



LEGAL NOTICE:

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

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

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

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

LEGAL DISCLAIMER:

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

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

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

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

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

NVM Express® Technical Proposal (TP)

Technical Proposal ID	TP4029a – Power Loss Signal Support
Change Date	2024-04-21
Builds on Specification	NVMe® Base Specification 2.0a NVMe over PCIe Transport Specification 1.0a NVMe Management Interface Specification 1.2a
Ratified Technical Proposals Referenced	TP4074 – Defining Scope for Features

Technical Proposal Author(s)

Name	Company
John Geldman, Paul Suhler	Kioxia
Yoni Shternhell	WDC
Peter Onufryk	Intel
Simon Kan, Lee Zaretsky	Dell EMC
Walt Hubis, Paul Suhler	Micron Technology

Technical Proposal Overview

This technical proposal provides a usage mechanism for the Power Loss Signaling function introduced by PCI-SIG for the Mini Express form factors.

Revision History

Revision Date	Change Description
2019-04-10	Initial version
2019-05-08	Added more interface details
2019-09-16	Added definition of unsafe shutdown referring to both CC.SHN and Power Loss Notification. Changed SMART / Health Information log page to refer to definition of unsafe shutdown. Modified section 10.4 to refer to unsafe shutdowns.
2019-09-20	Changed the definitions of the signals to use the PCIe signals as examples. Added 7.6.TBD (Power Loss Signaling).
2020-03-16	Major rewrite based on consensus discussions
2020-03-26	Comments and changes from 2020-03-26 TWG discussion Modified Unsafe Shutdown Count to domain scope (and added example)
2020-04-10	Changed to multiple FQ Time Settings (no energy)
2020-04-16	Changed some “shutdown” to “loss of main power” Added “Impl Spec” to new register bits.
2020-05-13	Added ‘n’ Recovery From Emergency Power Fail Times to Power Loss log & redid math Limited Shutdown description to CC.SHN references
2020-06-03	Reordered, recalculated Power Loss log Updated questions for Power Loss Header
2020-09-15	Revised to align with direction agreed in 2020-09-03 meeting of the Technical WG, and with subsequent discussions among the co-sponsors. Figure references: Updated to Base spec 1.4a. 1.9 PCI-SIG ECN citation: Updated to Revision B. 1.6.TBD3 unsafe shutdown definition: Deleted. TP 4082 covers unsafe shutdown adequately. Controller Status register: Changed Unsafe Shutdowns bit to Rebuild In Process bit. Power Loss log page: <ul style="list-style-type: none"> Increased time value fields to 16-bit mantissa and 7-bit exponent, to allow times longer than 40,950 seconds. Restructured from three arrays of time values, to one array of descriptors, each having five time values. Added new fields Normal Shutdown Time and Abrupt Shutdown Time. Renamed Recovery from Emergency Power Fail Time and changed to mantissa-and-exponent format. 5.14.1.2 SMART / Health Info log page: <ul style="list-style-type: none"> Changed Unsafe Shutdowns bit to Rebuild In Process bit. Tied bit definitions to settings of corresponding bits added to the CSTS register. 5.21.1.TBD: Added figure describing reporting of supported capabilities by the Get Features command. 7.6.2 Shutdown: Added paragraphs referring to the Normal Shutdown Time and Abrupt Shutdown Time fields. 8.TBD: <ul style="list-style-type: none"> Consolidated material from 7.6.TBD. Added requirement to support new CSTS bits and the Power Loss log page. Added details to Forced Quiescence.
2020-10-21	Changes from 2020-10-08 Technical WG meeting: <ul style="list-style-type: none"> Added Description for NVMe Specification Changes Document section. Deleted adding of bits to the Controller Status (CSTS) register. Deleted Normal Shutdown Time and Abrupt Shutdown Time. Moved status bit descriptions to the SMART / Health Information log page, which had previously had only brief references to the bits in CSTS. Renamed Rebuild in Progress to Recovery in Progress. Renamed PLANC bit to Analysis Not Complete (ANC). Farewell PLANC Epoch. Moved FQVT, EPFVT, and EPFRT to Power State Descriptor. Deleted new log page. Consolidated all new PLS-related capability indications in one byte in Identify Controller data structure.

2020-11-18	<p>Added comments from Yoni and Mike. Aligned Power Loss Signaling Config values to match the order in the Power Loss Signaling Information field.</p> <p>Changes from 2020-11-05 Technical WG Meeting:</p> <ul style="list-style-type: none"> Removed RNC and RIP bits from the SMART log page. They will go into TP 4084 as a single bit indicating all namespaces are ready for I/O commands. Added requirement that at least one supported power state will report non-zero times (Figure 252). Added signal transition sequence diagrams. Added description of ways to process in-flight commands.
2020-11-19	<p>Changes from 2020-11-19 Technical WG Meeting:</p> <ul style="list-style-type: none"> Removed requirements for Get Features command from the Power Loss Signaling Config feature. Reporting PLS mode support here is redundant with that in the Power Loss Signaling Information field. Revised definitions and use of the PLN and PLA signals to be transport-independent by using "property" terminology. Added a paragraph to 8.TBD about mapping to PCIe, which can be moved in Refactoring.
2020-12-08	<ul style="list-style-type: none"> Minor editorial changes Added note about command completion and atomic operation requirements. Added the new Attached Namespaces Ready status bit to the new Power Loss Signaling Information field in the Identify Controller data structure and to the new Power Loss Information field in the SMART / Health Information log page. Changed the "property" terminology to "method".
2021-01-04	<p>Comments from 2020-12-10 Technical WG Meeting and changes from e-mail discussions:</p> <ul style="list-style-type: none"> Deleted all changes to the SMART log page and removed that log page from this TP. Deleted the Attached Namespaces Ready (ANRDY) bit. Added IOI to the I/O Command Set Independent Identify Namespace data structure (TP 4105a).
2021-02-03	<p>Changes from 2021-01-14 Technical WG meeting:</p> <ul style="list-style-type: none"> Implemented some editorial suggestions by Yoni Shternhell. Added a change to NVMe-MI Shutdown command (from TP 4082). Annotated with comments by Mike Allison. Updated Phase 3 decisions. Removed Power Loss Signaling Scope from the new Power Loss Signaling Information field in the Identify Controller data structure. Added definition of PLS scope to section 8.TBD. Defined interactions between PLS processing and shutdowns and resets.
2021-03-03	<p>Changes from 2021-02-04, 2021-02-18 Technical WG meetings, etc.:</p> <ul style="list-style-type: none"> Updated Phase 2 decisions to be made. 5.15.2.TBD: Changed IOI to a two-bit field, to allow a validity indication. 5.21.1.TBD: Defined the scope of the PLS Config feature to be the same as the PLS scope. 8.TBD: Clarified behavior if device is in a power state for which no times are reported. 8.TBD: Clarified behavior if PLN is asserted when the current power state has EPFVT, FQVT, and EPFRT times that are zero. 8.TBD: Rewrote PLA method to cover cases in which a controller does not complete PLS processing.
2021-03-11	<p>Changes from 2021-03-04 and 2021-03-11 Technical WG meetings:</p> <ul style="list-style-type: none"> 5.15.2.TBD: Removed comments about support for IOI. It will remain in the Identify data structure already shown, and support will be required if PLS is enabled. 8.TBD: Replaced signal level figures with a state machine.
2021-03-23	<p>Changes from 2021-03-15 and 2021-03-22 NVMe-MI TG and 2021-03-18 Technical WG meetings:</p> <ul style="list-style-type: none"> NVMe-MI 9.1: Added note to Figure 143 about effects of EPF processing. Removed FQ requirement not to delete queues. 8.TBD: Corrected state machine figure. 8.TBD.1: Modified command processing for Forced Quiescence mode.

