



LEGAL NOTICE:

© **Copyright 2008 to 2024 NVM Express®**, Inc. **ALL RIGHTS RESERVED.**

This Technical Proposal is proprietary to the NVM Express, Inc. (also referred to as "Company") and/or its successors and assigns.

NOTICE TO USERS WHO ARE NVM EXPRESS, INC. MEMBERS: Members of NVM Express, Inc. have the right to use and implement this Technical Proposal subject, however, to the Member's continued compliance with the Company's Intellectual Property Policy and Bylaws and the Member's Participation Agreement.

NOTICE TO NON-MEMBERS OF NVM EXPRESS, INC.: If you are not a Member of NVM Express, Inc. and you have obtained a copy of this document, you only have a right to review this document or make reference to or cite this document. Any such references or citations to this document must acknowledge NVM Express, Inc. copyright ownership of this document. The proper copyright citation or reference is as follows: "© 2008 to 2024 NVM Express, Inc. ALL RIGHTS RESERVED." When making any such citations or references to this document you are not permitted to revise, alter, modify, make any derivatives of, or otherwise amend the referenced portion of this document in any way without the prior express written permission of NVM Express, Inc. Nothing contained in this document shall be deemed as granting you any kind of license to implement or use this document or the specification described therein, or any of its contents, either expressly or impliedly, or to any intellectual property owned or controlled by NVM Express, Inc., including, without limitation, any trademarks of NVM Express, Inc.

LEGAL DISCLAIMER:

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

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

The NVM Express® design mark is a registered trademark of NVM Express, Inc.

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

NVM Express® Technical Proposal for New Feature

Technical Proposal ID	TP4169 Shutdown Clarifications and Enhancements
Change Date	2024-07-19
Impacted Specifications	NVM Express Base Specification 2.0d NVM Express Management Interface Specification 1.2d TP6032 Out-of-Band Admin Command While Shutdown Enhancement
References	TP6032 Out-of-Band Admin Command While Shutdown Enhancement TP4082 Multiple Controller Shutdown Enhancements ECN111

Technical Proposal Author(s)

Name	Company
Judy Brock, Michael Allison	Samsung
Austin Bolen, David Black	Dell

Update the prior definition of NVM Subsystem Shutdown to correctly specify interactions between NVM Subsystem Shutdown and Controller Level Resets (CLRs) that are not caused by NVM Subsystem Reset.

Add field to indicate how long a host should wait for NVM Subsystem Shutdown to complete.

Improve specification of interactions between all forms of Controller Level Reset (e.g., Controller Reset) and both forms of shutdown (controller shutdown and NVM Subsystem Shutdown).

Remove use of the word “safe” to describe ready-to-be-powered-off conditions.

Specify NVMe-MI Ignore Shutdown functionality (TP6032) for the situations in which the media is in the shutdown state while CSTS.SHST is cleared to 00b.

Revision History

Revision Date	Change Description
2024-04-02 (J Brock)	<p>Initial version –</p> <ul style="list-style-type: none"> Changed reset behavior to preserve CSTS.ST and CSTS.SHST after both Controller Shutdown and after NVM Subsystem Shutdown Allow transition back to unsafe-to-power-off after Shutdown followed by CLR (except if caused by Controller Reset (CC.EN -> 0)). Unsafe to power off means media initialization has begun In CSTS.ST field, CSTS.SHST field, and PCIe Transport binding spec, clarified that transport-specific resets (e.g., PCIe Conventional Reset and FLR) have no effect on NVM Subsystem Shutdown that is in progress or has completed
2024-04-09 (D Black)	<ul style="list-style-type: none"> Remove ability for CSTS.SHST to survive Controller Level Reset in CC.EN text for subsequent editing, clean up CSTS.SHST description. Another alternative revision of the “safe to power off” text in section 3.6.1
2024-04-11 (D Black)	<ul style="list-style-type: none"> Reset change tracking Minor edits to resolve additional comments, remove resolved comments Remove ability of NVM Subsystem Shutdown to override PCIe resets. Consolidate section 3.6.1 changes.
2024-04-12 (D Black)	<ul style="list-style-type: none"> Reset change tracking (again) Remove NVM Subsystem Reset override of PCIe resets. Remove PCIe spec changes in favor of Base spec language to explicitly permit media initialization after the relevant PCIe resets.
2024-04-19 (D Black)	<ul style="list-style-type: none"> Reset change tracking (again) Initial changes to preserve both CSTS.ST and CSTS.SHST across CLR's other than NVM Subsystem Reset. Add portion of CAP figure, for later addition of incompatible change bit/field. Minor updates to power-off changes in 3.6.1, and add sentence about unsafe power off to CC.EN description.
2024-04-19 (J Brock)	<ul style="list-style-type: none"> Added text to allow controllers in the NVM subsystem or domain to abort any command with a status code of Commands Aborted due to Power Loss Notification while NVM Subsystem Shutdown is in progress
2024-04-22 (D Black)	<ul style="list-style-type: none"> Revise 3.6.1 power-off changes to remove use of “shall”. Minor edits to CSTS.ST-related text on resuming command processing in CSTS.SHST description (from Judy). Align text on resuming command processing across CSTS.SHST description, and section 3.6.1. Initial ToDo list to allow Controller Reset during NVM Subsystem Shutdown
2024-04-23 (D Black)	<ul style="list-style-type: none"> Revise 3.6.1 power-off changes to use state transition language. Add incomplete shutdown as CAP.TO example for max timeout Add text on safe power removal and how to start executing commands to that both chunks of text are present in all three shutdown subsections.
2024-04-24 (D Black)	<ul style="list-style-type: none"> Add ToDo item for NVM Subsystem Shutdown Timeout Add material from TP6032 that will require modification. Correct NVMe-MI Admin text to require media access and to assume that media is initialized if accessed. Initial version of text on how NVM Subsystem Shutdown interacts with Unsafe Shutdowns field. Consolidate events for both controller shutdown and NVM subsystem shutdown. Move permission to initialize memory to 3.7.2 (Controller Level Reset) Additional edits and comments.
2024-04-25 (D. Black)	<ul style="list-style-type: none"> Reset change tracking – clean version for Technical WG discussion. Minor updates & additional comments from Technical WG discussion Clarify text on how to start executing commands after a shutdown. Clarify 3.7.2 material on resetting controller transport properties/attributes
2024-04-26 (D. Black)	<ul style="list-style-type: none"> Copy 3.6.1 “to start executing commands” text for memory based controllers to 3.6.2 for message based controllers and make necessary modifications. Initial changes to TP 6032 material.

