



## Changes in NVM Express 2.0 Specification

This document is intended to help the reader understand changes in the refactored NVMe® 2.0 family of specifications.

NVMe 2.0 is a set of the following eight specifications:

- NVM Express® Base Specification Revision 2.0
- Command Set Specifications
  - NVM Express® NVM Command Set Specification Revision 1.0
  - NVM Express® Zoned Namespaces Command Set Revision 1.1
  - NVM Express® Key Value Command Set Revision 1.0
- Transport Specifications
  - NVM Express® over PCIe Transport Revision 1.0
  - NVM Express® RDMA Transport Revision 1.0
  - NVM Express® TCP Transport Revision 1.0
- NVM Express® Management Interface Revision 1.2

This document captures three categories of NVMe 2.0 changes that are independent of the refactoring of NVMe into the above noted set of eight specs.

The first category of changes is New Features that are optional capabilities. These features may be implemented by a revision 2.0 compatible device based on market need.

The second category is changes from past behavior. One set of changes here includes extensions of past functionality.

The final category of changes are incompatible changes from past behavior, or new mandatory requirements for how a feature shall be implemented (where there may have been multiple reasonable interpretations of previous specification language).

All the changes included in NVMe 2.0 were developed in Technical Proposals and ECNs.

### 1 New Features

This section describes the new features in NVMe Specification 2.0 Family.

- **Simple copy command**
  - Defines a simple copy command used by the host to copy data from one or more source logical block ranges to a single consecutive destination logical block range.
  - References
    - NVMe NVM Command Set Specification 1.0
      - Section 3.2.2, and 5.2.2.5
    - Technical Proposal 4065b
- **Domains and Partitioning**



- Expands NVMe architecture with Domains and defines capabilities that enable large scale environments with different domains and the possibilities of partitioning of domains.
- References
  - NVMe Base Specification 2.0
    - Section 3.2.4
  - Technical Proposal 4009
- **NVM Set and Endurance Group Management**
  - Defines an interface for interoperable management of Endurance Groups and NVM Sets.
  - References
    - NVMe Base Specification 2.0
      - Sections 3.2.2, 3.2.3, 3.8, 5.3, 5.16.1.16, 5.16.1.17, 5.17.2.18, and 8.3
    - Technical Proposal 4052c
- **Command and Feature Lockdown**
  - Defines a Command and Feature Lockdown feature that allows hosts and management entities to prohibit the execution of some commands and Set Feature command feature identifiers.
  - References
    - NVMe Base Specification 2.0
      - Section 5.16.1.20, 5.19, and 8.4
    - Technical Proposal 4046a
- **NVMe KV**
  - Defines a new NVMe KV command set.
  - References
    - NVMe Key Value Command Set Specification 1.0
    - Technical Proposal 4015a
- **Zoned Namespaces**
  - Defines Zoned Namespaces and the associated Zoned Namespace Command Set. A Zoned Namespace divides the logical address space of a namespace into zones. Each zone is an LBA range that shall be written sequentially, and if rewritten shall be explicitly reset. The interface allows zoned namespaces to expose natural boundaries for internal device structures, and offload management of internal mapping tables to the host.
  - References
    - NVMe Zoned Namespace Command Set Specification 1.1
    - Technical Proposal 4053a
- **Namespace Types**
  - Extends the NVMe architecture to support different types of namespaces e.g., namespaces that contain logical blocks that have special access rules (e.g., must be written sequentially), namespaces that contain logical blocks that are used to perform data transformations, or namespaces that contain things other than logical blocks (e.g., KV pairs).
  - References

- NVMe Base Specification 2.0
    - Section 2, 2.3.2, 5.17.2.5, and 5.27.1.21
  - NVMe NVM Command Set Specification 1.0
    - Section 4.1.5.3, 4.1.6
  - Technical Proposal 4056d
- **IO Command Set Independent Data Structure**
    - Adds a new Identify data structure that provides information that is consistent across all namespaces and information that is necessary before a namespace is attached to a controller.
    - References
      - NVMe Base Specification 2.0
        - Section 5.17.2.8
      - Technical Proposal 4105a
  - **Rotational Media**
    - Adds support for rotational media (e.g., HDD)
    - References
      - NVMe Base Specification 2.0
        - References 1.5.50, 1.5.55, 1.5.56, 3.5.3, 5.16.1.10, 5.16.1.22, 5.17.2.8, 5.27.1.22, and 8.20
      - Technical Proposal 4088

## 2 Feature Enhancements

This section describes optional feature enhancements which are changes from past NVMe behavior.

