



LEGAL NOTICE:

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

This NVM Express Base Specification 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 Base Specification 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 - 2021 NVM Express, Inc. ALL RIGHTS RESERVED.” When making any such citations or references to this document you are not permitted to revise, alter, modify, make any derivatives of, or otherwise amend the referenced portion of this document in any way without the prior express written permission of NVM Express, Inc. Nothing contained in this document shall be deemed as granting you any kind of license to implement or use this document or the specification described therein, or any of its contents, either expressly or impliedly, or to any intellectual property owned or controlled by NVM Express, Inc., including, without limitation, any trademarks of NVM Express, Inc.

LEGAL DISCLAIMER:

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

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

NVM Express Management Interface 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	4084 – Time-to-Ready Enhancements
Change Date	2021-4-08
Builds on Specification	NVMe Base Specification, Revision 1.4b NVMe over Fabrics Specification, Revision 1.1
References	TP 4052b Endurance Group Management TP 4082 Multiple Controller Shutdown Enhancements TP 4105a I/O Command Set Independent Identify Controller Data Structure TP 6026 Boot Partition Read Access via OOB

Technical Proposal Author(s)

Name	Company
Austin Bolen	Dell EMC
David Black	Dell EMC
Curtis Ballard	HPE

- This TP distinguishes two controller ready modes for attached namespaces during controller initialization:
 - Attached namespaces and media required to process Admin commands become ready by the time that the controller becomes ready (Controller Ready With Media mode)
 - Attached namespaces and media required to process Admin commands may become ready after the controller becomes ready (Controller Ready Independent of Media mode).
- Adds support bits and ready timeout values for both controller ready modes along with an indication of which mode is currently in use.
- Optional support for changing controller ready modes is specified.
- Specifies that CAP.TO shall be set to FFh if the actual timeout exceeds the 127.5 second limit of that field.
- Adds a status bit in I/O Command Set Independent Identify Namespace Data Structure to indicate if a namespace is ready or not.
- Adds a new error status code to indicate a namespace attached to the controller or media required to process an Admin command is not ready.
- Adds a Persistent Event Log entry to log when a namespace attached to a controller or media required to process an Admin command takes too long to become ready.

Revision History

Revision Date	Change Description
2020-07-08	• Initial version
2020-07-16	• Reworking based on workgroup feedback on 7/16.

2020-07-30	<ul style="list-style-type: none"> • Updates based on workgroup feedback on 7/30.
2020-08-10	<ul style="list-style-type: none"> • Added bit to Identify Namespace to indicate if the namespace is ready or not. • Add Persistent Event Log entry for when namespaces aren't ready in time.
2020-08-13	<ul style="list-style-type: none"> • Reworked based on workgroup feedback on 8/13: <ul style="list-style-type: none"> ○ Move namespace ready changes out of CAP.TO definition. ○ Rename new namespace ready bits and fields ○ Namespace Status change during shutdown does not generate an async event ○ Add new controller register (now named NRT – Namespace Ready Timeouts) to Fabrics spec as a property ○ Many additional editorial changes
2020-08-18	<ul style="list-style-type: none"> • Add new 7.6.1.1-2 subsections on namespace ready initialization modes and timing. This moves some content from earlier in the TP, including the definition of namespace ready and the potentially problematic modifications to CSTS.RDY. • Add text on logging error if namespaces are not ready in time. • Change from “active NSIDs” to “namespaces attached to the controller” • Make both timeout fields in new register 16 bits in size for consistency.
2020-08-20	<ul style="list-style-type: none"> • Reset change tracking to make doc readable again. • Make CAP.TO max value (FFh) usage backwards compatible (based on text from Christoph Hellwig). • Rename NRT register to RDYTO (Ready Timeouts) and rename Controller Ready Timeout field to Controller Only Ready Timeout (CORT) • Shorten “namespace ready behavior mode” to “namespace ready mode” • Remove Time to Namespace Ready field from Identify Namespace data structure. Async events generated when namespaces become ready suffice. • Added missing sentence to indicate that changing namespace ready mode does not affect the ready timeouts in the RDYTO register. • Explain that CORT field indicates the (likely shorter) time to controller ready if the controller is allowed to become ready before namespaces, even if controller is operating in the mode where they both have to be ready together. • Added requirement for critical warning (SMART / Health log) if namespaces becoming ready late is part of a larger NVM subsystem reliability problem.
2020-08-25	<ul style="list-style-type: none"> • Reset change tracking again. • Pick up changes from extensive review by Mike Allison • Consistently refer to namespaces attached to the controller, not active namespaces or active NSIDs. • Rename new register again to CRTO (Controller Ready Timeouts) • Extensive renaming of modes, fields, bits to make it clear that the Controller is the primary entity involved. Result should be more intuitive. • Delete resolved comments.
2020-08-27	<ul style="list-style-type: none"> • After reset, controller is required to be in mode where all namespaces are ready when the controller is ready. • Remove text about setting Critical Warning bit if namespace ready delay is indicative of more serious subsystem reliability problems. • Additional editorial changes, including picking up a few things missed in the 8-25 version.
2020-09-17	<ul style="list-style-type: none"> • Added Persistent Event Log additional data. • Added support for handling the situation where Admin commands aren't ready because the media is not ready. • Fixed all the spec versioning text. • Fixed how CRINE is set in the case where the CAP.CNRCAP field is not set to 11b. • Added a time limit on how long Namespace Not Ready can be returned with the Do Not Retry bit cleared to '0' after the controller is enabled. • Fixed/made consistent the names of a bunch of fields. • Clarified that the requirements for controller to become ready by CRWNT if there's an error with the namespace(s) that can't be resolved only applies in Controller Ready With Namespaces mode.