2024-04-29 (D. Black)	<ul style="list-style-type: none"> • Added power scope sentence to 3.6.1 on controller shutdown. • Reduce amount of change TP 6032 text based on comments from Judy Brock • Remove shutdown text in 3.6.* that links to Unsafe Shutdowns field • Media initialization is not allowed between NVM Subsystem Shutdown and NVM Subsystem Reset (3.7.2) • Remove disabling of Controller Reset during NVM Subsystem Shutdown.
2024-04-30 (D. Black)	<ul style="list-style-type: none"> • Initial sentence on clearing CC.SHST to 00b when CC.EN and CC.SHN are both written to enable a controller without a Controller Level Reset. • Add TP 6032 table text for CSTS.SHST=00b.
2024-05-01 (D. Black)	<ul style="list-style-type: none"> • Explain what happens after an event occurs that ends “safe to power off” • Remove Unsafe Shutdowns material, separate TP will be used for any changes • Clarify multi-domain NVM Subsystem Shutdown text to always be domain-specific.
2024-05-02 (D. Black)	<ul style="list-style-type: none"> • Minor edits from Judy Brock. A few open issues in TP6032 material resulted. • Accept all changes to provide baseline for Technical WG discussion. • Comments and changes from Technical WG discussion: <ul style="list-style-type: none"> ○ Clearing CC.SHN also clears CSTS.SHST after a controller shutdown. ○ Edit TP6032 text to include clearing CSTS.SHST to zero. ○ Additional minor edits • Added Estimated Time for NVM Subsystem Shutdown to Identify Controller
2024-05-03 (D. Black)	<ul style="list-style-type: none"> • Explain CSTS.SHST behavior on all controllers when NVM Subsystem Shutdown completes. • NVMe-MI material edits from Austin.
2024-05-05 (D. Black)	<ul style="list-style-type: none"> • Rewrite “power off” material to not use either “safe” or “unsafe” • Put back Unsafe Shutdowns material and modify to match “power off” changes • Rewrite “permitted to prepare media for use” in Controller Level Reset • Consistently use CSTS.SHST value descriptions. • Additional minor edits, mostly from Austin
2024-05-06 (D. Black)	<ul style="list-style-type: none"> • Initial Description of Changes text. • Add section 3.6.3 material on difference in NVM Subsystem Shutdown between this TP and revision 2.0 of the Base Spec. • Change the support bit for NSM Subsystem Shutdown changes in this TP to be an extension of the existing CAP.NSSS support bit for the original NVM Subsystem Shutdown.
2024-05-07 (D. Black)	<ul style="list-style-type: none"> • While NVM Subsystem Shutdown is occurring or complete, a CLR does not affect that shutdown. • Add potential text changes to require CSTS.SHST=00b when CSTS.RDY transitions from ‘0’ to ‘1’. • Change “prepared for power off” to “ready to be powered off” to avoid confusion. • Add 3.6.3 guidance to host on how to deal with both old and new NVM Subsystem Shutdown behaviors (keep calm and use NVM Subsystem Reset). • Additional minor edits.
2024-05-10 (D. Black)	<ul style="list-style-type: none"> • Merge in changes from 5/9 Technical WG discussion. • In section 3.6.3 discussion of NVM Subsystem Shutdown changes: <ul style="list-style-type: none"> ○ New functionality: Clarify CSTS.ST/CSTS.SHST preservation. ○ Old functionality: Remove bullet on transport specific resets and CLR. • Add shutdown processing overview material in section 3.6. • Rename Unsafe Shutdowns field to avoid confusion.
2024-05-13 (D. Black)	<ul style="list-style-type: none"> • Correct TP6032 changes to avoid allowing media prep for use during NVM Subsystem Shutdown. • Scope 3.6.3 discussion of NVM Subsystem Shutdown changes to CLR interactions.
2024-05-14 (D. Black)	<ul style="list-style-type: none"> • Lots of edits from Mike Allison, including pulling in all of section 3.7.2 base text. • Shorten explanation of CC.EN != CSTS.RDY cases in section 3.6 overview • Add a sentence to 3.6 recommending against NVM Subsystem Shutdown support on single controller NVM subsystems. • Only use “shutdown state” in MI spec material, not Base spec material • Rewrite CAP bit descriptions and section 3.6.3 based on material from Judy Brock and Mike Allison

2024-05-15 (D. Black)	<ul style="list-style-type: none"> • Remove use of “Power Loss Protection” in Unexpected Power Losses field. • Change shutdown overview text in 3.6 to use shutdown state for media instead of shutdown condition. • Use shutdown state for media everywhere with references to section 3.6. • Add title page summary (in box).
2024-05-16 (D. Black)	<ul style="list-style-type: none"> • Clean up title page summary. • Rewrite power scope sentence for controller shutdown • Remove transition to 00b and back from CSTS.SHST description. • Change 01b value of CSTS.SHST to “in progress” from “occurring”. • Additional minor edits.
2024-05-21 (D. Black)	<ul style="list-style-type: none"> • Changes based on reviews from Austin Bolen and Randy Jennings. • Accept a lot of editorial changes, remove a lot of comments. • Additional minor edits.
2024-05-23 (D. Black)	<ul style="list-style-type: none"> • In new shutdown overview (section 3.6), change “media is available for use” to “media is usable.” • In new shutdown overview (section 3.6), whether the controller is ready to be powered off after a reset is implementation specific because media initialization is allowed.
2024-05-28 (D. Black)	<ul style="list-style-type: none"> • Remove CC.SHN change, add comment. • Correct copy/paste mistakes in original NVM Subsystem Shutdown material from 2.0d Base Spec. • Add NVMe-MI changes to remove use of “safe” • Additional minor edits.
2024-05-29 (D. Black)	<ul style="list-style-type: none"> • Minor edits from review in Errata meeting and from a review by Yoni Shternhell
2024-05-30 (D. Black)	<ul style="list-style-type: none"> • Final changes for Member Review • Correct description of CSTS.ST and CSTS.SHST behavior for old version of NVM Subsystem Shutdown.
2024-06-03 (D. Black)	<ul style="list-style-type: none"> • Changing field name from Unsafe Shutdowns to Unexpected Power Losses affects additional TP 6032 material, including two figures added to this TP.
2024-06-11 (Mike Allison)	<ul style="list-style-type: none"> • Corrected the NSSL byte offsets.
2024-07-16	<ul style="list-style-type: none"> • Post Member Review version – no additional changes

Description for Changes Document for NVM Express Base Specification 2.0d

New Features/Feature Enhancements/Required Changes:

- **Feature Enhancement:**

- NVM Subsystem Shutdown allows all forms of Controller Level Reset to occur, both while an NVM Subsystem Shutdown is in progress, and after its completion until a subsequent NVM Subsystem Reset. For the transport-specific resets (e.g., PCIe Conventional Reset and Function Level Reset), this is a feature enhancement.
- While a Controller Shutdown is in progress and after its completion, all forms of Controller Level Reset permit (but do not require) the controller to both initialize media and cease being ready to be powered off. For the transport-specific resets and NVM Subsystem Reset, this is a feature enhancement.
- Corrected the specification of the Unsafe Shutdowns field in the SMART / Health Information log page to state that it counts unexpected power losses. Removed usage of the word “safe” for clarity and changed the field name accordingly.
- Documented requirements to start executing commands after controller shutdown is complete for Fabrics (message-based) controllers based on corresponding requirements for memory-based controllers.
- Added an NVM Subsystem Shutdown Latency (NSSL) field to Identify Controller to indicate how long host software should wait for an NVM Subsystem Shutdown to complete.
- Added the additional CSTS.SHST=00b case for the NVMe-MI Ignore Shutdown (ISH) bit - processing an out-of-band NVM Admin command that requires media access and has the ISH bit set to ‘1’ is able to make the media accessible. This modifies TP6032 in addition to the NVM Express Base Specification.
- If a controller that is ready to be powered off processes an Admin Command that requires access to the media (refer to Figure 105) and specifies the Ignore Shutdown bit set to ‘1’ via the out-of-band mechanism (refer to the NVM Express Management Interface Specification), then that controller ceases to be ready to be powered off.

- **Incompatible changes:**

- NVM Subsystem Shutdown allows all forms of Controller Level Reset to occur, both while an NVM Subsystem Shutdown is in progress, and after its completion until a subsequent NVM Subsystem Reset. For Controller Reset, this is an incompatible change.
- While a Controller Shutdown is in progress and after its completion, all forms of Controller Level Reset permit (but do not require) the controller to initialize media and cease being ready to be powered off. For Controller Reset, this is an incompatible change.
- After an NVM Subsystem Shutdown, the CSTS.ST and CSTS.SHST values are preserved until an NVM Subsystem Reset. Any other Controller Level Reset does not change these values.
- While controller shutdown processing is reported as in progress or is reported as complete (i.e., CSTS.ST is cleared to ‘0’ and CSTS.SHST is set to 01b or is set to 10b), if the controller becomes ready (i.e., sets CSTS.RDY to ‘1’ from ‘0’) as a consequence of the host enabling the controller (i.e., the host sets CC.EN to ‘1’ from ‘0’), then any shutdown processing stops and CSTS.SHST is cleared to 00b before the controller becomes ready.

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

New Features/Feature Enhancements/Required Changes:

- **Feature Enhancement:**
 - Added a case for the Ignore Shutdown (ISH) bit defined in TP6032 - processing an out-of-band NVM Admin command that requires media access and has the Ignore Shutdown (ISH) bit set is also able to cause the media to transition out of the shutdown state if CSTS.SHST is cleared to 00b.