2021-03-29	<p>Changes from 2021-03-25 Technical WG meeting:</p> <ul style="list-style-type: none"> 8.TBD.1: Agreed that Forced Quiescence processing “should complete commands”. 8.TBD.2: Modified to clearly state that a power cycle is required to make ports active and resume command processing. <p>Changes from 2021-03-29 NVMe-MI TG meeting:</p> <ul style="list-style-type: none"> Noted that “part of” terminology may change to “is contained in” in a post-NVMe-MI 1.2 ECN.
2021-04-05	<p>Changes from 2021-04-04 NVMe-MI TG meeting:</p> <p>NVMe-MI Figure 143: Reworded notes 3 and 4 to match NVMe-MI usage.</p>
2021-04-15	<p>Changes from 2021-04-15 Technical WG meeting:</p> <ul style="list-style-type: none"> Split the PLN method Asserted-EPF value into Asserted-EPF-Active and Asserted-EPF-Inactive to allow controller to specify that transport layer will leave NVM subsystem ports active or make them not active. Moved the Power Loss Signaling state machine into its own level-3 section. 8.TBD Figure FigSM: Added branch to accommodate both of the Asserted-EPF-* values. Clarified that entry to the Idle state requires the controller to be Initialized. <p>NVMe-MI Figure 143: Modified footnotes 3 and 4 to show that EPF processing may or may not affect access.</p>
2021-04-19	<p>Final changes for Phase 2 exit revision:</p> <ul style="list-style-type: none"> 8.TBD.1 Figure FigSM: Added additional transitions for power cycle. <p>NVMe-MI Fig 143: Corrected note 4.</p>
Phase 3 Changes	
2021-05-03	<p>Initial Phase 3 revision:</p> <ul style="list-style-type: none"> Aligned to refactored specifications. Replaced the term “power loss signaling scope” with domain, because all NVM subsystems have at least one domain. 5.17.1.1 Figure 276: Packed the three new four-bit scale fields together to recover eight bits into the reserved space at the top of the structure.
2021-05-18	<p>Added comments from Swapna Galireddy.</p> <p>Changes from 2021-05-06 meeting of Technical WG.</p>
2021-06-03	<p>Changes from 2021-05-20 meeting of Technical WG, subsequent e-mail exchanges, and a general scrub of the TP:</p> <ul style="list-style-type: none"> 3.6 and 3.7: Added mention of Power Loss Processing interactions with shutdowns and resets. 5.17.1.1: Added references to support bits to EPFVTS and FQVTS fields. 5.17.1.1: Added model section references to EPFVT, FQVT, and EPFRT fields. 8.TBD.1: Modified Figure FigSM to ZNS spec style. Removed all state acronyms. 8.TBD.1: Put the normative per-state description of PLA state value and port communications in a table. 8.TBD.1: Put the normative description of conditions causing transitions in a table. PCIe Transport Spec: Aligned description of assertion and deassertion of PLN# and PLA# with the usage in the PCI-SIG ECN.
2021-06-09	<p>Changes from 2021-06-07 meeting of NVMe-MI TG:</p> <ul style="list-style-type: none"> Changed notes 3 and 4 to refer to NVM subsystem only, and not domain. Domain is not a concept in NVMe-MI yet. Changed note 4 to refer to SMBus/I2C VPD and SMBus/I2C Mux, because there may not be a management endpoint (original wording).
2021-06-16	<p>Changes from 2021-06-10 meeting of Technical WG:</p> <ul style="list-style-type: none"> 8.TBD.1: Created a subsection for each state, including a table of state transitions out of that state. 8.TBD.2: Clarified wording for fetching and processing of commands in FQ Processing.
2021-06-17	<p>Changes from 2021-06-17 meeting of Technical WG:</p> <ul style="list-style-type: none"> 8.TBD.1.x: A transition with multiple independent conditions is shown in one row with an “or” bullet list. 8.TBD.1.x: A transition with multiple required conditions is shown as an “and” bullet list.

2021-06-24	<p>Changes from 2021-06-24 meeting of Technical WG:</p> <ul style="list-style-type: none"> • 8.TBD.1: Use one bullet for transition conditions that are an AND of multiple items. • 8.TBD.2: Revised explanation of actions if power is lost and restored. • 8.TBD.2: Added requirements for commands received out-of-band. • 8.TBD.2: Added entering complete state to list of actions. • 8.TBD.3: Added requirements for commands received out-of-band. • 8.TBD.3: Added entering complete state to list of actions. • NVMe-MI: Added footnote describing handling commands received over MCTP while FQ Processing is active.
2021-07-08	<p>Changes from 2021-06-28 meeting of NVMe-MI TG:</p> <ul style="list-style-type: none"> • 8.TBD.2: In the list of FQ Processing actions, replaced details of MI actions with a reference to the new section there. • NVMe-MI: Moved notes from Figure 145 to a new Power Loss Signaling Interactions section (8.1.TBD). • NVMe-MI: FQ Processing causes Command Messages to be aborted. • NVMe-MI: Rewrote descriptions of EPF processing.
2021-07-12	<p>Changes from 2021-07-12 meeting of NVMe-MI TG:</p> <ul style="list-style-type: none"> • NVMe-MI: Rewrote Power Loss Signaling Interactions section.
2021-07-15	<p>Changes from 2021-07-15 meeting of Technical WG:</p> <ul style="list-style-type: none"> • Corrected some section and figure references. • 5.17.1.1: Cleaned up descriptions of new fields in Power State Descriptor Data Structure. • 5.27.1.TBD: Added description of attributes returned by a successful Get Features command. • 8.TBD.1: Updated titles for state transition figures to match section headings. • Changed PLN and PLA from “states” to “variables”. • 8.TBD.2, .3: Specify setting of PLA variables in the lists of actions.
2021-07-19	<p>Changes from 2021-07-19 <i>ad hoc</i> meeting:</p> <ul style="list-style-type: none"> • 5.17.1.1: Rewrote new fields in Figure 276. • 5.17.1.1: Added new Figure FigScales to define time scales. • 5.27.1.TBD: Removed text stating that the PLS Config feature shall be saveable. Other features generally have no such requirement, and the “Yes” in the “Current Setting Persists” column of Figure 316 is sufficient. • 8.TBD: Clarified that all controllers in the domain support the same PLS modes. • 8.TBD.3: Added requirement that all controllers in the domain implement the same EPF Processing state. • PCIe Transport 3.8.TBD: Edited for clarity.
2021-07-21	<p>Changes from 2021-07-21 discussion with Mike Allison:</p> <ul style="list-style-type: none"> • Made various clarifications. • 8.TBD: Decomposed the Idle state into the Not Ready state and the Ready state to reduce ambiguity in state transition conditions. • 8.TBD: Changed shutdown wording to use CSTS.SHST field. • 8.TBD.2: During FQ Processing, require aborting a Set Features command specifying the PLS Config feature.
2021-07-29	<p>Cumulative changes:</p> <ul style="list-style-type: none"> • 8.TBD.1: Added a statement that a power cycle causes a transition from any state to the Not Ready state. • 8.TBD.1.7: Removed statement that EPF Processing Active cannot be interrupted by resets or shutdowns; that was added in error. • Re-based to Base 2.0a Specification. No changes needed for PCIe Transport 1.0a Specification or anticipated for NVMe-MI 1.2a Specification.