10-22-2020	<ul style="list-style-type: none"> • Renamed a bunch of fields based on workgroup feedback. • Removed redundant text in CAP.RDY. • Clarified that the Admin Command Media Not Ready status code only applies if the Controller Ready Independent of Media Enable bit is set to '1'. • Added a list of Admin commands and log pages permitted to return a status code of Admin Command Media Not Ready. • Specified that the controller shall not complete AERs if the log page needed to clear the AER is not ready. • Added additional info about the command that caused the error to Persistent Event Log. • Wordsmithing.
10-23-2020	<ul style="list-style-type: none"> • Updates based discussion in TWG on 10/22/2020. • Only log an error to the PEL when namespaces attached to the controller or media required to process Admin commands is not ready in time. Do not log a PEL entry each time a command is issued that results in Namespace Not Ready or Admin Command Media Not Ready error status. • Clarified that any commands that don't require namespace or media access shall be ready by the time the controller is ready in both modes. • Clarified that CAP.TO is dependent on the mode of the controller ready mode. • Specified the behavior for the cases where controller ready mode is modified during or after controller enable. • Addressed feedback from Mike Allison and Fred Knight. • Wordsmithing.
10-29-2020	<ul style="list-style-type: none"> • Corrected the CRIMS and CRIMT definitions. • Clarified CRIME definition. • Removed changes to Namespace Not Ready. • Moved NSTAT field to I/O Command Set Independent Identify Namespace data structure. • Added info to the Controller Ready Modes During Initialization section on how the timeouts work when the new controller ready timeout modes are not supported so that hosts have a consolidated set of guidelines for old and new devices. • Fixed incorrect definition of CAP.TO in Controller Ready Modes During Initialization section. • Merged changes to RDY bit with changes to that bit from TP 4082. • Fixed typos. • Wordsmithing and cleanup.
12-9-2020	<ul style="list-style-type: none"> • Corrected file name.
12-10-2020	<ul style="list-style-type: none"> • Editorial updates.
1-8-2021	<ul style="list-style-type: none"> • Added additional commands and log pages from TP 4052b and TP 6026 that are permitted to return Admin Command Media Not Ready. • Updated list of references.
1-22-2021	<ul style="list-style-type: none"> • Editorial updates based on member review feedback from Jim Hatfield. • Added Security Send and Security Receive to the list of commands permitted to return Admin Command Media Not Ready. • Added a bit to the Persistent Event Log to indicate the controller was not ready. • Clarified that if Admin Command Media Not Ready is returned for an Admin command tunneled out-of-band then a More Processing Required Response is also returned.
1-27-2021	<ul style="list-style-type: none"> • Removed clarification from previous revision that if Admin Command Media Not Ready is returned for an Admin command tunneled out-of-band then a More Processing Required Response is also returned. The NVMe-MI interactions will be handled in a separate TP.

1-28-2021	<ul style="list-style-type: none"> Accept all changes and mark comments resolved.
2-9-2021	<ul style="list-style-type: none"> Editorial updates based on member review feedback.
4-05-2021	<ul style="list-style-type: none"> Integrated into the NVMe Base Specification and the NVMe-over Fabrics Specification, Revision 1.1.
4-8-2021	<ul style="list-style-type: none"> Removed comments, accepted all changes, and converted references/cross-references to text.

Description for NVMe Base Specification Changes Document

- Feature Enhancements
 - Controller Ready Timeouts Enhancements
 - **New requirement / incompatible change** in section 3.1:
 - Adds new capability bits in the Controller Capabilities register that allow a controller to advertise to the host if namespaces attached to the controller and media required to process Admin commands are required to become ready by the time the controller becomes ready (the default) or after the controller becomes ready. (mandatory)
 - **New requirement / incompatible change** in section 3.1.1:
 - Specifies that CAP.TO shall be set to FFh if the actual timeout exceeds the 127.5 second limit of that field. (mandatory)
 - **New requirement / incompatible change** in section 3.1.5:
 - Adds a new read/write bit in the Controller Configuration register that allows a host to dynamically change controller ready modes. After a reset, namespaces attached to the controller and media required to process Admin commands become ready by the time the controller becomes ready. Changing this bit modifies the controller ready mode to allow namespaces attached to the controller and media required by Admin commands to become ready after the controller is ready. (optional)
 - Adds a new register (Controller Ready Timeouts) that advertises the worst-case times for the controller and attached namespaces to become ready. (mandatory)
 - **New requirement / incompatible change** in section 5.2:
 - Add a requirement that AERs shall not be completed if the logs needed to clear them are not ready to be read. (mandatory)
 - **New requirement / incompatible change** in section 5.2.1:
 - Add a requirement that the controller shall not send a Namespace Attribute Changed event for Namespace Status has changed and shutdown processing is either occurring or complete. (mandatory)
 - **New requirement / incompatible change** in section TBDa:
 - Adds a new status field to the I/O Command Set Independent Identify Namespace data structure to indicate whether the namespace is ready or not. (mandatory)
 - **New requirement / incompatible change** in section 4.6.1.2.1:
 - Adds a new error status code to indicate a controller, namespace attached to the controller, and/or media required to process an Admin command is not ready. (mandatory)
 - **New requirement / incompatible change** in section 5.14.1.13.1.5:
 - Adds a Persistent Event Log entry to log when a controller, namespace attached to a controller, or media required to process an Admin

command takes too long to become ready. (mandatory if the Persistent Event Log is supported)