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
Purple	Text moved without change.
<Green Bracketed:>	Notes to editor
Orange:	Text from a reference specification (e.g., another ratified TP)
Orange Strikethrough:	Text deleted by a reference specification (e.g., another ratified TP)

Modify portions of NVM Express Base Specification 2.0d as follows:

Modify Figure 36 as follows:

Figure 36: Offset 0h: CAP – Controller Capabilities

Bits	Type	Reset	Description										
63: 62 64	RO	0h	Reserved										
61	RO	Impl Spec	<p>NVM Subsystem Shutdown Enhancements Supported (NSSES): This bit indicates whether the controller supports enhancements to the NVM Subsystem Shutdown feature.</p> <p>If the controller supports the enhancements to the NVM Subsystem Shutdown feature as defined in section 3.6.3, then this bit shall be set to '1' and the NSSS bit shall be set to '1'. If a controller compliant with a revision of the NVM Express Base Specification later than revision 2.0 sets the NSSS bit to '1', then that controller shall set this bit to '1'.</p> <p>If this bit is cleared to '0', then the controller does not support the enhancements to the NVM Subsystem Shutdown feature as defined in section 3.6.3.</p> <p>If the NSSRS bit is cleared to '0' or the NSSS bit is cleared to '0', then this bit shall be cleared to '0'.</p>										
...										
58	RO	Impl Spec	<p>NVM Subsystem Shutdown Supported (NSSS): This bit indicates whether the controller supports the NVM Subsystem Shutdown feature defined in section 3.6.3.</p> <p>If this bit is set to '1', then if the controller supports the NVM Subsystem Shutdown feature, then this bit is set to '1'. If the NSSES bit is set to '1', then this bit shall be set to '1'.</p> <p>If this bit is cleared to '0', then the controller does not support the NVM Subsystem Shutdown feature, then this bit is cleared to '0'. If the NSSRS bit is cleared to '0', then this bit shall be cleared to '0'.</p> <p>Refer to section 3.6.3 for a description of the NVM Subsystem Shutdown feature and the behavioral enhancements associated with the NSSES bit being set to '1'.</p>										
...										
47:46	RO	Impl Spec	<p>Controller Power Scope (CPS): This field indicates scope of controlling the main power for this controller.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Power Scope</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Not Reported</td> </tr> <tr> <td>01b</td> <td>Controller scope</td> </tr> <tr> <td>10b</td> <td>Domain scope (i.e., the NVM subsystem supports multiple domains (refer to section 3.2.4).</td> </tr> <tr> <td>11b</td> <td>NVM subsystem scope (i.e., the NVM subsystem does not support multiple domains).</td> </tr> </tbody> </table> <p>If the NSSS bit is set to '1', then this field shall not be cleared to 00b.</p>	Value	Power Scope	00b	Not Reported	01b	Controller scope	10b	Domain scope (i.e., the NVM subsystem supports multiple domains (refer to section 3.2.4).	11b	NVM subsystem scope (i.e., the NVM subsystem does not support multiple domains).
Value	Power Scope												
00b	Not Reported												
01b	Controller scope												
10b	Domain scope (i.e., the NVM subsystem supports multiple domains (refer to section 3.2.4).												
11b	NVM subsystem scope (i.e., the NVM subsystem does not support multiple domains).												
...										

Figure 36: Offset 0h: CAP – Controller Capabilities

Bits	Type	Reset	Description
31:24	RO	Impl Spec	<p>Timeout (TO): This is the worst-case time that host software should wait for the CSTS.RDY bit to transition from:</p> <ul style="list-style-type: none"> a) '0' to '1' after the CC.EN bit transitions from '0' to '1'; or b) '1' to '0' after the CC.EN bit transitions from '1' to '0'. <p>This worst-case time may be experienced after events such as an abrupt shutdown, loss of main power without shutting down the controller, or activation of a new firmware image; typical times are expected to be much shorter.</p> <p>This field is in 500 millisecond units. The maximum value of this field is FFh, which indicates a 127.5 second timeout.</p> <p>If the Controller Ready Independent of Media Enable (CC.CRIME) bit is cleared to '0' and the worst-case time for the CSTS.RDY bit to change state is due to enabling the controller after the CC.EN bit transitions from '0' to '1', then this field shall be set to:</p> <ul style="list-style-type: none"> a) the value in the Controller Ready With Media Timeout (CRTO.CRWMT) field; or b) FFh if the value in the CRTO.CRWMT field is greater than FFh. <p>If the Controller Ready Independent of Media Enable (CC.CRIME) bit is set to '1' and the worst-case time for the CSTS.RDY bit to change state is due to enabling the controller after the CC.EN bit transitions from '0' to '1', then this field shall be set to:</p> <ul style="list-style-type: none"> a) the value in the Controller Ready Independent of Media Timeout (CRTO.CRIMT) field; or b) FFh if the value in the CRTO.CRIMT field is greater than FFh. <p>Controllers that support the CRTO register (refer to Figure 62) are able to indicate larger timeouts for enabling the controller. Host software should use the value in the CRTO.CRWMT field or the CRTO.CRIMT field depending on the controller ready mode indicated by the CC.CRIME bit to determine the worst-case timeout for the CSTS.RDY bit to transition from '0' to '1' after the CC.EN bit transitions from '0' to '1'. Host software that is based on revisions earlier than NVM Express Base Specification, Revision 2.0 is not required to wait for more than 127.5 seconds for the CSTS.RDY bit to transition.</p> <p>Refer to sections 3.5.3 and 3.5.4 for more information.</p>
...

Modify Figure 46 as follows:

Figure 46: Offset 14h: CC – Controller Configuration

Bits	Type	Reset	Description										
...										
15:14	RW	00b	<p>Shutdown Notification (SHN): This field is used to initiate a controller shutdown when a power down condition is expected. For a normal controller shutdown, it is expected that the controller is given time to process the controller shutdown. For an abrupt shutdown, the host may not wait for the controller shutdown to complete before power is lost.</p> <p>The controller shutdown notification values are defined as:</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>No notification; no effect</td> </tr> <tr> <td>01b</td> <td>Normal shutdown notification</td> </tr> <tr> <td>10b</td> <td>Abrupt shutdown notification</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table> <p>This field should be written by host software prior to any power down condition and prior to any change of the PCI power management state. It is recommended that this field also be written prior to a warm reset (refer to the PCI Express Base Specification). To determine when the controller shutdown processing is complete, refer to the definition of the CSTS.ST bit and the definition of the CSTS.SHST field. Refer to sections 3.6.3, 3.6.4 and 3.6.2 for additional shutdown processing details.</p> <p>Other fields in this property (including the EN bit) may be modified as part of updating this field to 01b or 10b to initiate a controller shutdown. If the EN bit is cleared to '0' such that the EN bit transitions from '1' to '0', then both a Controller Reset and a controller shutdown occur. If an NVM Subsystem Shutdown is reported as in progress or is being reported as completed (i.e., CSTS.ST is set to '1', and CSTS.SHST is set to 01b or 10b), then writes to this field modify the field value but have no effect. Refer to section 3.6.3 for details.</p>	Value	Definition	00b	No notification; no effect	01b	Normal shutdown notification	10b	Abrupt shutdown notification	11b	Reserved
Value	Definition												
00b	No notification; no effect												
01b	Normal shutdown notification												
10b	Abrupt shutdown notification												
11b	Reserved												
...										

00	RW	0b	<p>Enable (EN): When set to '1', then the controller shall process commands. When cleared to '0', then the controller shall not process commands nor post completion queue entries to Completion Queues. When the host modifies CC to clear this bit from '1' to '0', the controller is reset (i.e., a Controller Reset, refer to section 3.7.2). That reset deletes all I/O Submission Queues and I/O Completion Queues, resets the Admin Submission Queue and Completion Queue, and brings the hardware to an idle state. That reset does not affect transport specific state (e.g. PCI Express registers including MMIO MSI-X registers), nor the Admin Queue properties (AQA, ASQ, or ACQ). All other controller properties defined in this section and internal controller state (e.g., Feature values defined in section 5.27.1 that are not persistent across power states) are reset to their default values. The controller shall ensure that there is no impact (e.g., data loss) caused by that Controller Reset to the results of commands that have had corresponding completion queue entries posted to an I/O Completion Queue prior to that Controller Reset. Refer to section 3.6 3.7.2.</p> <p>When this bit is cleared to '0', the CSTS.RDY bit is cleared to '0' by the controller once the controller is ready to be re-enabled. When this bit is set to '1', the controller sets the CSTS.RDY bit to '1' when it is ready to process commands. The CSTS.RDY bit may be set to '1' before namespace(s) are ready to be accessed.</p> <p>Setting this bit from a '0' to a '1' when the CSTS.RDY bit is a '1' or clearing this bit from a '1' to a '0' when the CSTS.RDY bit is cleared to '0' has undefined results. The Admin Queue properties (AQA, ASQ, and ACQ) are only allowed to be modified when this bit is cleared to '0'.</p> <p>If an NVM Subsystem Shutdown is reported as in progress or is reported as completed (i.e., the CSTS.ST bit is set to '1', and the CSTS.SHST field is set to 01b or 10b), then:</p> <ul style="list-style-type: none"> • writes to setting this field bit from '0' to '1' modify modifies the field value but have has no effect (e.g., the controller does not respond by setting the CSTS.RDY bit to '1'); and • clearing this bit from '1' to '0' resets the controller as defined by this field. <p>Refer to section 3.6.3 for details.</p>
----	----	----	---

Modify Figure 47 as follows:

Figure 47: Offset 1Ch: CSTS – Controller Status

Bits	Type	Reset ¹	Description
...
06	RO	Impl-Spec HwInit	<p>Shutdown Type (ST): When If CSTS.SHST is set to a non-zero value, then this bit indicates the type of shutdown reported by CSTS.SHST.</p> <p>If this bit is set to '1', then CSTS.SHST is reporting the state of an NVM Subsystem Shutdown and this bit remains set to '1' until an NVM Subsystem Reset occurs.</p> <p>If this bit is cleared to '0', then CSTS.SHST is reporting the state of a controller shutdown.</p> <p>An NVM Subsystem Reset shall clear this bit to '0'. All other Controller Level Resets shall not change the value of this bit.</p> <p>If CSTS.SHST is cleared to 00b, then this bit should be ignored by the host.</p>
...