2021-08-05	<p>Changes from 2021-08-05 Technical WG meeting:</p> <ul style="list-style-type: none"> • Corrected “* Processing Complete” names to match those in the state machine, i.e., “* Complete”. • Corrected font sizes in tables. • Added CLR/Shutdown arc from EPF Processing Active state. • Modified three statements “... ports in the domain are inactive ...” to say “... all ports in the domain are not active ...”. • Use Commands Aborted due to Power Loss Notification status code for aborting Set Features command while in FQ Processing state.
Member Review	
2021-08-06	<p>Revision for 30-day Member Review:</p> <ul style="list-style-type: none"> • Deleted all comments. • Accepted all changes.
2021-08-09	<p>Final revision for 30-day Member Review:</p> <ul style="list-style-type: none"> • Accepted comments from Silicon Motion (Edward Hsieh).
2021-08-11	<p>Member Review comment resolution:</p> <ul style="list-style-type: none"> • Global: Replaced four references to PCI-SIG ECN with M.2 Specification R4.0 V1.0. • 5.17.2.1: Added missing “in seconds” to EPFVT description. • 8.TBD.1.4 FigFQC: Added missing transition from FQ Complete to Ready.
2021-08-17	<p>Member Review comment resolution:</p> <ul style="list-style-type: none"> • Change PCI reference from ECN to M.2 spec (approved in 2021-08-12 Technical WG meeting). • Four comments closed.
2021-08-26	<p>Changed filename to new standard. Resolved comments from Samsung:</p> <ul style="list-style-type: none"> • Modified Ready state to include SHST=00b while RDY = ‘0’. • Added transition from Ready to Not Ready. • Clarified that controller in Ready state may or may not respond to PLN = ‘1’ if RDY = ‘0’. • Designated some actions during PLS processing as able to be performed in parallel. • Various editorial fixes and clarifications.
Member Review 2	
2021-09-09	<p>Changes from 2021-09-09 Technical WG meeting to prepare for second Member Review:</p> <ul style="list-style-type: none"> • Deleted closed comments. • Renamed Not Ready and Ready states in the Power Loss Signaling Processing State Machine. • Removed one inappropriate reference to the CSTS.RDY bit, to align with state machine change. • Clarified introduction to list of possible command processing impacts of a power cycle that occurs after completion of PLS processing.
2021-09-10	<p>Final revision for second Member Review:</p> <ul style="list-style-type: none"> • Clarified yesterday’s comment above about state machine state name change. • Two comments remain to highlight changes from 2021-09-09 Technical WG meeting. • All changes accepted.
2021-10-11	<p>Member Review comment resolution:</p> <ul style="list-style-type: none"> • Revised spec change descriptions to match new TP template. • Time Scale Values: Change Dh, Eh, and Fh to Reserved. • Corrected font size problem. • Annotated with Intel comments which have not been accepted. • Incorporated comments from Samsung and resolved. All comments are editorial.
2021-10-12	<p>Member Review comment resolution:</p> <ul style="list-style-type: none"> • Incorporated comments from Samsung and resolved. All comments are editorial.

Member Review 3	
2021-10-22	Editorial changes from 2021-10-14 Technical WG meeting: <ul style="list-style-type: none"> Deleted comments on reviewed changes. Changed port active/inactive terminology to enabled/disabled. In Base Spec model section, changed ports to singular. Changes from 2021-10-21 Technical WG meeting to prepare for third Member Review: <ul style="list-style-type: none"> Deleted comments on reviewed changes. 5.TBD.1: Modified state machine figure to place two transition conditions in notes. PCIe Transport Spec: Added requirement about deassertion of PLA# in multi-controller NVM subsystems. Revised NVMe-MI section.
2021-10-28	Changes from 2021-10-28 Technical WG meeting to prepare for third Member Review: <ul style="list-style-type: none"> PCIe Transport section: Made clarifying changes to the first paragraph. NVMe-MI section: Removed changes in last revision, per recommendation by NVMe-MI TG. Accepted all changes.
2021-12-01	Resolution of comments from Mike Allison. <ul style="list-style-type: none"> Added TP4074 as a ratified referenced TP. Updated Figure 316 to include changes from TP4074 as black text. Other editorial changes
2021-12-02	Changes from 2021-10-28 Technical WG meeting to prepare for Integration: <ul style="list-style-type: none"> Removed unaffected notes cells from figures. Accepted all changes and turned off tracking.
2022-01-11	Integration
2022-03-14	Resolution of comments from John Geldman <ul style="list-style-type: none"> Remove the "Draft Version" watermark Fix the copyright years Update to ® logo
TP4029a	
2024-01-31	Added Power Loss Signaling Config feature to three figures in Base Spec to specify feature support for each controller type.
2024-02-01	Added Power Loss Signaling Config feature to Figure 126 in NVMe-MI.
2024-02-08	Corrected NVMe Enclosure requirement from P to O. Corrected filename from Member Review to Phase 3.
2024-02-08	Renamed file for Member Review.
2024-04-21	Integrated

Description for Changes Document for NVMe Base Specification 2.0a

New Feature:

- Power Loss Signaling (optional)
 - Defines how a controller responds to indications that power loss is imminent.
 - Adds requirements and behavior in a new model section.
 - Adds two Power Loss Signaling variables to define communication with the transport layer.
 - Adds support indication to the Identify Controller data structure.
 - Adds the new Power Loss Signaling Config feature.
 - Adds three pairs of time fields to the Power State Descriptor.
 - References:
 - Technical Proposal 4029

Description for Changes Document for NVMe PCIe Transport Specification 1.0a

New Feature:

- Power Loss Signaling (optional)
 - Maps the behavior of the Power Loss Signaling variables (NVMe Base Specification) to the Power Loss Notification (PLN#) signal and the Power Loss Acknowledge (PLA#) signal (PCI Express M.2 Specification).
 - References:
 - Technical Proposal 4029

Description for Changes Document for NVMe Management Interface Specification 1.2a

New Feature:

- Power Loss Signaling (optional)
 - Adds requirements to the description of operations supported during NVM subsystem power states.
 - Adds a new section defining how Power Loss Signaling processing in a controller affects operations of a Management Endpoint, the SMBus/I2C VPD, and the SMBus/I2C Mux.
 - References:
 - Technical Proposal 4029

Markup Conventions:

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

Modify portions of the NVMe Base Specification 2.0a as follows:

1 Introduction

...

1.5 Definitions

...

1.5.TBD1 Power Loss Acknowledge (PLA)

The transport-specific variable that is used by the controller to inform the host of the controller's current Power Loss Signaling processing (refer to section 8.TBD).

1.5.TBD2 Power Loss Notification (PLN)

The transport-specific variable that is used to inform the controller that a main power loss event is expected to occur (refer to section 8.TBD).

...

1.8 References

...

PCI Express M.2 Specification, Revision 4.0, Version 1.0, November 5, 2020. Available from <https://www.pcisig.com>.

...

3 NVM Express Architecture

3.1 NVM Controller Architecture

...

3.1.2 Controller Types

...

3.1.2.1 I/O Controller

...

3.1.2.1.3 Features Support

...

Figure 25: I/O Controller – Feature Support

Feature Name	Feature Support Requirements ¹	Logged in Persistent Event Log ¹
...		
Power Loss Signaling Config	O	O
...		

...

3.1.2.2 Administrative Controller

...

3.1.2.2.3 Features Support

...

Figure 30: Administrative Controller – Feature Support

Feature Name	Feature Support Requirements ¹	Logged in Persistent Event Log ¹
...		
Power Loss Signaling Config	O	O
...		

...

3.1.2.3 Discovery Controller

...

3.1.2.3.4 Features Support

...

Figure 34: Discovery Controller – Feature Support

Feature Name	Feature Support Requirements ¹	Logged in Persistent Event Log ¹
...		

Figure 34: Discovery Controller – Feature Support

Feature Name	Feature Support Requirements ¹	Logged in Persistent Event Log ¹
Power Loss Signaling Config	P	P
...		

...

3.6 Shutdown Processing

...

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.

Interactions between NVM Subsystem Shutdown and Power Loss Signaling processing are described in section 8.TBD.

...

3.7 Resets

3.7.1 NVM Subsystem Reset

Interactions between NVM Subsystem Reset and Power Loss Signaling processing are described in section 8.TBD.

3.7.1.1 Single Domain NVM Subsystems

...

5 Admin Command Set

...

5.15 Get Features command

...

Figure 194: Get Features – Feature Identifiers

Description	Section Defining Format of Attributes Returned
...	
Power Loss Signaling Config	5.27.1.TBD
...	

...

5.17 Identify command

...

5.17.2 Identify Data Structures

5.17.2.1 Identify Controller data structure (CNS 01h)

The Identify Controller data structure (refer to Figure 275) is returned to the host for the controller processing the command.