- References
 - NVMe Base Specification, Revision 2.0 section [3.1](#), [3.1.1](#), [3.1.5](#), [4.6.1.2.1](#), [5.14.1.13.1.5](#), [5.2](#), [5.2.1](#), [7.6.1.1](#), [7.6.2.2](#), TP 4052b, TP 4082, TP 4105a, TP 6026

Description for NVMe over Fabrics Specification Changes Document

- Feature Enhancements
 - Controller Ready Timeouts Enhancements (mandatory)
 - **New requirement / incompatible** change in section [3.6.1](#):
 - Adds a new property (Controller Ready Timeouts) that advertises the worst-case times for the controller and attached namespaces to become ready.
 - References:
 - NVMe over Fabrics Specification, Revision 1.1 section [3.6.1](#).

Description of Specification Changes to NVMe Base Specification

Markup Conventions:

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

Add a new row to Figure 68 (Register Definition) as follows:

Start	End	Symbol	Description
...
68h	6Bh	CRTO	Controller Ready Timeouts
...

Modify a portion of Figure 69 (Offset 0h: CAP – Controller Capabilities) as follows:

Bits	Type	Reset	Description
31:24	RO	Impl Spec	<p>Timeout (TO): This is the worst-case time that host software shall <u>should</u> wait for CSTS.RDY to transition from:</p> <ul style="list-style-type: none"> a) '0' to '1' after CC.EN transitions from '0' to '1'; or b) '1' to '0' after CC.EN transitions from '1' to '0'. <p>This worst-case time may be experienced after events such as an abrupt shutdown or activation of a new firmware image; typical times are expected to be much shorter.</p> <p>This field is in 500 millisecond units. <u>The maximum value of this field is FFh, which indicates a 127.5 second timeout.</u></p> <p><u>If the Controller Ready Independent of Media Enable (CC.CRIME) bit is cleared to '0' and the worst-case time for CSTS.RDY to change state is due to enabling the controller after CC.EN transitions from '0' to '1', then this field shall be set to:</u></p> <ul style="list-style-type: none"> a) <u>the value in Controller Ready With Media Timeout (CRTO.CRWMT); or</u> b) <u>FFh if CRTO.CRWMT is greater than FFh.</u> <p><u>If the Controller Ready Independent of Media Enable (CC.CRIME) bit is set to '1' and the worst-case time for CSTS.RDY to change state is due to enabling the controller after CC.EN transitions from '0' to '1', then this field shall be set to:</u></p> <ul style="list-style-type: none"> a) <u>the value in Controller Ready Independent of Media Timeout (CRTO.CRIMT); or</u> b) <u>0FFh if CRTO.CRIMT is greater than FFh.</u> <p><u>Controllers that support the CRTO register (refer to Figure TBD4) are able to indicate larger timeouts for enabling the controller. Host software should use the value in CRTO.CRWMT or CRTO.CRIMT depending on the controller ready mode indicated by CC.CRIME to determine the worst-case timeout for CSTS.RDY to transition from '0' to '1' after CC.EN transitions from '0' to '1'. Host software that is based on versions 1.4 or earlier of this specification is not required to wait for more than 127.5 seconds for CSTS.RDY to transition.</u></p> <p><u>Refer to sections 7.6.1.1 and 7.6.1.2 for more information.</u></p>
...

Bits	Type	Reset	Description						
60:59	RO	Impl Spec	Controller Ready Modes Supported (CRMS) : This field indicates the ready capabilities of the controller. Refer to sections 7.6.1.1 and 7.6.1.2 for more detail.						
			<table><tr><th>Bits</th><th>Description</th></tr><tr><td>60</td><td>Controller Ready Independent of Media Support (CRIMS): If this bit is set to '1', then the controller supports the Controller Ready Independent of Media mode. If this bit is cleared to '0', then the controller does not support Controller Ready Independent of Media mode.</td></tr><tr><td>59</td><td>Controller Ready With Media Support (CRWMS): If this bit is set to '1', then the controller supports the Controller Ready With Media mode. If this bit is cleared to '0', then the controller does not support Controller Ready With Media mode. This bit shall be set to '1' on controllers compliant with versions later than version 1.4 of this specification.</td></tr></table>	Bits	Description	60	Controller Ready Independent of Media Support (CRIMS) : If this bit is set to '1', then the controller supports the Controller Ready Independent of Media mode. If this bit is cleared to '0', then the controller does not support Controller Ready Independent of Media mode.	59	Controller Ready With Media Support (CRWMS) : If this bit is set to '1', then the controller supports the Controller Ready With Media mode. If this bit is cleared to '0', then the controller does not support Controller Ready With Media mode. This bit shall be set to '1' on controllers compliant with versions later than version 1.4 of this specification.
			Bits	Description					
60	Controller Ready Independent of Media Support (CRIMS) : If this bit is set to '1', then the controller supports the Controller Ready Independent of Media mode. If this bit is cleared to '0', then the controller does not support Controller Ready Independent of Media mode.								
59	Controller Ready With Media Support (CRWMS) : If this bit is set to '1', then the controller supports the Controller Ready With Media mode. If this bit is cleared to '0', then the controller does not support Controller Ready With Media mode. This bit shall be set to '1' on controllers compliant with versions later than version 1.4 of this specification.								