- **Command and Effect Log (CEL) Enhancements**
  - Defines the following
    - A new log page to report the list of supported Log pages.
    - A new log page to describe the behaviors of all possible feature identifiers (FID) 0-255.
    - A new log page to report the list of supported NVMe-MI commands.
    - A field to the Commands Supported and Effects data structure and to the new log pages proposed to identify the scope (e.g., NVM subsystem, Domain, controller, Endurance Group, NVM Set, namespace, etc.).
  - References
    - NVMe Base Specification 2.0
      - Section 5.16.1.1, 5.16.1.6, 5.16.1.18, and 5.16.1.19
    - Technical Proposal 4071b
- **Non-MDT Command Size Limits**
  - Defines mechanisms to limit data size of commands that aren't subject to MDTs (Maximum Data Transfer Size) in Identify Controller – Dataset Management, Verify, Write Uncorrectable, and Write Zeroes.
  - References
    - NVMe NVM Command Set Specification 1.0
      - Section 3.2.3.1
    - Technical Proposal 4040a



- **CMB Write Elasticity Status**
  - Defines a mechanism for a controller to indicate to a host information used to prevent congestion in a PCI Express fabric due to CMB PCIe write requests.
  - References
    - NVMe Base Specification 2.0
      - Section 3.1.3.18, 3.1.3.19, and 8.5
    - Technical Proposal 4059a
  
- **Multiple Controller Firmware Update**
  - Defines the process of a firmware update to a domain that contains multiple controllers by updating the specification to state that overlapping downloads results in undefined behavior and allow a domain to optionally report the detection of overlapping downloads to a host.
  - References
    - NVMe Base Specification 2.0
      - Section 1.5.23, 3.11, 8.2.2
    - Technical Proposal 4060
  
- **Telemetry Log Size Changes**
  - Changes the current size limitation of the Telemetry feature. As currently specified, the logs are limited to 32MB and the goal is to allow larger telemetry log sizes.
  - References
    - NVMe Base Specification 2.0
      - Section 5.16.1.8, 5.16.1.9, 5.27.1.18, and 8.24
    - Technical Proposal 4079
  
- **Telemetry Enhancements**
  - Allows a controller to indicate to the host if the Controller-initiated log page can be overwritten by the controller.
  - References
    - NVMe Base Specification 2.0
      - Section 5.16.1.8 and 5.16.1.9
    - Technical Proposal 4063
  
- **Namespace Attachment Limits**
  - Defines the limit on the number of namespaces that can be attached to each and all controllers within a domain.
  - References
    - NVMe Base Specification 2.0
      - Section 5.17.2.1, 5.22
    - Technical Proposal 4078
  
- **Time To Ready Enhancements**
  - Distinguishes two controller ready modes for attached namespaces during controller initialization namely Controller Ready with Media Mode and Controller Ready Independent of Media Mode. Adds support bits and ready timeout values for both controller ready modes along with an indication of which mode is currently in use.
  - References