Figure 275: Identify – Identify Controller Data Structure, I/O Command Set Independent

Bytes	I/O ¹	Admin ¹	Disc ¹	Description								
Controller Capabilities and Features												
...												
110	O	O	R	Power Loss Signaling Information: This field indicates information about Power Loss Signaling processing capabilities.								
				<table><tr><th>Bits</th><th>Description</th></tr><tr><td>7:2</td><td>Reserved</td></tr><tr><td>1</td><td>PLS Forced Quiescence (PLSFQ): If set to '1', then the controller supports Power Loss Signaling with Forced Quiescence (refer to section 8.TBD.2). If cleared to '0', then the controller does not support Power Loss Signaling with Forced Quiescence.</td></tr><tr><td>0</td><td>PLS Emergency Power Fail (PLSEPF): If set to '1', then the controller supports Power Loss Signaling with Emergency Power Fail (refer to section 8.TBD.3). If cleared to '0', then the controller does not support Power Loss Signaling with Emergency Power Fail.</td></tr></table>	Bits	Description	7:2	Reserved	1	PLS Forced Quiescence (PLSFQ): If set to '1', then the controller supports Power Loss Signaling with Forced Quiescence (refer to section 8.TBD.2). If cleared to '0', then the controller does not support Power Loss Signaling with Forced Quiescence.	0	PLS Emergency Power Fail (PLSEPF): If set to '1', then the controller supports Power Loss Signaling with Emergency Power Fail (refer to section 8.TBD.3). If cleared to '0', then the controller does not support Power Loss Signaling with Emergency Power Fail.
				Bits	Description							
				7:2	Reserved							
1	PLS Forced Quiescence (PLSFQ): If set to '1', then the controller supports Power Loss Signaling with Forced Quiescence (refer to section 8.TBD.2). If cleared to '0', then the controller does not support Power Loss Signaling with Forced Quiescence.											
0	PLS Emergency Power Fail (PLSEPF): If set to '1', then the controller supports Power Loss Signaling with Emergency Power Fail (refer to section 8.TBD.3). If cleared to '0', then the controller does not support Power Loss Signaling with Emergency Power Fail.											
...												

< **Note to Editor:** Suggest assigning the Power Loss Signaling Information byte to byte 110. That will leave the contiguous 8-byte block 109:102 as Reserved. >

Figure 276 defines the power state descriptor that describes the attributes of each power state. For more information on how the power state descriptor fields are used, refer to section 8.15 on power management.

Figure 276: Identify – Power State Descriptor Data Structure

Bits	Description								
255:184 220	Reserved								
219:216	<p>Emergency Power Fail Vault Time Scale (EPFVTS): This field indicates the scale for the Emergency Power Fail Vault Time field as defined in Figure FigScale.</p> <p>If the EPFVT field is cleared to 0h, then this field:</p> <ul style="list-style-type: none"> shall be cleared to 0h; and is ignored by the host. 								
215:212	<p>Forced Quiescence Vault Time Scale (FQVTS): This field indicates the scale for the Forced Quiescence Vault Time field as defined in Figure FigScale.</p> <p>If the FQVT field is cleared to 0h, then this field:</p> <ul style="list-style-type: none"> shall be cleared to 0h; and is ignored by the host. 								
211:208	<p>Emergency Power Fail Recovery Time Scale (EPFRTS): This field indicates the scale for the Emergency Power Fail Recovery Time field as defined in Figure FigScale.</p> <p>If the EPFRT field is cleared to 0h, then this field:</p> <ul style="list-style-type: none"> shall be cleared to 0h; and is ignored by the host. 								
207:200	<p>Emergency Power Fail Vault Time (EPFVT): This field, with the EPFVTS field, indicates the worst-case time required for the controller to complete Emergency Power Fail Processing (refer to section 8.TBD.3) in the absence of failures. The time is equal to the value in this field multiplied by the scale indicated by the EPFVTS field.</p> <p>If Power Loss Signaling with Emergency Power Fail is not supported (i.e., the PLSEPF bit is cleared to '0' in the Power Loss Signaling Information field of the Identify Controller data structure (refer to Figure 275)), then this field shall be cleared to 0h.</p> <table border="1"> <thead> <tr> <th>Value</th><th>Definition</th></tr> </thead> <tbody> <tr> <td>0</td><td>Not reported</td></tr> <tr> <td>1 to 99</td><td>Time value</td></tr> <tr> <td>100 to 255</td><td>Reserved</td></tr> </tbody> </table>	Value	Definition	0	Not reported	1 to 99	Time value	100 to 255	Reserved
Value	Definition								
0	Not reported								
1 to 99	Time value								
100 to 255	Reserved								
199:192	<p>Forced Quiescence Vault Time (FQVT): This field, with the FQVTS field, indicates the worst-case time required for the controller to complete Forced Quiescence Processing (refer to section 8.TBD.2). The time is equal to the value in this field multiplied by the scale indicated in the FQVTS field.</p> <p>If Power Loss Signaling with Forced Quiescence is not supported (i.e., the PLSFQ bit is cleared to '0' in the Power Loss Signaling Information field of the Identify Controller data structure (refer to Figure 275)), then this field shall be cleared to 0h.</p> <table border="1"> <thead> <tr> <th>Value</th><th>Definition</th></tr> </thead> <tbody> <tr> <td>0</td><td>Not reported</td></tr> <tr> <td>1 to 99</td><td>Time value</td></tr> <tr> <td>100 to 255</td><td>Reserved</td></tr> </tbody> </table>	Value	Definition	0	Not reported	1 to 99	Time value	100 to 255	Reserved
Value	Definition								
0	Not reported								
1 to 99	Time value								
100 to 255	Reserved								
191:184	<p>Emergency Power Fail Recovery Time (EPFRT): This field, with the EPFRTS field, indicates the worst-case time required for the controller to complete the first initialization (as described in section 8.TBD.3) after an Emergency Power Fail process which completed before power loss. The time is equal to the value in this field multiplied by the scale indicated in the EPFRTS field.</p> <p>If Power Loss Signaling with Emergency Power Fail is not supported (i.e., the PLSEPF bit is cleared to '0' in the Power Loss Signaling Information field of the Identify Controller data structure (refer to Figure 275)), then this field shall be cleared to 0h.</p> <table border="1"> <thead> <tr> <th>Value</th><th>Definition</th></tr> </thead> <tbody> <tr> <td>0</td><td>Not reported</td></tr> <tr> <td>1 to 99</td><td>Time value</td></tr> <tr> <td>100 to 255</td><td>Reserved</td></tr> </tbody> </table>	Value	Definition	0	Not reported	1 to 99	Time value	100 to 255	Reserved
Value	Definition								
0	Not reported								
1 to 99	Time value								
100 to 255	Reserved								

Figure 276: Identify – Power State Descriptor Data Structure

Bits	Description										
183:182	<p>Active Power Scale (APS): This field indicates the scale for the Active Power field. If an Active Power Workload is reported for a power state, then the Active Power Scale shall also be reported for that power state.</p> <table> <tr> <th>Value</th><th>Definition</th></tr> <tr> <td>00b</td><td>Not reported for this power state</td></tr> <tr> <td>01b</td><td>0.0001 W</td></tr> <tr> <td>10b</td><td>0.01 W</td></tr> <tr> <td>11b</td><td>Reserved</td></tr> </table>	Value	Definition	00b	Not reported for this power state	01b	0.0001 W	10b	0.01 W	11b	Reserved
Value	Definition										
00b	Not reported for this power state										
01b	0.0001 W										
10b	0.01 W										
11b	Reserved										
...											

Figure FigScale describes the time scales indicated by the EPFVTS field, the FQVTS field, and the EPFRTS field (refer to Figure 276).

Figure FigScale: Time Scale Values

Value	Definition
0h	1 microsecond
1h	10 microseconds
2h	100 microseconds
3h	1 millisecond
4h	10 milliseconds
5h	100 milliseconds
6h	1 second
7h	10 seconds
8h	100 seconds
9h	1,000 seconds
Ah	10,000 seconds
Bh	100,000 seconds
Ch	1,000,000 seconds
Dh	Reserved
Eh	Reserved
Fh	Reserved

...

5.17.2.8 I/O Command Set Independent Identify Namespace data structure (CNS 08h)

...