03:02	RO	<p>Shutdown Status (SHST): This field indicates the status of shutdown processing that is initiated by the host setting the CC.SHN field, the host setting the NSSD.NSSC field property, or a Management Endpoint has processed an NVMe-MI Shutdown command (refer to the NVM Express Management Interface specification). Shutdown processing is able to occur on this controller as a consequence of a host setting the NSSD.NSSC field on another controller to initiate an NVM Subsystem Shutdown that affects this controller.</p> <p>The shutdown status values are defined as:</p> <table border="1" data-bbox="581 443 1386 590"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Normal operation (no shutdown has been requested)</td> </tr> <tr> <td>01b</td> <td>Shutdown processing occurring in progress</td> </tr> <tr> <td>10b</td> <td>Shutdown processing complete</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table> <p>If this field is set to 01b (i.e., shutdown processing in progress), then:</p> <ul style="list-style-type: none"> an NVM Subsystem Reset aborts both a controller shutdown and an NVM Subsystem Shutdown; and any other type of Controller Level Reset (CLR): <ul style="list-style-type: none"> may or may not abort a controller shutdown; and shall not abort an NVM Subsystem Shutdown. <p>If this field is cleared to 00b (i.e., normal operation) when a CLR is initiated, then that CLR shall not change the value of this field.</p> <p>If this field is set to 01b when a Controller Level Reset CLR is initiated, and the shutdown processing is not aborted by that CLR, then that CLR shall not change the value of this field transitions to 00b on the reset and then to 01b to indicate the shutdown is still in progress and host software may or may not observe this transition.</p> <p>If this field is set to 01b when a CLR is initiated and shutdown processing is aborted by that CLR, then that CLR shall clear this field to 00b.</p> <p>If this field is set to 10b (i.e., shutdown processing complete) when a CLR is initiated by NVM Subsystem Reset, then that CLR shall clear this field to 00b.</p> <p>If this field is set to 10b when a CLR is initiated by a method other than NVM Subsystem Reset and:</p> <ul style="list-style-type: none"> the CSTS.ST bit is set to '1', then that CLR shall not change the value of this field; and the CSTS.ST bit is cleared to '0', then that CLR shall clear this field to 00b. <p>If the CSTS.ST bit is cleared to '0' and this field is set to 10b (i.e., controller shutdown processing is reported as complete), then to start executing commands on the controller:</p> <ul style="list-style-type: none"> If if the CC.EN bit is set to '1', after a shutdown operation (CSTS.SHST set to 10b), then a Controller Level Reset CLR (e.g., a Controller Reset) followed by enabling the controller (i.e., host sets the CC.EN bit from '0' to '1') is required (refer to section 3.6.1). If a host software submits commands to the controller without a prior Controller Level Reset CLR, then the behavior is undefined; and If if the CC.EN bit is cleared to '0', then to start executing commands on the controller: <ul style="list-style-type: none"> a Controller Level Reset CLR followed by enabling the controller is required (refer to sections 3.6.1 and 3.6.2); or the CC.EN bit is required to be set to '1' and the CC.SHN field is required to be cleared to 00b with the same write to the CC property (refer to sections 3.6.1 and 3.6.2). <p>If the CSTS.ST bit is set to '1' and this field is set to 10b (i.e., NVM Subsystem Shutdown processing is reported as complete), then an NVM Subsystem Reset followed by enabling the controller is required to start executing commands (refer</p>	Value	Definition	00b	Normal operation (no shutdown has been requested)	01b	Shutdown processing occurring in progress	10b	Shutdown processing complete	11b	Reserved
Value	Definition											
00b	Normal operation (no shutdown has been requested)											
01b	Shutdown processing occurring in progress											
10b	Shutdown processing complete											
11b	Reserved											

Figure 47: Offset 1Ch: CSTS – Controller Status

Bits	Type	Reset ¹	Description
			to section 3.6.3). If a host submits commands to the controller without a prior NVM Subsystem Reset, then the behavior is undefined. Refer to section 3.6.3 on the reset behavior of this field when CAP.CPS is set to 40b or 11b.
...
00	RO	0b	<p>Ready (RDY): This bit is set to '1' when the controller is ready to process submission queue entries after the CC.EN bit is set to '1'. This bit shall be cleared to '0' when the CC.EN bit is cleared to '0' once the controller is ready to be re-enabled. Commands should 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 produces undefined results. Refer to the definition of the CAP.TO field, sections 3.5.3, and 3.5.4 for timing information related to this field.</p> <p>If an NVM Subsystem Shutdown that affects this controller is reported as in progress or is reported as has completed that affects this controller (i.e., the CSTS.ST bit is set to '1' and the CSTS.SHST field is set to 01b or is set to 10b), then an NVM Subsystem Reset is required before this bit is allowed to be set to '1' from '0'. Refer to section 3.6.3.</p> <p>If a controller shutdown is reported as in progress or is reported as complete (i.e., the CSTS.ST bit is cleared to '0' and the CSTS.SHST field is set to 01b or is set to 10b), then before this bit is allowed to be set to '1' from '0', controller shutdown processing shall stop (e.g., complete or be terminated) and the CSTS.SHST field shall be cleared to 00b.</p>
<p>NOTE:</p> <p>1. During a Controller Level Reset, the field and bit values may transition to values other than the reset value prior to indicating the reset value.</p>			

Modify Section 3.1.3.20 as follows:

3.1.3.20 Offset 64h: NSSD – NVM Subsystem Shutdown

This optional property provides host software with the capability to initiate a normal or an abrupt NVM sSubsystem sShutdown.

Support for this property is indicated by the state of the NVM Subsystem Shutdown Supported (CAP.NSSS) field. If the property is not supported, then the address range occupied by the register is reserved.

The NVM Subsystem Shutdown Enhancements Supported (CAP.NSSES) bit affects the functionality invoked by host modification of this property (refer to section 3.6.3).

...

Modify Section 3.6 as follows:

3.6 Shutdown Processing

This section describes the recommended procedures for shutdown processing prior to a power-off condition.

There are two shutdown processing mechanisms, controller shutdown (refer to sections 3.6.1 and 3.6.2) and NVM Subsystem Shutdown (refer to section 3.6.3). The CSTS.ST bit indicates the shutdown mechanism that is in progress, if any (refer to Figure 47). A host requests a controller shutdown by modifying the CC.SHN field (refer to Figure 46). A host requests an NVM Subsystem Shutdown by modifying the NSSD property (refer to section 3.1.3.20) or by issuing an NVMe-MI Shutdown command to a Management Endpoint (refer to the NVMe Express Management Interface Specification).

At most one shutdown processing mechanism is able to be in progress for a controller at any time. If an NVM Subsystem Shutdown is requested while a controller shutdown is in progress, then the NVM Subsystem Shutdown overrides the controller shutdown. The progress and completion of shutdown processing is indicated by the CSTS.SHST field (refer to Figure 47).

NVM Subsystem Shutdown should not be supported by any NVM subsystem that does not support more than one controller, without counting virtual controllers (e.g., NVM Subsystem Shutdown should not be supported by an NVM subsystem that supports one primary controller and multiple secondary controllers) (refer to section 8.26).

Figure SSS describes the interactions of the shutdown processing state indicated by the CSTS.SHST field with the state of the controller indicated by the CC.EN bit (refer to Figure 46) and by the CSTS.RDY bit (refer to Figure 47). The four possible media states in Figure SSS are: shutdown, shutdown in progress, usable, and initialization in progress.