- NVMe Base Specification 2.0
  - Section 3.1, 3.1.1, 3.1.3, 3.1.5, 4.6.1.2.1, 5.14.1.13.1.5, 5.2, 5.2.1, 7.6.1.1, 7.6.2.2
- Technical Proposal 4084
  - Additional References: TP 4052b, TP 4082, TP 4105a, TP 6026
- **Graceful Disconnection**
  - Defines adding a Disconnect command to delete an NVMe I/O Queue. This command also provides a basic mechanism for an association to survive the deletion or termination of an individual I/O Queue and the loss of the associated NVM Transport connection.
  - References
    - NVMe Base Specification 2.0
      - Section 3.3.2.4 and 6.4
    - Technical Proposal 8001
- **Persistent Event Log Context Counter**
  - Adds a generation number to the Persistent Event Log so that while a host or Management Endpoint is reading the log page, it can detect if the persistent event log reporting context changed. This change is patterned after the change made to the Telemetry Enhancement (TP 4063).
  - References
    - NVMe Base Specification 2.0
      - Section 5.16.1.14
    - Technical Proposal 4083
- **SGL Optimizations**
  - Defines a backward-compatible optional threshold for any single command for the number of SGL Data Block Descriptors and SGL Bit Bucket Descriptors that have a non-zero length.
  - References
    - NVMe Base Specification 2.0
      - Section 4.1.2, and 5.17.2.1
    - Technical Proposal 4064
- **Protection Information Enhancement**
  - Defines a new 16 bytes protection information with two separate formats to provide a 32b Guard field and a 64b Guard Field. The new Storage Tag and Logical Block Reference Tag fields are also defined that have viable size.
  - References
    - NVMe NVM Command Set Specification 1.0
      - Section 4.1.2, 4.1.5.3, 4.1.3.4, 4.1.6, 5.2.1, 5.2.2.5
    - NVMe Zoned Namespace Command Set Specification 1.0
      - Section 4.1.5.1, and 4.5
    - Technical Proposal 4068b
- **ZNS Status Code and Async Notification**
  - Allows the Zone Management command to return new status codes if the request new zone state is not achieved due to a controller excursion, allows more than



one time for reporting controller recommendations, and allow the host to get a list of zones with pending controller recommendations.

- References
  - Zoned Namespace Command Set Specification 1.0
    - Section 4.1.5.1, and 3.4.2.2.3
  - Technical Proposal 4102
- **Aborting Commands During Format NVM Command**
  - Adds a list of allowed commands that may be processed during a Format NVM operation and defines the appropriate status code to be returned to notify that Admin commands are aborted.
  - References
    - NVM Command Set Specification 1.0
      - Section 5, 8.21.1
    - Technical Proposal 4047b
- **Multiple Controller Shutdown Enhancement**
  - Updates the process of shutting down a domain that contains multiple controllers.
  - References
    - NVMe Base Specification 2.0
      - Section 3.1.3.1, 3.1.3.6, 3.1.3.20, 3.6.3, 5.2.1, 5.16.1.3
    - NVM Command Set Specification 1.0
      - Section 2.1.4.2.1
    - Technical Proposal 4082
- **Persistent Event Log (PEL) State and Source Reporting**
  - Adds bits/fields that allow a host to determine if a Persistent Event Log reporting context exists, and if so, the interface (NVMe, NVMe-MI, or NVMe-oF) and Port Identifier from which the reporting context was established.
  - References
    - NVMe Base Specification 2.0
      - Section 5.16.1.1, 5.16.1.14, and 5.17.2.13
    - Technical Proposal 4086
- **Get Log Page Index Offsets**
  - Defines the scope of each Feature where the scoping is implied. In addition, Features that require the option of multiple scoping requirements are defined. Allows the Get Log Page command to specify an index offset to a specific data structure in the list of data structures.
  - References
    - NVMe Base Specification 2.0
      - Section 5.16, and 5.16.1.13
    - Technical Proposal 4075a
- **In-band Authentication Protocol for NVMe over Fabrics**
  - Defines the Diffie-Hellman HMAC-CHAP (DH-HMAC-CHAP) in-band authentication protocol for NVMe over Fabrics.
  - References
    - NVMe Base Specification 2.0
      - Section 6.3, 8.13