Figure 280: Identify – I/O Command Set Independent Identify Namespace Data Structure

Bytes	O/M ¹	Description																		
...																				
14	M	Namespace Status (NSTAT): This field indicates the status of the namespace with the specified NSID. <table><tr><th>Bits</th><th>Description</th></tr><tr><td>7:4³</td><td>Reserved</td></tr><tr><td>2:1</td><td>I/O Impacted (IOI): This field indicates whether I/O performance is currently degraded (e.g., I/O command processing is delayed or slow).<table><tr><th>Value</th><th>Description</th></tr><tr><td>00b</td><td>I/O performance degradation is not reported</td></tr><tr><td>01b</td><td>Reserved</td></tr><tr><td>10b</td><td>I/O performance is not currently degraded</td></tr><tr><td>11b</td><td>I/O performance is currently degraded</td></tr></table></td></tr><tr><td>0</td><td>Namespace Ready (NRDY): A value of '1' indicates that the namespace is ready (refer to section 3.5.3). A value of '0' indicates that the namespace is not ready.</td></tr></table>	Bits	Description	7:4 ³	Reserved	2:1	I/O Impacted (IOI): This field indicates whether I/O performance is currently degraded (e.g., I/O command processing is delayed or slow). <table><tr><th>Value</th><th>Description</th></tr><tr><td>00b</td><td>I/O performance degradation is not reported</td></tr><tr><td>01b</td><td>Reserved</td></tr><tr><td>10b</td><td>I/O performance is not currently degraded</td></tr><tr><td>11b</td><td>I/O performance is currently degraded</td></tr></table>	Value	Description	00b	I/O performance degradation is not reported	01b	Reserved	10b	I/O performance is not currently degraded	11b	I/O performance is currently degraded	0	Namespace Ready (NRDY): A value of '1' indicates that the namespace is ready (refer to section 3.5.3). A value of '0' indicates that the namespace is not ready.
Bits	Description																			
7:4 ³	Reserved																			
2:1	I/O Impacted (IOI): This field indicates whether I/O performance is currently degraded (e.g., I/O command processing is delayed or slow). <table><tr><th>Value</th><th>Description</th></tr><tr><td>00b</td><td>I/O performance degradation is not reported</td></tr><tr><td>01b</td><td>Reserved</td></tr><tr><td>10b</td><td>I/O performance is not currently degraded</td></tr><tr><td>11b</td><td>I/O performance is currently degraded</td></tr></table>	Value	Description	00b	I/O performance degradation is not reported	01b	Reserved	10b	I/O performance is not currently degraded	11b	I/O performance is currently degraded									
Value	Description																			
00b	I/O performance degradation is not reported																			
01b	Reserved																			
10b	I/O performance is not currently degraded																			
11b	I/O performance is currently degraded																			
0	Namespace Ready (NRDY): A value of '1' indicates that the namespace is ready (refer to section 3.5.3). A value of '0' indicates that the namespace is not ready.																			
4095:15		Reserved																		

...

5.27 Set Features command

...

5.27.1 Feature Specific Information

...

Figure 316: Set Features – Feature Identifiers

Feature Identifier	Current Setting Persists Across Power Cycle and Reset ²	Uses Memory Buffer for Attributes	Feature Name	Scope ⁶
00h	Reserved			
...				
1Bh	Yes	No	Power Loss Signaling Config	Domain
...				

...

5.27.1.TBD Power Loss Signaling Config (Feature Identifier 1Bh)

This Feature configures the behavior of Power Loss Signaling (refer to section 8.TBD). The scope of this Feature is as described in Figure 316. The attributes are specified in Command Dword 11 (refer to Figure FigFeat).

If a Get Features command is successfully completed for this Feature, then the attribute described in Figure FigFeat is returned in Dword 0 of the completion queue entry for that command.

If a Set Features command is submitted for this Feature and the Power Loss Signaling Mode field specifies a Power Loss Signaling mode that is not supported (refer to the Power Loss Signaling Information field in [Figure 275](#)), then that command shall be aborted with a status code of Invalid Field in Command.

If a Set Features command is processed while the controller is in the FQ Processing state, then that command is aborted as described in section [8.TBD.2](#).

Figure FigFeat: Power Loss Signaling Config – Command Dword 11

Bits	Description	
31:02	Reserved	
01:00	Power Loss Signaling Mode (PLSM): Specifies the Power Loss Signaling mode of operation.	
	Value	Description
	00b	Power Loss Signaling not enabled
	01b	Power Loss Signaling with Emergency Power Fail enabled
	10b	Power Loss Signaling with Forced Quiescence enabled
	11b	Reserved

...

8 Extended Capabilities

...

8.TBD Power Loss Signaling (Optional)

Power Loss Signaling (PLS) is a capability that the host uses to inform all controllers in a domain of impending power loss, and that each controller uses to inform the host that the controller is preparing for that power loss.

There are two modes of Power Loss Signaling processing, Forced Quiescence Processing (refer to section [8.TBD.2](#)) and Emergency Power Fail Processing (refer to section [8.TBD.3](#)). All controllers in a domain support the same modes (i.e., all controllers report the same value in the PLSFQ bit and all controllers report the same value in the PLSEPF bit; refer to [Figure 275](#)).

Not more than one Power Loss Signaling mode is active at any time. The host uses the Power Loss Signaling Config feature (refer to section [5.27.1.TBD](#)) to select the mode of operation or to disable Power Loss Signaling. All controllers in a domain use the same mode (i.e., the scope of the Power Loss Signaling Config feature is domain). The selection persists across power cycles as defined in [Figure 316](#).

Each controller contains two variables which are used by Power Loss Signaling to perform communication between the host and controller:

- The Power Loss Notification (PLN) variable is set by the NVMe Transport and has two values, Asserted and Deasserted.
- The Power Loss Acknowledge (PLA) variable is set by the controller and has four values, Asserted-FQ, Asserted-EPF-Enabled, Asserted-EPF-Disabled, and Deasserted.

The values of the variables are described in [Figure FigVars](#).

Figure FigVars: Power Loss Signaling Variables

Variable Name	Reset ¹	Description
PLN	Deasserted	PLN Value: The PLN variable values are defined as follows:

Variable Name	Reset ¹	Description										
PLA	Deasserted	PLA Value: The PLA variable values are defined as follows:										
		<table><tr><th>Value</th><th>Description</th></tr><tr><td>Asserted-FQ</td><td>The controller is performing Forced Quiescence Processing (refer to section 8.TBD.2).</td></tr><tr><td>Asserted-EPF-Enabled</td><td>The controller is performing Emergency Power Fail processing and the port is enabled (refer to section 8.TBD.3).</td></tr><tr><td>Asserted-EPF-Disabled</td><td>The controller is performing Emergency Power Fail Processing and the port is disabled (refer to section 8.TBD.3).</td></tr><tr><td>Deasserted</td><td>The controller is not performing Power Loss Signaling processing.</td></tr></table>	Value	Description	Asserted-FQ	The controller is performing Forced Quiescence Processing (refer to section 8.TBD.2).	Asserted-EPF-Enabled	The controller is performing Emergency Power Fail processing and the port is enabled (refer to section 8.TBD.3).	Asserted-EPF-Disabled	The controller is performing Emergency Power Fail Processing and the port is disabled (refer to section 8.TBD.3).	Deasserted	The controller is not performing Power Loss Signaling processing.
		Value	Description									
		Asserted-FQ	The controller is performing Forced Quiescence Processing (refer to section 8.TBD.2).									
		Asserted-EPF-Enabled	The controller is performing Emergency Power Fail processing and the port is enabled (refer to section 8.TBD.3).									
		Asserted-EPF-Disabled	The controller is performing Emergency Power Fail Processing and the port is disabled (refer to section 8.TBD.3).									
Deasserted	The controller is not performing Power Loss Signaling processing.											
Note:												
1. Following a Controller Level Reset, the variable shall be set to this value.												

Transport-specific details of the PLN variable and the PLA variable, including effects on communication connectivity between host and controller, are described in the Power Loss Signaling Support section of the appropriate NVMe Transport specification and in the Power Loss Signaling Interactions section of the NVMe Management Interface Specification.