Figure SSS: Shutdown Processing Interactions

CC.EN	CSTS.RDY	CSTS.SHST	Controller able to process Admin and I/O commands ⁴	Media state	Controller ready to be powered off
0	0	00b	no	any	implementation specific ¹
0	0	01b	no ³	shutdown in progress	no
0	0	10b	no	shutdown	yes
1	1	00b	yes	initialization in progress or usable ²	no
1	1	01b	no ³	shutdown in progress	no
1	1	10b	no	shutdown	yes

Notes:

1. In some cases (e.g., following initial application of power, or following a Controller Level Reset that occurred while shutdown processing was reported as complete), the controller is permitted to initialize the media and cease being ready to be powered off as a consequence.
2. If the CC.CRIME bit is cleared to '0', then the media is usable. If the CC.CRIME bit is set to '1', then either media initialization is in progress or the media is usable (refer to Figure 46).
3. While shutdown processing is in progress, the controller may abort any command with a status code of Commands Aborted due to Power Loss Notification.
4. I/O commands are only able to be processed by a controller that supports I/O commands. Fabrics commands are always able to be processed by a controller that supports Fabrics commands.

Figure SSS does not include transition conditions for a controller that is becoming ready or is undergoing reset. During these transitions, the CC.EN bit and the CSTS.RDY bit have different values (refer to Figure 46 and Figure 47). The media may or may not be usable during these transitions.

Figure SSS does not include the NVMe-MI effects of processing an Admin Command that requires access to the media (refer to Figure 105) and specifies the Ignore Shutdown bit set to '1' is processed by the controller via the out-of-band mechanism (refer to the NVM Express Management Interface Specification). Processing of such a command causes the media to become usable, after which the media may or may not be returned to its previous condition.

Modify Section 3.6.1 as follows:

3.6.1 Memory-based Transport Controller Shutdown

...

It is recommended that the host wait a minimum of the RTD3 Entry Latency reported in the Identify Controller data structure (refer to Figure 276) for the shutdown operations to complete; if the value reported in RTD3 Entry Latency is 0h, then the host should wait for a minimum of one second. While shutdown processing is in progress on a controller, ~~it~~ it is not recommended to disable ~~the that~~ that controller via the CC.EN bit field. This causes a Controller Reset which may impact the time required to complete shutdown processing. While shutdown processing is in progress on a controller, ~~the that~~ that controller may abort any command with a status code of Commands Aborted due to Power Loss Notification.

~~It is safe to power off the~~ The controller is ready to be powered off (e.g., the media is in the shutdown state (refer to Figure SSS)) when the CSTS.ST bit is cleared to '0', and the CSTS.SHST field indicates that controller shutdown processing is complete (i.e., the CSTS.SHST field is set to 10b), ~~(regardless of the value of the CC.EN bit).~~ ~~It remains safe to power off the~~ The controller remains ready to be powered off (e.g., the media remains in the shutdown state) until:

- A. the controller is enabled (i.e., the CC.EN bit transitions from '0' to '1');
- B. the controller is reset by a Controller Level Reset; or
- C. an Admin Command that requires access to the media (refer to Figure 105) and specifies the Ignore Shutdown bit set to '1' is processed by the controller via the out-of-band mechanism (refer to the NVM Express Management Interface Specification).

If a Controller Level Reset occurs while controller shutdown processing is reported as complete (i.e., the CSTS.ST bit is cleared to '0' and the CSTS.SHST field is set to 10b), then the controller may remain ready to be powered off (e.g., the media remains in the shutdown state) or the controller may cease being ready to be powered off (e.g., because the controller is preparing the media for use (refer to section 3.7.2)).

If the power scope for the controller includes multiple controllers (e.g., the CAP.CPS field is set to 10b or is set to 11b), and any controller included in that power scope is not ready to be powered off, then the portion of the NVM subsystem included in that power scope is not ready to be powered off.

To start executing commands on the controller after that controller reports controller shutdown processing complete (i.e., the CSTS.ST bit is cleared to '0' and the CSTS.SHST field is set to 10b) utilizing the CC.EN bit:

- if the CC.EN bit is set to '1', then a Controller Level Reset is required to clear the CC.EN bit to '0' on that controller and the CC.EN bit is subsequently required to be set to '1' as part of the initialization sequence (refer to section 3.5); and
- if the CC.EN bit is cleared to '0', then: ~~the controller is required to be enabled (i.e., CC.EN is set to '1' from '0'). The CC.SHN is required to be cleared to 00b in the same write to the CC property~~
 - a Controller Level Reset is required and the CC.EN bit is subsequently required to be set to '1' as part of the initialization sequence (refer to section 3.5); or
 - the CC.EN bit is required to be set to '1' and the CC.SHN field is required to be cleared to 00b with the same write to the CC property (refer to Figure 46). The controller clears the CSTS.SHST field to 00b in response to that write.

The initialization sequence (refer to section 3.5) should then be executed on that controller.

It is an implementation choice whether the host aborts all outstanding commands to the Admin Queue prior to the controller shutdown. The only commands that should be outstanding to the Admin Queue when the controller reports shutdown processing complete are Asynchronous Event Request commands.

Modify Section 3.6.2 as follows:

3.6.2 Message-based Transport Controller Shutdown

...

After the CC.EN bit transitions to '0' (i.e., due to Controller Level Reset), the association between the host and controller shall be preserved for at least 2 minutes. After this time, the association may be removed if the controller has not been re-enabled.

To start executing commands on the controller after that controller reports controller shutdown processing complete (i.e., the CSTS.ST bit is cleared to '0' and the CSTS.SHST field is set to 10b) utilizing the CC.EN bit:

- if the CC.EN bit is set to '1', then a Controller Level Reset is required to clear the CC.EN bit to '0' on that controller and the CC.EN bit is subsequently required to be set to '1' as part of the initialization sequence (refer to section 3.5); and
- if the CC.EN bit is cleared to '0', then:
 - a Controller Level Reset is required and the CC.EN bit is subsequently required to be set to '1' as part of the initialization sequence (refer to section 3.5); or
 - the CC.EN bit is required to be set to '1' and the CC.SHN field is required to be cleared to 00b with a single Property Set command (refer to section 6.6) that changes the CC property (refer to Figure 46). The controller clears the CSTS.SHST field to 00b in response to that write.

The initialization sequence (refer to section 3.5) should then be executed on that controller.

Modify Section 3.6.3 as follows:

3.6.3 NVM Subsystem Shutdown

An NVM Subsystem Shutdown initiates a shutdown of all controllers in a domain or NVM subsystem from a single controller.

A controller indicates support for the NVM Subsystem Shutdown Feature by setting the CAP.NSSS bit to '1' (refer to [Figure 36](#)).

The NVM Subsystem Shutdown Feature defined in this revision of the NVM Express Base Specification includes some functionality that differs from the functionality of the NVM Subsystem Shutdown Feature defined in revision 2.0 of the NVM Express Base Specification. A controller indicates support for these functionality differences by setting the NVM Subsystem Shutdown Enhancements Supported (CAP.NSSES) bit to '1' (refer to [Figure 36](#)).

If a controller sets the CAP.NSSES bit to '1', then while an NVM Subsystem Shutdown is reported as in progress or is reported as complete (i.e., while the CSTS.ST bit is set to '1' and the CSTS.SHST field is set to 01b or is set to 10b):

- a. a Controller Reset initiates a Controller Level Reset (CLR) (refer to section [3.7.2](#)); and
- b. the values of both the CSTS.ST bit and the CSTS.SHST field (refer to [Figure 47](#)) are not changed by a CLR initiated by any method other than an NVM Subsystem Reset.

If a controller clears the CAP.NSSES bit to '0', then, as defined in revision 2.0 of the NVM Express Base Specification:

- a. while an NVM Subsystem Shutdown is reported as in progress or is reported as complete, a Controller Reset does not initiate a CLR (i.e., Controller Reset is disabled); and
- b. while an NVM Subsystem Shutdown is reported as complete, any CLR initiated by any transport-specific reset type may clear the value of the CSTS.ST bit to '0' and may clear the value of the CSTS.SHST field to 00b.

A host is able to support NVM Subsystem Shutdown functionality both on controllers that set the CAP.NSSES bit to '1' and on controllers that clear the CAP.NSSES bit to '0' by ensuring that any NVM Subsystem Shutdown is followed by an NVM Subsystem Reset regardless of the value of the CSTS.ST bit and the value of the CSTS.SHST field.