- Technical Proposal 8006
- **Transport Layer Security (TLS) 1.3 Support**
  - Adds support for NVMe/TCP TLS 1.3.
  - References
    - NVMe TCP Transport Specification 1.0
      - Sections 3.6.1.1-3.6.1.6
    - Technical Proposal 8011
- **Set/Get Features Ported from NVMe-MI to the NVMe 2.0 Base Spec**
  - Ports the Set/Get Features from the NVMe-MI 1.1 spec to the NVMe 1.4 spec and extends Namespace Metadata.
  - References
    - NVMe Base Specification 2.0
      - Section 5.17.2.1, 5.27.1, 5.27.1.23
    - Technical Proposal 6016a
- **Boot Partition Read Access using OOB**
  - Enables access of Boot Partitions via a log page to allow SMBus access in the event PCIe is not accessible
  - References
    - NVMe-NVM Express Base Specification 2.0
      - Section 5.16.1.21
    - Technical Proposal 6026
- **Additional Controller Metadata Types**
  - Adds additional controller metadata types useful for creating an installation record and a host-determined failure record in the Set Features event in the Persistent Event log page. Adds the new Enhanced Controller Metadata feature
  - References
    - NVMe Base Specification 2.0
      - Section 5.17.2.1, 5.27.1, 5.27.1.23
    - Technical Proposal 6009

### 3 Required Changes

This section describes mandatory behavior changes required to comply with NVM Express Specification 2.0.

- **Errata from ECN 001**
  - Required Changes
    - In the SMART / Health Information log page the Data Units Written field is not impacted by the Write Zeroes command (Reference NVMe Base Spec 2.0 section 5.16.1.3)
    - The Additional Hardware Error Information field starts at offset 04h and not 03h in the NVM Subsystem Hardware Error Event data structure (reference NVMe Base Spec 2.0 section 5.16.1.14.1.5)
    - The Namespace Capacity field in the Changed Namespace Event Data Format moved to bytes 31:24 instead of 23:16 to match the size of NCAP



in the Identify Namespace data structure in the NVM Command Set Specification 1.0 (reference NVMe Base Spec 2.0 section 16.1.14.1.6)

- Incompatible Changes
  - The CMB Queue Dword Alignment (CQDA) field in the Controller Memory Buffer Location register added the requirement that size granularity of the physically contiguous memory shall be the size of SQE or CQE.
  - The Namespace Capacity field in the Change Namespace Event Data Format (refer to Figure 222) was corrected to be 8B. Technical input submitted to the NVM Express™ Workgroup is subject to the terms of the NVM Express™ Participant's agreement. Copyright © 2014-2020 NVMe™ Corporation.
  - The reset value of fields in the AER Correctable Error Mask register changes to reflect the PCIe Base Specification.
  - A recommendation is added that the value of FFFFh for a SQE CID field should not be used as this value is reserved by the Error Information log page to indicate the error is not associated with a command.
- References
  - ECN 001
- **Errata from ECN 003**
  - Required Changes
    - The Namespace Attribute Changed notice was updated to reflect that reading the Namespace List log page with RAE bit cleared to '0' clears the event (reference NVMe Base Spec 2.0 section 5.2.1).
    - Error Count field rolls over to 1h from FFFFFFFFh (reference NVMe Base Spec 2.0 section 5.16.1.2).
    - NSID of 0h is also supported (reference NVMe base spec 2.0 section 5.16.1.3).
    - If the Endurance Group log page is supported, then the Host Read Commands field in that log page is updated by a Verify command (reference NVMe base spec 2.0 section 5.16.1.10)
    - Removed the requirement that when a persistent event log reporting context is established that a number of dword of 0 does not return a log page. This means that the original definition of NUMDU and NUMDL shall be used (reference NVMe base spec 2.0 section 5.16.1.14).
    - A Sanitize command specifying an Action set to 001b shall be successful if the most recent sanitize operation did not fail (reference NVMe base spec 2.0 section 5.24).
  - Incompatible Changes
    - In the Error Count field of the Error Information Log Entry,
      - NVMe 1.4 did not specify what happens if the value of the field is FFFFFFFFh and increments again.
      - This ECN specifies that it shall wrap to 00000001h.
    - In the Establish Context and Read Log Data field in Command Dword 10 of the Get Log Page command:
      - NVMe 1.4 specified that no log data shall be returned if the number of dwords to be returned is set to zero
      - This ECN deletes that requirement
    - In the Sanitize command:





- NVMe 1.4 did not specify how to process Sanitize Sanitize Action set to 001b (i.e., Exit Failure Mode) when Sanitize is not in process and has not failed.
    - This ECN requires that the command shall complete with a status of Successful Completion and perform no other action
  - References
    - ECN 003
- **Errata from ECN 004**
  - Required Changes
    - The Persistent Memory Region Time Units field in the PMR Capabilities register reserved values changes from “01b to 11b” to “10b to 11b” to not overlap with other defined values (reference NVMe base spec 2.0 section 3.1.3.22).
    - The following New SGL Descriptor Type and SGL Descriptor Sub Type requirements are added:
    - For SGL Descriptor Type 01h then SGL Descriptor Sub Type shall be 0h or abort the command. For SGL Descriptor Type 04h then SGL Descriptor Sub Type shall be 0h or abort the command. The status code to use is SGL Descriptor Type Invalid (reference NVMe base spec 2.0 section 4.1.2 and 3.3.3.2.1.1).
    - The Endurance Group Event Aggregate Log Page Change notification occurs due to the Endurance Group Event Aggregate log and opposed to the Predictable Latency Event Aggregate Log (reference NVMe base spec 2.0 section 5.2.1).
    - If the NOWS field in the Identify Namespace data structure is supported, then there is a requirement for a namespace associated with a NVM Set that the Optimal Write Size field shall be utilized and set appropriately. If not supported, then it is optional is the Optimal Write Size field is utilized (reference NVM command set spec 1.0 section 4.1.5).
    - The definition of the No-Deallocate Inhibited bit in the Identify Controller data structure was updated. While these was not a new requirement, previous reader may think there is and a validation should be done (reference NVMe base spec 2.0 section 5.17.2.1).
  - Incompatible Changes
    - In figure 115 (SGL Descriptor Sub Type Values), the controller is now required to report an error if the SGL Descriptor Sub Type is 1h and the SGL Descriptor Type is 1h or 4h. The handling of these cases was previously unspecified.
    - In figure 128, the description of the status code SGL Descriptor Type Invalid has been extended to include the combination of type and subtype as well.
    - In figure 106 (Command Format – Admin and NVM Command Set) handles a previously unspecified case: namespace identifier is not user for the command and the host specifies a NSID set to FFFFFFFFh.
  - References
    - ECN 004
- **Errata from ECN 006**
  - Required Changes

- The PMRMSC 64-bit register was assigned an offset not 64-bit aligned. To fix the already released information, the register was split into two 32-bit registers: PMRMSCL and PMRMSCU that must be accessed with 32-bit accesses (reference NVMe base spec 2.0 section 3.1.3.27 and 3.1.3.28).
  - Incompatible change
    - The PMCMSC register is split into two 32-bit registers.
  - References
    - ECN 006
- **Errata from ECN 007**
  - Required Changes
    - Invalid SGL Segment Descriptor status code may include an SGL Segment descriptor, or an SGL Last Segment descriptor contains an invalid address (e.g., an address that is not qword aligned) (reference NVMe base spec 2.0 section 3.3.3.2.1.1).
    - If the controller does not maintain the timestamp across the most recent Controller Level Reset, then the Timestamp field is cleared to 0h due to that Controller Level Reset (reference NVMe base spec 2.0 section 5.27.1.11).
    - If a volatile write cache is not supported or enabled, then a Flush command is successful even during a sanitize operation (reference NVM Command Set spec 1.0 section 7.1)
  - Incompatible Changes
    - Section 8.15.1 clarified that not all I/O commands shall be aborted when a sanitize operation is in progress as this conflicts with the Flush command. Section 8.3.1.5 specifies changes to the Flush command.
  - References
    - ECN 007
- **Errata from ECN 008**
  - Required Changes
    - Asymmetric Namespace Access Change was changed to remove redundant notifications due to detachment of a namespaces (reference NVMe base spec 2.0 section 5.2.1).
  - Incompatible Changes
    - The description of the Asymmetric Namespace Access Change Asynchronous Event Notification (03h) indicated that the controller shall not send that event for cases where the Namespace Attribute Change is already sent (to reduce duplicate Asynchronous Event Notifications for the same event – “as the Namespace Attribute Changed event is sent for these changes”). The text specifically mentions excluding the event if it was as a result of creation of the namespace or the deletion of the namespace. However, those two statements conflict. The general statement (to reduce duplicate Asynchronous Event Notifications was not fully expanded to list the exact events that caused the duplicate Asynchronous Event Notifications). In addition, a NS Creation didn't change the ANA Log page, so it would never have generated an ANA AEN anyway.