Add a new field to Figure 78 (Offset 14h: CC – Controller Configuration) as follows:

Bits	Type	Reset	Description						
...						
24	RW/RO	0b	<p><u>Controller Ready Independent of Media Enable (CRIME):</u> This field controls the controller ready mode. The controller ready mode is determined by the state of this bit at the time the controller is enabled by transitioning the CC.EN bit from '0' to '1'.</p> <p>If the CAP.CRMS field is set to 11b, then this bit is RW. If the CAP.CRMS field is not set to 11b, then this bit is RO and shall be cleared to '0'. Refer to sections 7.6.1.1 and 7.6.1.2 for more detail.</p> <p><u>Changing the value of this field may cause a change in the time reported in the CAP.TO field. Refer to the definition of CAP.TO for more details.</u></p> <table><tr><th>Value</th><th>Definition</th></tr><tr><td>0b</td><td>Enabling the controller (i.e., CC.EN transitions from '0' to '1') when this bit is cleared to '0' enables Controller Ready With Media mode.</td></tr><tr><td>1b</td><td>Enabling the controller when this bit is set to '1' enables Controller Ready Independent of Media mode.</td></tr></table>	Value	Definition	0b	Enabling the controller (i.e., CC.EN transitions from '0' to '1') when this bit is cleared to '0' enables Controller Ready With Media mode.	1b	Enabling the controller when this bit is set to '1' enables Controller Ready Independent of Media mode.
Value	Definition								
0b	Enabling the controller (i.e., CC.EN transitions from '0' to '1') when this bit is cleared to '0' enables Controller Ready With Media mode.								
1b	Enabling the controller when this bit is set to '1' enables Controller Ready Independent of Media mode.								

Add new section under section 3.1 as follows:

Section 3.1. **TBD14** Offset **68h**: CRTO – Controller Ready Timeouts

This register indicates the controller ready timeout values. This register is mandatory for controllers compliant with versions later than 1.4 of this specification.

Figure **TBD4: Offset **68h**: CRTO – Controller Ready Timeouts**

Bits	Type	Reset	Description
31:16	RO	Impl Spec	<p>Controller Ready Independent of Media Timeout (CRIMT): If the CAP.CRMS.CRIMS bit is cleared to '0', then this field is not applicable and shall be cleared to 0h.</p> <p>If the CAP.CRMS.CRIMS bit is set to '1', then this field contains the worst-case time that host software should wait after CC.EN transitions from '0' to '1' for the controller to become ready and be able to successfully process all commands that do not access attached namespaces and Admin commands that do not require access to media when the controller is in Controller Ready Independent of Media mode (i.e., the CC.CRIME bit is set to '1'). Attached namespaces and media required to process Admin commands may or may not be ready within this time period (refer to section 7.6.1.1, section 7.6.1.2, and Figure TBD11).</p> <p>This worst-case time may be experienced after events such as an abrupt shutdown or activation of a new firmware image; typical times are expected to be much shorter. This field is in 500 millisecond units.</p> <p>The value of this field should not exceed FFh (i.e., 127.5 seconds).</p>
15:0	RO	Impl Spec	<p>Controller Ready With Media Timeout (CRWMT): This field contains the worst-case time that host software should wait after CC.EN transitions from '0' to '1' for:</p> <ul style="list-style-type: none"> a) the controller to become ready and be able to successfully process all commands; and b) for all attached namespaces and media required to process Admin commands to become ready, <p>independent of which ready mode (refer to CC.CRIME) the controller is in (refer to section 7.6.1.1 and section 7.6.1.2).</p> <p>This worst-case time may be experienced after events such as an abrupt shutdown or activation of a new firmware image; typical times are expected to be much shorter. This field is in 500 millisecond units.</p> <p>The value of this field shall be greater than or equal to the value of the CRTO.CRIMT field and may be significantly larger than the value of the CRTO.CRIMT field.</p>

Add the following to Figure 128 (Status Code – Generic Command Status Values):

Value	Description
<p>24h</p> <p><Note to the editor: Select a value in the range 0h to 7Fh></p>	<p>Admin Command Media Not Ready: The Admin command requires access to media and the media is not ready. The Do Not Retry bit indicates whether re-issuing the command at a later time may succeed. This status code shall only be returned:</p> <ul style="list-style-type: none"> a) for Admin commands; and b) if the controller is in Controller Ready Independent of Media mode (CC.CRIME is bit is set to '1'). <p>This status code shall not be returned with the Do Not Retry bit cleared to '0' after the amount of time indicated by the Controller Ready With Media Timeout field (CRO.CRWMT) after the controller is enabled (i.e., CC.EN transitions from '0' to '1').</p> <p>Refer to Figure TBD11 for the list of Admin commands permitted to return this status code.</p>