3.6.3.1 NVM Subsystem Shutdown in a Single Domain NVM Subsystem

...

While NVM Subsystem Shutdown processing is in progress, any controller in the NVM subsystem may abort any command with a status code of Commands Aborted due to Power Loss Notification.

It is recommended that the host wait a minimum of the NVM Subsystem Shutdown Latency reported in the Identify Controller data structure (refer to [Figure 276](#)) for NVM Subsystem Shutdown processing to complete; if the reported NVM Subsystem Shutdown Latency value is 0h, then the host should wait for a minimum of 30 seconds. While an NVM Subsystem Shutdown is reported as in progress, it is not recommended to reset the NVM subsystem via an NVM Subsystem Reset or a power cycle (which causes an NVM Subsystem Reset). This aborts the NVM Subsystem Shutdown which may impact the subsequent time required for the NVM subsystem to become ready to perform I/O (e.g., after power is reapplied following a power cycle).

For either a normal **shutdown** or an abrupt NVM Subsystem Shutdown, ~~it is safe to power off~~ the NVM subsystem ~~is ready to be powered off~~ (e.g., the media is in the shutdown state (refer to [Figure SSS](#))) when the CSTS.ST bit is set to '1' and the CSTS.SHST field indicates that shutdown processing is complete (i.e., the CSTS.SHST field is set to 10b) on any controller in the NVM subsystem. The NVM subsystem shall not set the CSTS.SHST field to 10b on any controller in the NVM subsystem until the entire NVM subsystem is ready to be powered off. The NVM Subsystem shall indicate that NVM Subsystem Shutdown processing is complete by setting the CSTS.SHST field to 10b on all controllers in the NVM subsystem. ~~It remains safe to power off the~~ The NVM subsystem ~~remains~~ ready to be powered off (e.g., the media remains in the shutdown state) until:

- A. an NVM Subsystem Reset; ~~occurs~~ or
- B. an Admin command that requires access to the media (refer to [Figure 105](#)) and specifies the Ignore Shutdown bit set to '1' is processed by any controller via the out-of-band mechanism (refer to the NVM Express Management Interface Specification).

If a normal or an abrupt NVM Subsystem Shutdown is ~~reported as in progress~~ ~~being processed~~ or ~~is reported as completed~~ within the NVM subsystem (i.e., the CSTS.ST bit is set to '1' and the CSTS.SHST field is set to either 01b or 10b on all controllers in the NVM subsystem, indicating that an NVM Subsystem Reset has not occurred since initiation of that NVM Subsystem Shutdown), then:

- an NVM Subsystem Reset:
 - shall abort any in progress NVM Subsystem Shutdown;
 - clears the CSTS.SHST field to 00b in all controllers in the NVM subsystem; and
 - clears the CSTS.ST bit to '0' in all controllers in the NVM subsystem;and
- a **Controller Level Reset** of any controller in the NVM subsystem that is initiated by any other method (refer to section [3.7.2](#)):
 - shall not abort any in progress NVM Subsystem Shutdown;
 - does not change the values of the CSTS.ST bit and the CSTS.SHST field, as described in [Figure 47](#); and
 - shall not cause that NVM subsystem to cease being ready to be powered off (e.g., shall not transition the media out of the shutdown state) if that NVM Subsystem was ready to be powered off when that Controller Level Reset was initiated.
- ~~any other type of Controller Level Reset has no effect on the processing of that NVM Subsystem Shutdown.~~

To start executing commands on the controller after that controller reports NVM Subsystem Shutdown processing complete (i.e., the CSTS.ST bit is set to '1' and the CSTS.SHST field is set to 10b):

- regardless of the value of the CC.EN bit, an NVM Subsystem Reset is required; and
- the CC.EN bit is subsequently required to be set to '1' as part of the initialization sequence (refer to section [3.5](#)).

The initialization sequence (refer to section [3.5](#)) should then be executed on that controller.

3.6.3.2 Domain Shutdown in a Multiple Domain NVM Subsystem

...

While NVM Subsystem Shutdown processing is in progress, any controller in the domain may abort any command with a status code of Commands Aborted due to Power Loss Notification.

It is recommended that the host wait a minimum of the NVM Subsystem Shutdown Latency reported in the Identify Controller data structure (refer to [Figure 276](#)) for NVM Subsystem Shutdown processing on a domain to complete; if the reported NVM Subsystem Shutdown Latency value is 0h, then the host should wait for a minimum of 30 seconds. While an NVM Subsystem Shutdown is reported as in progress, it is not recommended to reset the domain via either an NVM Subsystem Reset on the domain or power cycling the domain (which causes an NVM Subsystem Reset on the domain). This aborts the NVM Subsystem Shutdown which may impact the subsequent time required for the domain to become ready to perform I/O (e.g., after power is reapplied following a power cycle).

For either a normal or an abrupt NVM Subsystem Shutdown on the domain, ~~it is safe to power off~~ the domain is ready to be powered off (e.g., the media is in the shutdown state (refer to [Figure SSS](#))) when the CSTS.ST bit is set to '1' and the CSTS.SHST field indicates that shutdown processing is complete (i.e., the CSTS.SHST field is set to 10b) on any controller in the domain. The NVM subsystem shall not set the CSTS.SHST field to 10b on any controller in the domain until the entire domain is ready to be powered off. The NVM Subsystem shall indicate that NVM Subsystem Shutdown processing is complete by setting the CSTS.SHST field to 10b on all controllers in the domain. ~~It remains safe to power off the~~ The domain remains ready to be powered off (e.g., the media remains in the shutdown state) until:

- A. an NVM Subsystem Reset occurs on that domain; or
- B. an Admin command that requires access to the media (refer to [Figure 105](#)) and specifies the Ignore Shutdown bit set to '1' is processed by any controller in the domain via the out-of-band mechanism (refer to the NVM Express Management Interface Specification).

If a normal or an abrupt NVM Subsystem Shutdown is reported as in progress ~~being processed~~ or is reported as completed on a domain within a domain (i.e., the CSTS.ST bit is set to '1' and the CSTS.SHST field is set to either 01b or 10b on all controllers in the domain, indicating that an NVM Subsystem Reset on the domain has not occurred since initiation of that NVM Subsystem Shutdown), then:

- an NVM Subsystem Reset on the domain:
 - shall abort any in progress NVM Subsystem Shutdown on the domain; and
 - clears the CSTS.SHST field to 00b in all controllers in the domain; and
 - clears the CSTS.ST bit to '0' in all controllers in the domain, ; and
 - a Controller Level Reset of any controller in the domain that is initiated by any other method (refer to section [3.7.2](#))
 - shall not abort any in progress NVM Subsystem Shutdown on the domain;
 - does not change the values of the CSTS.ST bit and the CSTS.SHST field, as described in [Figure 47](#); and
 - shall not cause the domain to cease being ready to be powered off (e.g., shall not transition the media out of the shutdown state) if the domain was ready to be powered off when that Controller Level Reset was initiated.
- ~~• any other type of Controller Level Reset has no effect on the processing of that shutdown.~~

To start executing commands on the controller after that controller reports NVM Subsystem Shutdown processing complete (i.e., the CSTS.ST bit is set to '1' and the CSTS.SHST field is set to 10b):

- regardless of the value of CC.EN, an NVM Subsystem Reset on that domain is required; and
- the CC.EN bit is subsequently required to be set to '1' as part of the initialization sequence (refer to section [3.5](#)).

The initialization sequence (refer to section [3.5](#)) should then be executed on that controller.

Modify Section 3.7.2 as follows:

3.7.2 Controller Level Reset

...

The following methods initiate a Controller Level Reset

- NVM Subsystem Reset;
- Controller Reset (i.e., [the CC.EN bit](#) transitions from '1' to '0'); and
- Transport specific reset types (refer to the applicable NVMe Transport binding specification), if any

A Controller Level Reset consists of the following actions:

- The controller stops processing any outstanding Admin or I/O commands;
- All I/O Submission Queues are deleted;
- All I/O Completion Queues are deleted;
- The controller is brought to an idle state. When this is complete, [the CSTS.RDY bit](#) is cleared to '0'; and
- All NVMe controller properties defined in [either](#) section 3.1.3 [or the applicable NVMe Transport binding specification](#) and [all](#) internal controller state are reset, with the following exceptions:

- for controllers using a memory-based transport:

- the Admin Queue properties (AQA, ASQ, or ACQ) are not reset as part of a Controller Reset;
- the Controller Memory Buffer Memory Space Control property (CMBMSC) is reset as part of neither a Controller Reset nor a Function Level Reset; and
- the Persistent Memory Region Memory Space Control Upper property (PMRMSCU) and the Persistent Memory Region Memory Space Control Lower property (PMRMSCL) are not reset as part of a Controller Reset;

and

- for controllers using a message-based transport:

- there are no exceptions.