If the controller supports Power Loss Signaling, then the controller:

- shall support the PLN variable as specified in this section;
- may support the PLA variable as specified in this section;
- shall support Forced Quiescence Processing (refer to section 8.TBD.2), Emergency Power Fail Processing (refer to section 8.TBD.3), or both;
- if Forced Quiescence Processing is supported, shall report a non-zero value in the Forced Quiescence Vault Time field in the Power State Descriptor of one or more of the supported power states (refer to Figure 276);
- if Emergency Power Fail Processing is supported, shall report non-zero values in the Emergency Power Fail Vault Time field and the Emergency Power Fail Recovery Time field in the Power State Descriptor of one or more of the supported power states (refer to Figure 276);
- shall support reporting of whether I/O performance is degraded in the I/O Impacted (IOI) field (i.e., reports values 10b and 11b) in the I/O Command Set Independent Identify Namespace data structure (refer to Figure 280); and
- shall support the Power Loss Signaling Config feature (refer to section 5.27.1.TBD).

If the PLN variable is set to Asserted, then the controller performs either Forced Quiescence or Emergency Power Fail Processing, as determined by the setting of the Power Loss Signaling Config feature.

The controller shall ignore transitions in the PLN variable if:

- a) a Controller Level Reset (refer to section 3.7.2) is in process; or
- b) the CSTS.SHST field is not cleared to 00b (i.e., the controller is in the process of shutting down or has completed shutdown).

If the PLN variable is set to Asserted and the controller is in a power state for which:

- a) the Emergency Power Fail Vault Time field is cleared to 0h;
- b) the Forced Quiescence Vault Time field is cleared to 0h; or
- c) the Emergency Power Fail Recovery Time field is cleared to 0h,

then the time to perform an action for which the corresponding value is cleared to 0h is vendor specific.

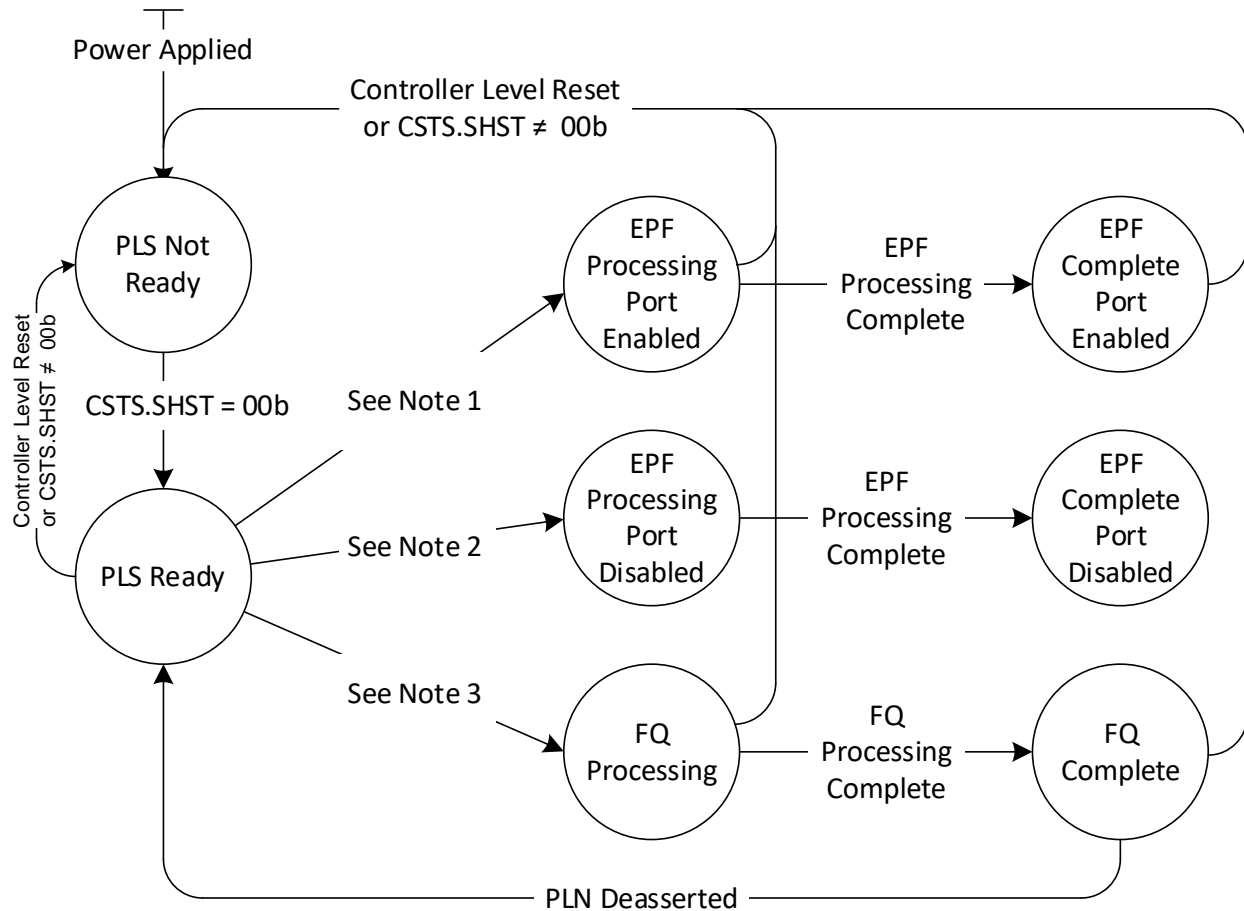
If the controller is in the EPF Complete Port Enabled state, the EPF Complete Port Disabled state, or the FQ Complete state and power is lost, then during the first restoration of power following the power loss, processing of commands may be affected while the controller performs internal recovery operations. Examples of these effects include:

- a) a namespace not being ready (i.e., the NRDY bit is cleared to '0'; refer to [Figure 280](#)); and
- b) commands to a namespace being processed at reduced performance, as indicated by the IOI field (refer to [Figure 280](#)).

8.TBD.1 Power Loss Signaling Processing State Machine

[Figure FigSM](#) illustrates how transitions in the PLN variable initiate Power Loss Signaling processing by the controller. Each circle represents a processing state.

Figure FigSM: Power Loss Signaling Processing State Machine



Note 1: EPF is Enabled & PLN is Asserted & EPF Processing Port Enabled state is supported.

Note 2: EPF is Enabled & PLN is Asserted & EPF Processing Port Disabled state is supported.

Note 3: FQ is Enabled & PLN is Asserted.

For all states, a power cycle causes a transition to the PLS Not Ready state.

The Power Loss Signaling processing state of the controller determines the value of the PLA variable and whether communications on the port are processed (refer to [Figure FigStates](#)).

Figure FigStates: PLS States

State Name	PLA Variable Value (if supported)	Port Communication Processed
PLS Not Ready	Deasserted	Yes
PLS Ready	Deasserted	Yes
FQ Processing	Asserted-FQ	Yes
FQ Complete	Deasserted	Yes
EPF Processing Port Disabled ¹	Asserted-EPF-Disabled	No
EPF Complete Port Disabled	Deasserted	No
EPF Processing Port Enabled ¹	Asserted-EPF-Enabled	Yes
EPF Complete Port Enabled	Deasserted	Yes
Notes: 1. The controller shall not implement both the EPF Processing Port Disabled state and the EPF Processing Port Enabled state.		

Note: I/O command processing in all of the PLS states, other than the PLS Not Ready state, complies with atomic operation requirements for power fail, if any, as specified in the appropriate I/O Command Set specification.

The conditions which trigger state transitions are described in the following sections. If a transition between two states can be caused by any one of multiple conditions, then those conditions are shown in a bullet list with an “or” (e.g., the transition from the FQ Processing state to the PLS Not Ready state, refer to [Figure FigFQP](#)). If a transition is caused by multiple conditions which all must occur, then those conditions are shown as a single-item bullet list with an “and” (e.g., any of the transitions from the PLS Ready state, refer to [Figure FigReady](#)).