Modify a portion of Section 5.2 (Asynchronous Event Request command) as follows:

Asynchronous events are grouped into event types. The event type information is indicated in Dword 0 of the completion queue entry for the Asynchronous Event Request command. When the controller posts a completion queue entry for an outstanding Asynchronous Event Request command and thus reports an asynchronous event, subsequent events of that event type are automatically masked by the controller until the host clears that event. An event is cleared by reading the log page associated with that event using the Get Log Page command (refer to section 5.14). If that log page is not accessible because media is not ready (i.e., the controller returns a status code of Admin Command Media Not Ready for the Get Log Page command requesting the log page), then the controller shall not post a completion queue entry for that asynchronous event until the controller is able to successfully return the log page that is required to be read to clear the asynchronous event.

Modify Figure 149 (Asynchronous Event Information – Notice) as follows:

Figure 149: Asynchronous Event Information – Notice

Value	Description
00h	<p>Namespace Attribute Changed: Indicates a change to one or both of the following:</p> <ul style="list-style-type: none"> the Identify Namespace data structure (refer to Figure 249) for one or more namespaces; or the Namespace List returned when the Identify command is issued with the CNS field set to 02h. <p>To clear this event, host software issues a Get Log Page command for the Changed Namespace List log page (Log Identifier 04h - refer to section 5.14.1.4) with the Retain Asynchronous Event bit cleared to '0'.</p> <p>A controller shall not send this event if:</p> <ol style="list-style-type: none"> Namespace Utilization (refer to Figure 249) has changed, as this is a frequent event that does not require action by the host; the ANAGRPID field (refer to Figure 249) has changed; or capacity information (i.e., the NUSE field and the NVMCAP field) returned in the Identify Namespace data structure (refer to Figure 249) changed as a result of an ANA state change; or <u>Namespace Status (refer to Figure TBDa <note to editor: Figure TBDa is part of TP 4105 and is called Identify – I/O Command Set Independent Identify Namespace Data Structure>) has changed and shutdown processing is either occurring (i.e., CSTS.SHST is set to 01b) or complete (i.e., CSTS.SHST is set to 10b).</u> <p>A controller shall only send this event for changes to the Format Progress Indicator field when bits 6:0 of that field transition from a non-zero value to 0h, or from 0h to a non-zero value.</p>

Add the following to Figure 224 (NVM Subsystem Hardware Error Event Codes):

Code	Description
<u>0Bh</u>	<p><u>Controller Ready Timeout Exceeded: Indicates that:</u></p> <ol style="list-style-type: none"> <u>the controller was not ready to process at least one command without error as described in section 7.6.1.2.1 within up to and including the amount of time indicated by:</u> <ul style="list-style-type: none"> <u>the Controller Ready With Media Timeout field (CRTO.CRWMT) in Controller Ready With Media mode (CC.CRIME is cleared to '0'); or</u> <u>the Controller Ready Independent of Media Timeout field (CRTO.CRIMT) in Controller Ready Independent of Media mode (CC.CRIME is set to '1'); and/or</u> <u>at least one namespace attached to the controller or media required to process at least one Admin command was not ready within up to and including the amount of time indicated by the Controller Ready With Media Timeout field (CRTO.CRWMT) after the controller was enabled by transitioning CC.EN from '0' to '1'.</u> <p><u>Refer to Figure TBD10 for the format of the Additional Hardware Error Information field.</u></p>

Add the following new table after Table 225 (Additional Hardware Error Information for correctable and uncorrectable PCIe errors):

Figure TBD10: Additional Hardware Error Information for Controller Ready Timeout Exceeded errors