In all Controller Level Reset cases except a Controller Reset, the controller properties [defined by the transport \(e.g., the PCIe registers defined by the PCIe Base Specification\)](#) are reset as defined by the applicable NVMe Transport binding specification (e.g., [the NVM Express NVMe over PCIe Transport Specification](#)).

[In all Controller Level Reset cases, if the media is not usable and an NVM Subsystem Shutdown that includes the controller is neither reported as in progress nor reported as complete \(i.e., the CSTS.ST bit is cleared to '0' or the CSTS.SHST field is cleared to 00b\), then the controller is permitted to initialize the media for use upon completion of the Controller Level Reset.](#)

To continue after a Controller Level Reset, the host should:

- update transport specific state and controller property state as appropriate;
- set [the CC.EN bit](#) to '1';
- wait for [the CSTS.RDY bit](#) to be set to '1';
- configure the controller using Admin commands as needed;
- create I/O Completion Queues and I/O Submission Queues as needed; and
- proceed with normal I/O operations.

Note that all Controller Level Reset cases except a Controller Reset result in the controller immediately losing communication with the host. In all these cases, the controller is unable to indicate any aborts or update any completion queue entries.

Modify Figure 148 as follows:

Figure 148: Asynchronous Event Information – Notice

Value	Description
00h	<p>Namespace Attribute Changed: Indicates a change to one or more of the following:</p> <ul style="list-style-type: none"> • the Identify Namespace data structure (refer to the applicable NVMe I/O Command Set specification) for one or more namespaces; • the I/O Command Set Independent Identify Namespace data structure; • the Namespace List returned when the Identify command is issued with the CNS field set to 02h; or • other data structures as specified in applicable NVMe I/O Command Set specifications. <p>To clear this event, host software issues a Get Log Page command for the Changed Namespace List log page (refer to section 5.16.1.5) 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"> a) Namespace Status (refer to Figure 281) has changed and shutdown processing is either reported as in progress occurring (i.e., CSTS.SHST is set to 01b) or is reported as complete (i.e., CSTS.SHST is set to 10b); b) the ANAGRPID field (refer to Figure 281) has changed; or c) an I/O Command Set specific change occurs (refer to the applicable I/O Command Set specification). <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>
...	...

Modify Figure 208 as follows:

Figure 208: SMART / Health Information Log Page

Bytes	Description
...	...
159:144	<p>Unsafe Shutdowns Unexpected Power Losses (UPL): Contains the number of unsafe shutdowns a count of the number of unexpected power losses. This count shall be is incremented if, and only if, main power is lost when:</p> <ul style="list-style-type: none"> a) the controller does not report it is safe to ready to be powered off power-down (i.e., the CSTS.SHST field is not set to 10b) prior to loss of main power; or b) media is not in a the shutdown state (refer to Figure SSS) because an Admin command that accesses media as defined by Figure 103 Figure 105 was processed via the out-of-band mechanism with the Ignore Shutdown bit set to '1' (refer to the NVM Express Management Interface Specification) while shutdown processing was reported as in progress occurring or had was reported as completed (i.e., while the CSTS.SHST field was set to 01b or set to 10b). <p>Note that this field was previously named Unsafe Shutdowns.</p> <p>If power is lost when the CSTS.SHST field is not set to 10b, then subsequent controller initialization (i.e., the amount of time from when the CC.EN bit transitions from '0' to '1' until the CSTS.RDY bit transitions from '0' to '1') may take longer for any NVM subsystem, and data corruption may occur for any NVM subsystem that is not protected against power loss.</p> <p>If the Controller Power Scope (i.e., CAP.CPS) field is cleared to 00b (i.e., Not Reported) or set to 01b (i.e., Controller scope), then the controller reports that it the controller is safe to ready to be powered off power-down the controller when the controller is shutdown (i.e., CSTS.SHST field is set to 10b).</p> <p>If the CAP.CPS field is set to 10b (i.e., Domain scope), then the controller reports that it the domain is safe to ready to be powered off power-down the domain when all controllers in that domain are shutdown (e.g., NVM Subsystem Shutdown processing is complete).</p> <p>If the CAP.CPS field is set to 11b (i.e., NVM subsystem scope), then the controller reports that it the NVM subsystem is safe to ready to be powered off power-down the NVM subsystem when all controllers in the NVM subsystem are shutdown (e.g., NVM Subsystem Shutdown processing is complete).</p>
...	...

Modify Figure 234 as follows:

Figure 1: NVM Subsystem Hardware Error Event Codes

Code	Description
...	...
08h	<p>Unsafe Shutdown Unexpected Power Loss: Indicates that the controller incremented the Unsafe Shutdowns Unexpected Power Losses field value in the SMART / Health Information Log.</p> <p>Note that this field was previously named Unsafe Shutdown.</p> <p>The Additional Hardware Error Information field shall be set to the value from the Unsafe Shutdowns field in the SMART / Health Information log at the time of the event.</p> <p>Refer to Figure XXX for the format of the Additional Hardware Error Information field.</p>
...	...

Modify **Figure XXX** from TP6032 as follows:

Figure XXX: Additional Hardware Error Information for ~~Unsafe Shutdown~~ Unexpected Power Loss Errors

Bytes	Value						
15:0	<p>Unsafe Shutdowns Unexpected Power Loss: This field shall indicate the value from the Unsafe Shutdowns Unexpected Power Losses field in the SMART / Health Information log at the time of the event.</p>						
16	<p>Unsafe Shutdown Unexpected Power Loss Information: This field contains additional information about the unsafe shutdown unexpected power loss.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>7:1</td> <td>Reserved</td> </tr> <tr> <td>0</td> <td> <p>Unsafe Shutdown Unexpected Power Loss Due to Out-Of-Band Activity: If the shutdown was unsafe because media was not in a shutdown state when main power was lost because an Admin command that accesses media as defined by Figure 103 Figure 105 was processed via the out-of-band mechanism with the Ignore Shutdown bit set to '1' (refer to the NVM Express Management Interface Specification) while shutdown processing was reported as in progress occurring or had reported as completed (i.e., the value of the CSTS.SHST field was set to 01b or set to 10b), then this bit shall be set to '1'; otherwise, this bit shall be cleared to '0'.</p> </td> </tr> </tbody> </table>	Bit	Definition	7:1	Reserved	0	<p>Unsafe Shutdown Unexpected Power Loss Due to Out-Of-Band Activity: If the shutdown was unsafe because media was not in a shutdown state when main power was lost because an Admin command that accesses media as defined by Figure 103 Figure 105 was processed via the out-of-band mechanism with the Ignore Shutdown bit set to '1' (refer to the NVM Express Management Interface Specification) while shutdown processing was reported as in progress occurring or had reported as completed (i.e., the value of the CSTS.SHST field was set to 01b or set to 10b), then this bit shall be set to '1'; otherwise, this bit shall be cleared to '0'.</p>
Bit	Definition						
7:1	Reserved						
0	<p>Unsafe Shutdown Unexpected Power Loss Due to Out-Of-Band Activity: If the shutdown was unsafe because media was not in a shutdown state when main power was lost because an Admin command that accesses media as defined by Figure 103 Figure 105 was processed via the out-of-band mechanism with the Ignore Shutdown bit set to '1' (refer to the NVM Express Management Interface Specification) while shutdown processing was reported as in progress occurring or had reported as completed (i.e., the value of the CSTS.SHST field was set to 01b or set to 10b), then this bit shall be set to '1'; otherwise, this bit shall be cleared to '0'.</p>						

Modify Figure 276 as follows:

Figure 276: Identify – Identify Controller Data Structure, I/O Command Set Independent

Bytes	I/O ¹	Admin ¹	Disc ¹	Description
...
87:84	M	M	R	RTD3 Resume Latency (RTD3R): This field indicates the expected latency in microseconds to resume from Runtime D3 (RTD3). Refer to section 8.15.4. A value of 0h indicates RTD3 Resume Latency is not reported.
91:88	M	M	R	RTD3 Entry Latency (RTD3E): This field indicates the typical latency in microseconds to enter Runtime D3 (RTD3). Refer to section 8.15.4. A value of 0h indicates RTD3 Entry Latency is not reported.
...
107:104	O	O	R	NVM Subsystem Shutdown Latency (NSSL): This field indicates the typical latency in microseconds for an NVM Subsystem Shutdown to complete. Refer to section 3.6.3. A value of 0h indicates that NVM Subsystem Shutdown Latency is not reported.
...