8.TBD.1.1 PLS Not Ready State

In the PLS Not Ready state, the controller is not performing Power Loss Signaling processing. The controller enters the PLS Not Ready state following any Controller Level Reset or if the CSTS.SHST field is not cleared to 00b (i.e., the controller is in the process of shutting down or has completed shutdown).

Transitions out of this state are defined in [Figure FigNotReady](#).

Figure FigNotReady: PLS Not Ready State Transition Conditions

State Transition		Transition Condition
Starting	Ending	
PLS Not Ready	PLS Ready	<ul style="list-style-type: none"> CSTS.SHST field is 00b.

8.TBD.1.2 PLS Ready State

In the PLS Ready state, the controller is not performing Power Loss Signaling processing and is not performing shutdown processing. The controller enters the PLS Ready state when the CSTS.SHST field is 00b.

Transitions out of this state are defined in [Figure FigReady](#). If the controller is in this state and the CSTS.RDY bit is cleared to '0', then it is implementation specific whether the controller responds to a transition of the PLN variable from Deasserted to Asserted.

Figure FigReady: PLS Ready State Transition Conditions

State Transition		Transition Condition
Starting	Ending	
PLS Ready	FQ Processing	<ul style="list-style-type: none"> The controller is configured for FQ processing; and the PLN variable transitions from Deasserted to Asserted.
	EPF Processing Port Disabled	<ul style="list-style-type: none"> The controller is configured for EPF processing; the controller implements the EPF Processing Port Disabled state; and the PLN variable transitions from Deasserted to Asserted.
	EPF Processing Port Enabled	<ul style="list-style-type: none"> The controller is configured for EPF processing; the controller implements the EPF Processing Port Enabled state; and the PLN variable transitions from Deasserted to Asserted.
	PLS Not Ready	<ul style="list-style-type: none"> The controller processes a Controller Level Reset; or CSTS.SHST is not cleared to 00b.

8.TBD.1.3 FQ Processing State

In the FQ Processing state, the controller is performing Forced Quiescence Processing, as described in section 8.TBD.2.

Transitions out of this state are defined in Figure FigFQP.

Figure FigFQP: FQ Processing State Transition Conditions

State Transition		Transition Condition
Starting	Ending	
FQ Processing	FQ Complete	<ul style="list-style-type: none"> The controller completes FQ processing.
	PLS Not Ready	<ul style="list-style-type: none"> The controller processes a Controller Level Reset; or CSTS.SHST is not cleared to 00b.

8.TBD.1.4 FQ Complete State

In the FQ Complete state, the controller has completed Forced Quiescence Processing.

Transitions out of this state are defined in Figure FigFQC.

Figure FigFQC: FQ Complete State Transition Conditions

State Transition		Transition Condition
Starting	Ending	
FQ Complete	PLS Not Ready	<ul style="list-style-type: none"> The controller processes a Controller Level Reset; or CSTS.SHST is not cleared to 00b.
	PLS Ready	<ul style="list-style-type: none"> The PLN variable is set to Deasserted.

8.TBD.1.5 EPF Processing Port Disabled State

In the EPF Processing Port Disabled state, the controller is performing Emergency Power Fail Processing, as described in section 8.TBD.3. The port is disabled. The following are not able to be initiated through the port:

- Controller Level Reset;
- controller shutdown;

- c) NVM Subsystem Reset; and
- d) NVM Subsystem Shutdown.

Transitions out of this state are defined in [Figure FigEPFPI](#).

Figure FigEPFPI: EPF Processing Port Disabled State Transition Conditions

State Transition		Transition Condition
Starting	Ending	
EPF Processing Port Disabled	EPF Complete Port Disabled	<ul style="list-style-type: none"> The controller completes EPF processing.

8.TBD.1.6 EPF Complete Port Disabled State

In the EPF Complete Port Disabled state, the controller has completed Emergency Power Fail Processing. The port is disabled. The following are not able to be initiated through the port:

- a) Controller Level Reset;
- b) controller shutdown;
- c) NVM Subsystem Reset; and
- d) NVM Subsystem Shutdown.

Transitions out of this state are defined in [Figure FigEPFCI](#).

Figure FigEPFCI: EPF Complete Port Disabled State Transition Conditions

State Transition		Transition Condition
Starting	Ending	
EPF Complete Port Disabled	PLS Not Ready	<ul style="list-style-type: none"> Power is cycled.

8.TBD.1.7 EPF Processing Port Enabled State

In the EPF Processing Port Enabled state, the controller is performing Emergency Power Fail Processing, as described in section [8.TBD.3](#). The port is enabled.

Transitions out of this state are defined in [Figure FigEPFPA](#).

Figure FigEPFPA: EPF Processing Port Enabled State Transition Conditions

State Transition		Transition Condition
Starting	Ending	
EPF Processing Port Enabled	EPF Complete Port Enabled	<ul style="list-style-type: none"> The controller completes EPF processing.
	PLS Not Ready	<ul style="list-style-type: none"> The controller processes a Controller Level Reset; or CSTS.SHST is not cleared to 00b.

8.TBD.1.8 EPF Complete Port Enabled State

In the EPF Complete Port Enabled state, the controller has completed Emergency Power Fail Processing. The port is enabled.

Transitions out of this state are defined in [Figure FigEPFCA](#).

Figure FigEPFCA: EPF Complete Port Enabled State Transition Conditions

State Transition		Transition Condition
Starting	Ending	
EPF Complete Port Enabled	PLS Not Ready	<ul style="list-style-type: none"> The controller processes a Controller Level Reset; or CSTS.SHST is not cleared to 00b.

8.TBD.2 Forced Quiescence Processing

This section describes the behavior of each controller in the domain when the domain is configured for Power Loss Signaling with Forced Quiescence (refer to [section 5.27.1.TBD](#)).

If the controller enters the FQ Processing state, then the controller shall perform the following actions in sequence:

- 1) set the PLA variable, if supported, to Asserted-FQ (refer to [Figure FigStates](#));
- 2) stop fetching commands on all submission queues;
- 3) perform the following actions in parallel or in any sequence:
 - a) process commands received out-of-band on a Management Endpoint as described in the Power Loss Signaling Interactions section of the NVM Express Management Interface Specification;
 - b) if a command was fetched prior to entering this state, then process that command as described in this section; and
 - c) prepare for power loss;
- 4) enter the FQ Complete state (refer to [Figure FigFQP](#)); and
- 5) set the PLA variable to Deasserted (refer to [Figure FigStates](#)).

Entry to the FQ Processing state should cause the controller to complete processing of all previously fetched commands (e.g., to post a CQE with a status code of Successful Completion, to abort the command and post a CQE with an appropriate status code). If a background operation is in progress (e.g., a device self-test operation or a sanitize operation), then that background operation should be suspended until the PLN variable is set to Deasserted and fetching and processing commands resumes.

While in the FQ Processing state, the controller shall abort a previously-fetched Set Features command specifying the Power Loss Signaling Config feature identifier with a status code of Commands Aborted due to Power Loss Notification.

While in the FQ Processing state, if power is lost and subsequently restored, then resumption of command processing:

- a) may take more time than resumption of command processing after a normal shutdown completes and then power is cycled; and
- b) typically takes less time than resumption of command processing after an abrupt shutdown completes and then power is cycled.

If the controller transitions from the FQ Processing state to the PLS Not Ready state, then the controller shall abort Forced Quiescence Processing.

Forced Quiescence Processing shall not cause a loss of communication connectivity between the host and the controller.

If the PLN variable is set to Asserted and is then set to Deasserted without an intervening loss of main power, Controller Level Reset, or shutdown (refer to [Figure FigSM](#)), then the controller resumes fetching and processing commands.

8.TBD.3 Emergency Power Fail Processing

This section describes the behavior of each controller in the domain when the domain is configured for Power Loss Signaling with Emergency Power Fail (refer to [section 5.27.1.TBD](#)).