Bytes	Value												
0	Controller State: Indicates the state of the controller at the time the Controller Ready Timeout Exceeded error occurred.												
	<table><tr><th>Bits</th><th>Description</th></tr><tr><td>7:3</td><td>Reserved</td></tr><tr><td>3</td><td>Controller Not Ready: Indicates the controller was not ready to process at least one command without error as described in section 7.6.1.2.1 within up to and including the amount of time indicated by:<ul style="list-style-type: none">a) the Controller Ready With Media Timeout field (CRTO.CRWMT) in Controller Ready With Media mode (CC.CRIME is cleared to '0'); orb) the Controller Ready Independent of Media Timeout field (CRTO.CRIMT) in Controller Ready Independent of Media mode (CC.CRIME is set to '1').</td></tr><tr><td>2</td><td>Admin Command Media Not Ready: Indicates media required to process at least one Admin command was not ready within up to and including the amount of time indicated by the Controller Ready With Media Timeout field (CRTO.CRWMT) after the controller was enabled by transitioning CC.EN from '0' to '1'.</td></tr><tr><td>1</td><td>Namespace Not Ready: Indicates at least one namespace attached to the controller was not ready within up to and including the amount of time indicated by the Controller Ready With Media Timeout field (CRTO.CRWMT) after the controller was enabled by transitioning CC.EN from '0' to '1'.</td></tr><tr><td>0</td><td>Controller Ready Independent of Media Enable: Indicates the value of the CC.CRIME bit when the Controller Ready Timeout Exceeded error occurred.</td></tr></table>	Bits	Description	7:3	Reserved	3	Controller Not Ready: Indicates the controller was not ready to process at least one command without error as described in section 7.6.1.2.1 within up to and including the amount of time indicated by: <ul style="list-style-type: none">a) the Controller Ready With Media Timeout field (CRTO.CRWMT) in Controller Ready With Media mode (CC.CRIME is cleared to '0'); orb) the Controller Ready Independent of Media Timeout field (CRTO.CRIMT) in Controller Ready Independent of Media mode (CC.CRIME is set to '1').	2	Admin Command Media Not Ready: Indicates media required to process at least one Admin command was not ready within up to and including the amount of time indicated by the Controller Ready With Media Timeout field (CRTO.CRWMT) after the controller was enabled by transitioning CC.EN from '0' to '1'.	1	Namespace Not Ready: Indicates at least one namespace attached to the controller was not ready within up to and including the amount of time indicated by the Controller Ready With Media Timeout field (CRTO.CRWMT) after the controller was enabled by transitioning CC.EN from '0' to '1'.	0	Controller Ready Independent of Media Enable: Indicates the value of the CC.CRIME bit when the Controller Ready Timeout Exceeded error occurred.
	Bits	Description											
	7:3	Reserved											
	3	Controller Not Ready: Indicates the controller was not ready to process at least one command without error as described in section 7.6.1.2.1 within up to and including the amount of time indicated by: <ul style="list-style-type: none">a) the Controller Ready With Media Timeout field (CRTO.CRWMT) in Controller Ready With Media mode (CC.CRIME is cleared to '0'); orb) the Controller Ready Independent of Media Timeout field (CRTO.CRIMT) in Controller Ready Independent of Media mode (CC.CRIME is set to '1').											
	2	Admin Command Media Not Ready: Indicates media required to process at least one Admin command was not ready within up to and including the amount of time indicated by the Controller Ready With Media Timeout field (CRTO.CRWMT) after the controller was enabled by transitioning CC.EN from '0' to '1'.											
	1	Namespace Not Ready: Indicates at least one namespace attached to the controller was not ready within up to and including the amount of time indicated by the Controller Ready With Media Timeout field (CRTO.CRWMT) after the controller was enabled by transitioning CC.EN from '0' to '1'.											
0	Controller Ready Independent of Media Enable: Indicates the value of the CC.CRIME bit when the Controller Ready Timeout Exceeded error occurred.												
3:1	Reserved												

Insert new Sections 7.6.1.1 and 7.6.1.2 after Section 7.6.1, renumbering the existing Section 7.6.1.1 to 7.6.1.3

7.6.1.1 Initialization: Controller Ready Modes During Initialization

There are two possible controller ready modes:

- **Controller Ready With Media:** By the time the controller becomes ready (i.e., by the time that CSTS.RDY transitions from '0' to '1') after the controller is enabled (i.e., CC.EN transitions from '0' to '1'), then:
 - a) the controller shall be able to process all commands without error as described in section 7.6.1.2.1; and
 - b) all namespaces attached to the controller and all media required to process Admin commands shall be ready (i.e., commands are not permitted to be aborted with a status code of Namespace Not Ready with the Do Not Retry bit cleared to '0' or Admin Command Media Not Ready with the Do Not Retry bit cleared to '0').

- **Controller Ready Independent of Media:** After the controller is enabled, all namespaces attached to the controller and media required to process Admin commands may or may not become ready by the time the controller becomes ready. Any NVM command that specifies one or more namespaces attached to the controller is permitted to be aborted with a status code of Namespace Not Ready with the Do Not Retry bit cleared to '0' for only up to and including CRYPTO.CRWM amount of time after the controller is enabled.

Admin commands that require access to the media are permitted to be aborted with a status code of Admin Command Media Not Ready with the Do Not Retry bit cleared to '0' for only up to and including CRYPTO.CRWM amount of time after the controller is enabled. Refer to Figure TBD11 for a list of Admin commands that are permitted to be aborted with a status code of Admin Command Media Not Ready.

The controller shall be able to process without error as described in section 7.6.1.2.1:

- a) all Admin commands not listed in Figure TBD11 by the time the controller is ready;
- b) all Admin commands listed in Figure TBD11 within up to and including CRYPTO.CRWM amount of time after the controller is enabled; and
- c) all NVM commands within up to and including CRYPTO.CRWM amount of time after the controller is enabled.