Modify Section 8.15.4 as follows:

8.15.4 Runtime D3 (RTD3) Transitions

...

If CSTS.ST is cleared to '0', then the RTD3 Entry Latency is the expected elapsed time from the time CC.SHN is set to 01b by host software until CSTS.SHST is set to 10b by the controller. When CSTS.SHST is set to 10b, ~~it is safe for host software to remove power from~~ the controller ~~is ready for host software to remove power~~.

Modify portions of TP6032 as follows:

Description of Specification Changes for NVM Express Management Interface 1.2d

8.X Shutdown Impacts

If shutdown processing is reported as in progress ~~occurring~~ or is reported as complete (i.e., the value of the CSTS.SHST field as defined by the NVM Express Base Specification is set to 01b or is set to 10b) on a Controller, then the media may be in ~~a~~ the shutdown state. If the value of the CSTS.SHST field is cleared to 00b (i.e., normal operation), then the media is able to be in the shutdown state under specific conditions defined by the NVM Express Base Specification (e.g., after initial power application or after the value of the CSTS.SHST field has been cleared to 00b by a Controller Level Reset while shutdown processing is reported as complete).

If the media is in ~~a~~ the shutdown state, then Controller processing of any NVMe Admin Commands that is received over the out-of-band mechanism and that requires access to media (refer to the Admin Commands Permitted to Return a Status Code of Admin Command Media Not Ready section in the NVM Express Base Specification) ~~that are processed while shutdown processing is occurring or complete~~ may be impacted by the Ignore Shutdown bit (ISH) as specified in Figure YYY.

If an NVMe Admin Command that requires access to media and specifies ~~with~~ the ISH bit set to '1' ~~that accesses media~~ is processed by any Controller in the NVM Subsystem while shutdown processing is reported as in progress ~~occurring~~ or is reported as complete (i.e., the value of the CSTS.SHST field is set to 01b or 10b), then the NVM Subsystem ~~media~~ may transition the media out of the shutdown state. If the NVM Subsystem loses main power while the media is not in the shutdown state, then the ~~Unsafe Shutdown Unexpected Power Losses~~ field in the SMART / Health Information log page is incremented (refer to the NVM Express Base Specification).

If the media is transitioned out of the shutdown state ~~using~~ by any Controller in the NVM Subsystem processing an NVMe Admin Command that specifies ~~with~~ the ISH bit set to '1', then the NVM Subsystem is permitted to transition the media back into the shutdown state after processing of that ~~the~~ NVMe Admin Command is completed. If the media is not transitioned back into the shutdown state and main power is lost while the media is not in the shutdown state, then the ~~Unsafe Shutdown Unexpected Power Losses~~ field is incremented. Whether or not an NVM Subsystem transitions the media back into the shutdown state is ~~outside the scope of this specification~~ implementation specific.

In all cases where an NVMe Admin Command that ~~requires~~ ~~accesses~~ to media is processed by any Controller in the NVM Subsystem while shutdown processing is reported as in progress ~~occurring~~ or is reported as complete (i.e., the value of the CSTS.SHST field is set to 01b or 10b) in Figure YYY, if the NVMe Admin Command is aborted with the Status field in Completion Queue Entry Dword 3 set to a value of Commands Aborted due to Power Loss Notification or Admin Command Media Not Ready, then media shall not be transitioned out of the shutdown state.

If shutdown processing is reported as in progress ~~occurring~~ or is reported as complete (i.e., the value of the CSTS.SHST field is set to 01b or 10b), then:

- there shall be no impact on access to the FRU Information Device; and
- there shall be no impact on the out-of-band mechanism other than for NVMe Admin Commands that access media as described in this section.

If the Controller is in normal operation (i.e., the value of the CSTS.SHST field is cleared to 00b), then:

- there shall be no impact on access to the FRU Information Device; and
- there shall be no impact on the out-of-band mechanism.

If ~~the~~ an NVMe Admin Command does not require access to media, then the ISH bit shall have no effect on the processing of ~~the~~ that NVMe Admin Command.

Figure YYY: Shutdown Interactions with NVMe Admin Commands that Access Media

CSTS.SHST Field Value at the Time the Controller Started Processing the NVMe Admin Command	ISH Bit Value	
	ISH Bit Cleared to '0'	ISH Bit Set to '1'
<p>00b: Normal operation (i.e. e.g., no shutdown has been requested, or the CSTS.SHST field has been cleared to 00b by a Controller Level Reset after shutdown processing is reported as complete)</p>	<p>If a shutdown has been is requested (e.g., i.e., the value of the CC.SHN field as defined by the NVM Express Base Specification transitions from 00b is set to 01b for normal shutdown or from 00b to 10b for abrupt shutdown) while the NVMe Admin Command is being processed, then the NVMe Admin Command may be aborted by the eController. If the NVMe Admin Command is aborted, then the Status field in Completion Queue Entry Dword 3 shall be set to a value of Commands Aborted due to Power Loss Notification.</p>	
	<p>If media is in the shutdown state when the NVMe Admin Command is processed, then the NVMe Admin Command may be aborted by the Controller and the Status field in Completion Queue Entry Dword 3 shall be set to a value of Admin Command Media Not Ready.</p>	<p>The NVMe Admin Command shall not be aborted by the Controller with the Status field in Completion Queue Entry Dword 3 set to Commands Aborted due to Power Loss Notification or Admin Command Media Not Ready. If media is in the shutdown state when the NVMe Admin Command that requires media access is processed, then the media shall be transitioned out of the shutdown state and then the NVMe Admin Command shall be processed. Since the media is required to be transitioned out of the shutdown state, the NVMe Admin Command processing may take longer than normal. If the processing takes longer than the maximum Request-To-Response time, then a More Processing Required Response is transmitted as specified in section 4.1.2.3.</p>

Figure YYY: Shutdown Interactions with NVMe Admin Commands that Access Media

CSTS.SHST Field Value at the Time the Controller Started Processing the NVMe Admin Command	ISH Bit Value	
	ISH Bit Cleared to '0'	ISH Bit Set to '1'
01b: Shutdown processing in progress occurring	The NVMe Admin Command may be aborted by the eController. If the NVMe Admin Command is aborted, then the Status field in Completion Queue Entry Dword 3 shall be set to a value of Commands Aborted due to Power Loss Notification.	The NVMe Admin Command shall not be aborted by the eController with the Status field in Completion Queue Entry Dword 3 set to Commands Aborted due to Power Loss Notification or Admin Command Media Not Ready. If media is in the a shutdown state when the NVMe Admin Command that requires media access is processed, then the media shall be transitioned out of the shutdown state and then the NVMe Admin Command shall be processed. Since the media is required to be transitioned out of the shutdown state, the NVMe Admin Command processing may take longer than normal. If the processing takes longer than the maximum Request-To-Response time, then a More Processing Required Response is transmitted as specified in section 4.1.2.3.
10b: Shutdown processing complete	The NVMe Admin Command may be aborted by the eController. If the NVMe Admin Command is aborted, then the Status field in Completion Queue Entry Dword 3 shall be set to a value of Admin Command Media Not Ready.	If the value of the CSTS.SHST field at the time the Controller started processing the NVMe Admin Command was set to 01b (i.e. shutdown processing in progress occurring), then the value of the CSTS.SHST field shall transition to a value of 10b (i.e. shutdown processing complete) within the same amount of time as if no NVMe Admin Command with that specifies the ISH bit set to '1' had been processed (e.g., for controller shutdown, within the amount of time indicated by the RTD3 Entry Latency field as defined by the NVM Express Base Specification if the RTD3 Entry Latency is reported).

Modify portions of NVM Express Management Interface Specification 1.2d as follows:

Modify Section 5.11 as follows:

5.11 Shutdown

...

The Shutdown command completes successfully when all NVMe Controllers in the NVM Subsystem report shutdown process complete (i.e., CSTS.SHST is set to 10b and CSTS.ST is set to '1'). Refer to the [NVM Subsystem Shutdown section of the NVM Express Base Specification](#) ~~on~~ for the conditions ~~when it is safe~~ under which ~~to power down~~ the NVM Subsystem ~~is ready to be powered off~~.