If the controller enters the EPF Processing Port Disabled state or the EPF Processing Port Enabled state as described in [Figure FigReady](#), then the controller shall perform the following actions in sequence:

- 1) set the PLA variable, if supported, to the value specified in [Figure FigStates](#) for that state (i.e., Asserted-EPF-Disabled or Asserted-EPF-Enabled);
- 2) stop fetching commands on all submission queues;
- 3) perform the following actions in parallel or in any sequence:
 - a) process port communications as described in [Figure FigStates](#); and
 - b) prepare for power loss in a manner that may or may not allow command processing to resume quickly in the event of power loss and then power resumption;
- 4) enter either the EPF Complete Port Enabled state (refer to [Figure FigEPFPA](#)) or the EPF Complete Port Disabled state (refer to [Figure FigEPFPI](#)); and
- 5) set the PLA variable, if supported, to Deasserted (refer to [Figure FigStates](#)).

Entry to the EPF Processing Port Disabled state or the EPF Processing Port Enabled state shall cause the controller to discard commands that were fetched from a submission queue prior to entry to the state, and to discard commands that were received out-of-band on a Management Endpoint prior to entry to the state.

The controller may implement the EPF Processing Port Disabled state or the EPF Processing Port Enabled state. The controller shall not implement both states. All controllers in the domain shall implement the same state.

The Emergency Power Fail Recovery Time (EPFRT) field and the Emergency Power Fail Recovery Time Scale (EPFRTS) field (refer to [Figure 276](#)) indicate the time that the controller requires to complete recovery during the first initialization following successful completion of Emergency Power Fail Processing. It is implementation specific whether the controller begins recovery before or after the host sets the CC.EN bit to '1'. Recovery may or may not be complete when the controller sets the CSTS.RDY bit to '1'.

If Emergency Power Fail Processing does not complete successfully (e.g., main power was lost before completion of processing), then the recovery time may exceed the Emergency Power Fail Recovery Time.

< End of changes to the NVMe Base Specification >

Modify portions of the NVMe PCI Express® Transport Specification 1.0a as follows:

1 Introduction

...

1.5 References

...

PCI Express M.2 Specification, Revision 4.0, Version 1.0, November 5, 2020. Available from <https://www.pcisig.com>.

...

3 Transport Binding

...

3.8 Transport Specific Content

...

3.8.TBD Power Loss Signaling Support (Optional)

In domains (refer to the NVMe Base Specification) that contain ports that support the PCIe Transport and that support Power Loss Signaling, each controller supports a Power Loss Notification (PLN) variable and may support a Power Loss Acknowledge (PLA) variable (refer to the Power Loss Signaling section in the NVMe Base Specification).

The PLN and PLA variables interact with the Power Loss Notification signal (PLN#) and the Power Loss Acknowledge signal (PLA#), respectively (refer to the PCI Express M.2 Specification).

If the PLN# signal is asserted at the PCIe interface, then the PLN variable shall be set to Asserted in all controllers in the domain.

If the PLN# signal is deasserted at the PCIe interface, then the PLN variable shall be set to Deasserted in all controllers in the domain.

If the PLA variable is supported and one or more controllers in the domain set the PLA variable to Asserted-FQ, then the PLA# signal, if supported, shall be asserted at the PCIe interface.

If the PLA variable is supported and one or more controllers in the domain set the PLA variable to Asserted-EPF-Enabled, then the PLA# signal, if supported, shall be asserted at the PCIe interface.

If the PLA variable is supported and one or more controllers in the domain set the PLA variable to Asserted-EPF-Disabled, then the PLA# signal, if supported, shall be asserted at the PCIe interface and all PCIe links in the domain shall become not active.

If the PLA variable is supported and all controllers in the domain set the PLA variable to Deasserted, then the PLA# signal, if supported, shall be deasserted at the PCIe interface.

To avoid potential race conditions, if the PLA variable is supported then NVM subsystems with more than one controller shall ensure that after the most recent assertion of the PLN# signal, the PLA# signal is not deasserted until each controller has asserted its PLA variable (i.e., set the PLA variable to Asserted-FQ, Asserted-EPF-Enabled, or Asserted-EPF-Disabled) and then set the PLA variable to Deasserted.

< End of changes to NVMe PCI Express Transport Specification >

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

6 NVM Express Admin Command Set

...

6.5 Set Features and Get Features

...

Figure 126: Management Endpoint - Feature Support

Feature Name ²	Feature Identifier	Support Requirements ¹	
		NVMe Storage Device	NVMe Enclosure
...			
Power Loss Signaling Config	1Bh	O	O
...			

...

8 Management Architecture

8.15 Set Features and Get Features

In the out-of-band mechanism, the ability of a Management Endpoint to receive and process Request Messages outlined in this specification is dependent on the state of the Management Endpoint. This section enumerates Management Endpoint operational times and the operations supported in each of these operational times.

The NVM Subsystem power state is defined by the state of main power and auxiliary power. Main power consists of one or more voltage rails as defined by form factor. When main power consists of multiple voltage rails, main power is considered “on” when power is good on all main voltage rails. Auxiliary power is optionally supported by a form factor and enables SMBus/I2C communications in the absence of main power. Only the Powered On and Powered Off states are applicable in form factors and platforms that do not support auxiliary power. Figure 146 defines the power states of a Management Endpoint. Note that auxiliary power is described from the perspective of the NVM Subsystem and could be provided by any appropriate power rail in a host platform.

The operations supported in each NVM Subsystem power state are summarized in Figure 146. VPD SMBus/I2C access consists of processing read operations to the FRU Information Device. SMBus/I2C MCTP access consists of processing and responding to MCTP messages on the NVM Subsystem SMBus/I2C port. PCIe MCTP access consists of processing and responding to MCTP messages issued on any NVM Subsystem PCIe port. The behavior of an operation that is “Not Supported” in Figure 146 is undefined.

Figure 145: Operations Supported During NVM Subsystem Power States

Operation	Powered Off -All Power Rails Off	Powered On -All Power Rails On	Auxiliary Power Only ² -Main Power Off -Auxiliary Power On	Main Power Only ² -Main Power On -Auxiliary Power Off
SMBus/I2C VPD and SMBus/I2C Mux Access	Not Supported	Supported ³	Supported ³	Implementation Specific
SMBus/I2C MCTP Access	Not Supported	Supported ³	Optional ¹	Implementation Specific ³
PCIe MCTP Access	Not Supported	Supported ³	Not Supported	Supported ³
NOTES: 1. An implementation that supports SMBus/I2C MCTP Access during Auxiliary Power may support a subset of commands during this power state. The commands that are supported are implementation specific. 2. Auxiliary Power Only and Main Power Only columns are not applicable to form factors that do not define Auxiliary power. 3. For interactions with Power Loss Signaling processing, refer to section 8.1.TBD.				

...

If a PCIe Command is received that targets a Controller whose corresponding PCIe link is in a low power state (i.e., PCIe ASPM), then processing of the command may cause the link to temporarily exit the low power state.

8.1.TBD Power Loss Signaling Interactions

A Controller which responds to a Power Loss Signaling notification performs either Forced Quiescence Processing or Emergency Power Fail Processing (refer to the Power Loss Signaling section of the NVMe Base Specification).

If one or more Controllers in an NVM Subsystem are in the:

- a) FQ Processing state;
- b) FQ Complete state;
- c) EPF Processing Port Enabled state;
- d) EPF Complete Port Enabled state;
- e) EPF Processing Port Disabled state; or
- f) EPF Complete Port Disabled state,

then:

- a) all Command Slots in the Management Endpoint in that NVM Subsystem should:
 - a) behave as if an implicit Abort Control Primitive (refer to section 4.2.1.3) was received with the exception that the Management Endpoint shall not transmit the Abort Control Primitive Response Messages; and
 - b) drop (silently discard) Control Primitives;
- b) access to an SMBus/I2C VPD in that NVM Subsystem may or may not be supported; and

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

- c) access to the SMBus/I2C Mux in that NVM Subsystem may or may not be supported.

When all Controllers in an NVM Subsystem have transitioned out of the FQ Complete state because the PLN variable transitioned to Deasserted, then:

- a) the Management Endpoint shall service Request Messages;
- b) access to the SMBus/I2C VPD shall be supported; and
- c) access to the SMBus/I2C Mux shall be supported.

5.28 Vital Product Data

...

< End of changes to NVMe-MI 1.2 Specification >