aggregate	O
Asymmetric Namespace Access	P
Persistent Event	O
LBA Status Information	P
Endurance Group Event Aggregate	O
Command and Feature Lockdown	O
Notes: 1. O = Optional, M = Mandatory, P = Prohibited	

8.2 [Command and Feature Lockdown \(Optional\)](#)

The [Command and Feature Lockdown](#) feature is used to prohibit the execution of commands submitted to NVM Express controllers and/or Management Endpoints in an NVM subsystem. Within this feature, commands and Feature Identifiers are defined with the following scopes:

- [Admin Command Set](#) commands defined by the Opcode field;
- [Set Features](#) command features defined by the Feature Identifier field;
- [Management Interface Command Set](#) commands defined by the Opcode field (refer to the NVMe Management Interface Specification); and
- [PCIe Command Set](#) commands defined by the Opcode field (refer to the NVMe Management Interface Specification).

[Command Set](#) commands and [Feature Identifiers](#) are defined to be prohibitible by this feature, however it is vendor specific which of the Command Set commands and Feature Identifiers are prohibitible from execution, including the Lockdown command itself.

Once a command or Feature Identifier is prohibited, it is prohibited on all NVM Express controllers, if applicable, and all management endpoints, if applicable in the NVM subsystem. The applicability of prohibiting execution is based on the Interface field of the Lockdown command (refer to section 5.8).

The Lockdown command is used to specify commands that are prohibited from execution (i.e., locked down) and may be used further to then again allow that command to be executed.

The prohibiting of execution of a command as part of this feature shall persist until:

- a) power cycle of the NVM subsystem; or
- b) further being allowed by a follow-on Lockdown command.

If a prohibited Admin Command Set command or Feature Identifier is processed by a controller in the NVM subsystem, then that command shall be aborted with a status code of Command Prohibited by Command and Feature Lockdown.

If a prohibited Management Interface Command Set command or PCIe Command Set command is processed by a management endpoint in the NVM subsystem, then that command shall be aborted and send a Response Message with an Access Denied Error Response (refer to the NVMe Management Interface Specification).

The prohibition or allowance of the execution of commands using this feature treats from where the command was received, being an NVM Express controller Admin Submission Queue or an out-of-band Management Endpoint independently. This means that a command may be prohibited if received on an NVM Express controller Admin Submission Queue but allowed if received on an out-of-band Management Endpoint, if supported. The Interface field in the Lockdown command is used to specify this behavior.

A host may use the Command and Feature Lockdown log page (refer to section 5.14.1.4) to determine the commands and Feature Identifiers that are allowed to be prohibited from execution. A Get Log Page command specifying the Command and Feature Lockdown log returns a list of command opcodes or Feature Identifiers depending on the scope specified in the Get Log Page command. The returned list of opcodes or Feature Identifiers are the opcodes or Feature Identifiers that are:

- a) supported as being prohibitible from execution using the Lockdown command;
- b) currently prohibited from execution if received on an NVM Express controller submission queue; or
- c) currently prohibited from execution if received out-of-band on a Management Endpoint.

If the Command and Feature Lockdown feature is supported (i.e., bit TBD in the OACS field in Figure 247 is set to '1'), then the controller shall support the Lockdown command and the Command and Feature Lockdown log.

Description of Specification Changes to NVM Express Management Interface Revision 1.1

5 Management Interface Command Set

The Management Interface Command Set defines the Command Messages that may be submitted by a Requester when the NMIMT value is set to NVMe-MI Command. The Management Interface Command Set is applicable to both the out-of-band mechanism and the in-band tunneling mechanism. The processing of commands in the Management Interface Command Set may be affected by the Command and Feature Lockdown feature (refer to the NVM Express specification).

The NVMe-MI Message structure with all fields that are common to all NVMe-MI Messages are defined in section 3.1. The Response Message structure for the Management Interface Command Set is defined in section 4.1.2. The Message Body for the Management Interface Command Set is shown in Figure 54. Command specific fields for the Management Interface Command Set are defined in this section.

6 NVM Express Admin Command Set

[<<.. Add Lockdown command to Figure 110 ..>>](#)

Figure 110: List of NVMe Admin Commands Supported using the Out-of-Band Mechanism

Command	NVMe Storage Device O/M/P ¹	NVMe Enclosure O/M/P ¹
<<...>>	<<...>>	<<...>>
<u>Lockdown</u>	<u>0</u>	<u>0</u>
<<...>>	<<...>>	<<...>>

[<<.. Modify error response to Figure 26 ..>>](#)

Figure 26: Response Message Status Values

Value	Description	Error Response Format Section
<<...>>	<<...>>	<<...>>
07h	Access Denied: A Request Message was prohibited from being processed due to a vendor specific protection mechanism <u>or the Command and Feature Lockdown feature (refer to the NVM Express specification).</u>	4.1.2.1
<<...>>	<<...>>	<<...>>

7 PCIe Command Set (Optional)

The PCIe Command Set defines optional commands that a Management Controller may submit to access the memory, I/O, and configuration addresses spaces associated with a Controller in the NVM Subsystem. Only addresses mapped to the specified Controller may be accessed (e.g., these commands do not directly access memory on a host). The NMIMT field in the message header for PCIe Command Messages and Response Messages is set to 4h (PCIe Command). The PCIe Command Set is only applicable in the out-of-band mechanism and is prohibited in the in-band tunneling mechanism. [The processing of commands in the PCIe Command Set may be affected by the Command and Feature Lockdown feature \(refer to the NVM Express specification\).](#)

PCIe Commands over the out-of-band mechanism may interfere with host software. A Management Controller should coordinate with the host or issue only PCIe Commands that do not interfere with host software or in-band NVMe commands (e.g., PCIe Configuration Read). Coordination between a Management Controller and a host is outside the scope of this specification.