Figure TBD11: Admin Commands Permitted to Return a Status Code of Admin Command Media Not Ready

Admin Command	Additional Restrictions
Capacity Management <note to editor: The Capacity Management Command is defined in TP 4052b>	
Device Self-test	If the Device Self-Test would result in testing one or more namespaces, then returning a Status Code of Admin Command Media Not Ready is permitted. If the Device Self-Test would not result in testing any namespaces, then returning a Status Code of Admin Command Media Not Ready is not permitted.
Firmware Commit	
Firmware Image Download	
Get LBA Status	

Figure TBD11: Admin Commands Permitted to Return a Status Code of Admin Command Media Not Ready

<u>Admin Command</u>	<u>Additional Restrictions</u>
<u>Get Log Page</u>	<p>Get Log Page is only permitted to return a status code of Admin Command Media Not Ready for the following log pages:</p> <p><note to editor: In the list below, Media Unit Status and Supported Capacity Configuration List are defined in TP 4052b and Boot Partition is defined in TP 6026></p> <ul style="list-style-type: none"> • <u>Device Self-test</u> • <u>Firmware Slot Information</u> • <u>Telemetry Controller-Initiated</u> • <u>Telemetry Host-Initiated</u> • <u>Predictable Latency Per NVM Set</u> • <u>Predictable Latency Event Aggregate</u> • <u>Persistent Event Log</u> • <u>LBA Status Information</u> • <u>Endurance Group Event Aggregate</u> • <u>Media Unit Status</u> • <u>Supported Capacity Configuration List</u> • <u>Boot Partition</u> • <u>Reservation Notification</u> • <u>Vendor Specific</u>
<u>Namespace Attachment</u>	
<u>Namespace Management</u>	
<u>Format NVM</u>	
<u>Sanitize</u>	
<u>Security Receive</u> ¹	
<u>Security Send</u> ¹	
<u>Vendor Specific</u>	
<p>NOTES:</p> <p>1. A host may require discovery operations performed via Security Send/Receive (e.g., TCG Level 0 Discovery) to be processed prior to media being ready and therefore it is recommended that controllers not return Admin Command Media Not Ready for such discovery operations.</p>	

The Controller Ready Modes Supported (CAP.CRMS) field (refer to Figure 69) indicates which controller ready modes are supported. The CAP.CRMS field consists of two bits:

- the Controller Ready With Media Support (CAP.CRMS.CRWMS) bit; and
- the Controller Ready Independent of Media Support (CAP.CRMS.CRIMS) bit.

Controllers are required to set the CAP.CRMS.CRWMS bit to '1' (i.e., to set the CAP.CRMS field to 01b or 11b). The CAP.CRMS.CRWMS bit was not defined in version 1.4 and earlier of this specification. Controllers compliant with version 1.4 and earlier of this specification may have the CAP.CRMS.CRWMS field cleared to 00b.

The Controller Ready Independent of Media Enable (CC.CRIME) bit (refer to Figure 78) controls the controller ready mode based on the value of the CAP.CRMS field as follows:

- If the CAP.CRMS field is cleared to 00b, the controller ready mode is not able to be selected. In this case, the read-only CC.CRIME bit shall be cleared to '0' and should be ignored by host software;
- If the CAP.CRMS field is set to 01b (i.e., the CAP.CRMS.CRIMS bit is cleared to '0' and the CAP.CRMS.CRWMS bit is set to '1'), then the controller is in Controller Ready With Media mode and the read-only CC.CRIME bit shall be cleared to '0'; and

- c) If the CAP.CRMS field is set to 11b, then both controller ready modes are supported, and the host may select the controller ready mode by modifying the value of the CC.CRIME bit. In this situation, the host should set the controller ready mode by writing to the CC.CRIME bit before the controller is enabled (e.g., as part of step 4 of the initialization sequence of actions described in Section 7.6.1).

7.6.1.2 Initialization: Controller Ready Timeouts During Initialization

The CAP.CRMS field was not defined in version 1.4 and earlier of this specification. Controllers compliant with version 1.4 and earlier of this specification may clear the CAP.CRMS field to 00b'. This section is applicable to controllers that clear the CAP.CRMS field to 00b and controllers that set CAP.CRMS to a non-zero value.

There are three controller ready timeout fields:

1. CAP.TO (refer to Figure 69);
2. CRYPTO.CRWMT (refer to Figure TBD4); and
3. CRYPTO.CRIMT (refer to Figure TBD4).

The details regarding these timeouts during controller initialization are as follows:

- a) The CAP.TO field shall be set as described in Figure 69.
- b) If the CAP.CRMS field is cleared to 00b', then:
 - A. the Controller Ready Independent of Media Timeout (CRYPTO.CRIMT) field is reserved;
 - B. the Controller Ready Independent of Media Timeout (CRYPTO.CRWMT) field is reserved;
and
 - C. the worst-case time the host should wait after the controller is enabled (i.e., CC.EN transitions from '0' to '1') for the controller to become ready (CSTS.RDY transitions from '0' to '1') is indicated by CAP.TO.
- c) If the controller is in Controller Ready With Media mode (i.e., the CC.CRIME bit is cleared to '0'), then:
 - A. the Controller Ready Independent of Media Timeout (CRYPTO.CRIMT) field is not applicable;
and
 - B. the Controller Ready With Media Timeout (CRYPTO.CRWMT) indicates the worst-case time the host should wait after the controller is enabled for:
 - the controller to become ready and be able to successfully process all commands;
and
 - all attached namespaces and media required to process Admin commands to become ready.
and;
- d) If the controller is in Controller Ready Independent of Media mode (i.e., the CC.CRIME bit is set to '1'), then
 - A. the Controller Ready With Media Timeout (CRYPTO.CRWMT) field indicates the worst-case time that host software should wait for all attached namespaces and media required to process Admin commands to become ready after the controller is enabled; and
 - B. the Controller Ready Independent of Media Timeout (CRYPTO.CRIMT) field indicates the worst-case time the host should wait after the controller is enabled for the controller to become ready and be able to successfully process:
 - all commands that do not access attached namespaces; and
 - Admin commands that do not require access to media.

Changes to the value of the CC.CRIME bit shall have no effect on the values of the CRYPTO.CRWMT and CRYPTO.CRIMT fields. Changes to the value of the CC.CRIME bit may have an effect on the value of the CAP.TO field (refer to the definition of CAP.TO in Figure 69 for more details).

7.6.1.2.1 Initialization: Handling Errors During Initialization

If the CAP.CRMS field is non-zero and the controller has been enabled by transitioning CC.EN from '0' to '1' and the controller encounters a failure that prevents:

- a) at least one command from being able to be processed without error;
- b) at least one namespace attached to the controller from becoming ready within up to and including CRTO.CRWMT amount of time after the controller was enabled; or
- c) media required by at least one Admin command from becoming ready within up to and including CRTO.CRWMT amount of time after the controller was enabled.

then:

- a) if the controller has not become ready, then the controller shall become ready (i.e., set CSTS.RDY to '1') within up to and including CRTO.CRWMT amount of time after the controller was enabled; and
- b) if the Persistent Event Log is supported, then the controller shall record an NVM Subsystem Hardware Error Event with the NVM Subsystem Hardware Error Event code set to a value of Controller Ready Timeout Exceeded in the Persistent Event Log (refer to Figure 224).

Change section 7.6.1.1 (Software Progress Marker) to 7.6.1.3 as follows:

7.6.1.43 Software Progress Marker

Modify a portion of section 8.15 (8.15 Sanitize Operations) as follows:

The Sanitize Status log page shall be updated as described:

- Initialize before any controller in the NVM subsystem is ready as described in sections 7.6.1.1 and 7.6.1.2;
- Update before a Sanitize command that starts a sanitize operation is completed (i.e., prior to the completion queue entry being posted for the Sanitize command); and
- Update when a sanitize operation is complete (e.g., immediately prior to the completion queue entry being posted for the Sanitize Operation Completed asynchronous event or for the Sanitize Operation Completed With Unexpected Deallocation asynchronous event).

Description of Specification Changes to NVMe over Fabrics Specification

Add the following to Figure 29 (Property Definitions):

Figure 29: Property Definitions

<u>68h</u>	<u>6Bh</u>	<u>CRTO</u>	<u>Controller Ready Timeouts</u>
------------	------------	-------------	----------------------------------

Description of Specification Changes to TP 4105a

Add a new field to Figure TBDa (Identify – I/O Command Set Independent Identify Namespace Data Structure) as follows:

Bytes	O/M ¹	Description						
14	M	<u>Namespace Status (NSTAT): This field indicates the status of the namespace with the specified NSID.</u>						
		<table><tr><th>Bits</th><th>Description</th></tr><tr><td>7:1</td><td>Reserved</td></tr><tr><td>0</td><td><u>Namespace Ready (NRDY): A value of '1' indicates that the namespace is ready (refer to section 7.6.1.1). A value of '0' indicates that the namespace is not ready.</u></td></tr></table>	Bits	Description	7:1	Reserved	0	<u>Namespace Ready (NRDY): A value of '1' indicates that the namespace is ready (refer to section 7.6.1.1). A value of '0' indicates that the namespace is not ready.</u>
		Bits	Description					
		7:1	Reserved					
0	<u>Namespace Ready (NRDY): A value of '1' indicates that the namespace is ready (refer to section 7.6.1.1). A value of '0' indicates that the namespace is not ready.</u>							

Description of Specification Changes to TP 4082

Modify a portion of Figure 79 (Offset 1Ch: CSTS – Controller Status) as follows:

00	RO	0b	<p>Ready (RDY): This bit is set to '1' when the controller is ready to accept Submission Queue Tail doorbell writes after CC.EN is set to '1'. This bit shall be cleared to '0' when CC.EN is cleared to '0' once the controller is ready to be re-enabled. Commands should shall not be submitted to the controller until this bit is set to '1' after the CC.EN bit is set to '1'. Failure to follow this recommendation requirement produces undefined results. <u>Refer to the definition of CAP.TO and sections 7.6.1.1 and 7.6.1.2 for timing information related to this field.</u></p> <p>If shutdown processing is not occurring on the controller (CSTS.SHST is set to 00b), then the Hhost software shall wait a minimum of CAP.TO seconds for this bit to be set to '1' after setting CC.EN to '1' from a previous value of '0'.</p> <p>If a controller shutdown processing is occurring on the controller (i.e., CSTS.ST is cleared to '0' and CSTS.SHST is set to 01b), then the host software shall wait a minimum of RTD3 Entry Latency microseconds (refer to the Identify Controller data structure) plus CAP.TO seconds for this bit to be set to '1' after setting CC.EN to '1' from a previous value of '0'.</p> <p>If an NVM Subsystem Shutdown has completed that affects this controller (i.e., CSTS.ST is set to '1' and CSTS.SHST is set to 10b), then an NVM Subsystem Reset is be required before this bit is allowed to be set to '1'. Refer to section <u>7.6.New</u>.</p>
----	----	----	--