- The Namespace Attribute Changed event is sent for any change in any Identify Namespace data (including changes to CNS=0h data, CNS=02h data, and CNS=11h data). Therefore, the statement to reduce the duplicate Asynchronous Event Notifications and not send this the ANA Change Asynchronous Event Notifications if the Namespace Attribute Change event is sent would apply – which would prevent the ANA Change Asynchronous Event Notifications if it occurred as a result of an NS Create (CNS=11h change), an NS Attach (CNS=02h), a NS Detach (CNS=02h), or a NS Detach (CNS=11h change). However, the specific list does not include Attach or Detach.
  - This ECN clarifies this wording. This may be considered an incompatible change.
- **Non-MDTS Command Size Limits**
    - Required Changes
      - TP 4056d added the NVM Command Set Identify specific Controller data structure (required) and this TP defined the Verify Size Limit (VSL), Write Zeroes Size Limit (WZSL), Write Uncorrectable Size Limit (WUSL), Dataset Management Ranges Limit (DMRL), Dataset Management Range Size Limit (DMRSL), and Dataset Management Size Limit (DMSL) fields which need to be supported if the associated command is supported. The support of limiting sizes for these commands is optional, but supporting the new fields is a requirement to indicate if the command new limits are supported (reference NVM Command Set spec 1.0 section 4.1.5.4).
    - Reference
      - TP 4040a
- **Namespace Types**
    - Required Changes
      - A large set of changes that addressed the following issue: The Identify Namespace data structure was out of room and needed to support specific namespace and controller information per I/O command Set. To do that the following was done:
        - Many commands are updated to support a CSI value is 0h for NVM Command Set, but this new field is required to be validated to be 0h that was previously reserved. (2h for ZNS support)
        - Identify command has to support new CNS values to access the NVM Command Set Identify specific Namespace and Controller data structures. These are required to be supported.
        - The DW1 in the completion queue entry is no longer reserved as some commands are not using them.
        - If only supporting NVM command set then the new feature and some new CNS values are not required.
        - Namespace Identification Descriptor must include the Command Set Identifier
    - References
      - Technical Proposal 4056d
- **Telemetry Enhancements**



- Required Changes
  - Added the generation number to be incremented when a Telemetry Host-Initiated log is created (reference NVMe base spec 2.0 section 5.16.1.8).
- References
  - Technical Proposal 4063
- **Protection Information Enhancement**
  - Required Changes
    - The support for PI at the beginning of the metadata is no longer supported (reference NVM Command Set spec 1.0 section 4.1.5.1).
    - Must support the Extended LBA Formats in the NVM Command Set I/O Command Set Specific Identify Namespace Data Structure. Filled with 0h does select the PI supported by NVMe 1.4 (reference NVM Command Set spec 1.0 section 4.1.5.3).
    - Need to support the LBA Format Upper field if the Host Behavior Support feature is supported (reference NVMe base spec 2.0 section 5.14).
  - References
    - Technical Proposal 4068b
- **Command and Effect Log (CEL) Enhancements**
  - Required Changes
    - Adds mandatory log page of Supported log pages. In the future this is how the OS determines the log pages supported. No longer using the Identify Controller data structure (reference NVMe base spec 2.0 section 15.16.1.1).
    - Commands Supported and Effects log page added defining the scope of the commands (reference NVMe base spec 2.0 section 15.16.1.6).
    - Adds mandatory log page of supported Feature Identifiers Supported and the effects. In the future this is how the OS determines the FIDs supported. No longer using the Identify Controller data structure (reference 15.16.1.18).
    - Adds mandatory log page of supported in-band NVMe-MI commands using the NVMe-MI Send/Receive commands and the effects. In the future this is how the OS determines the in-band NVMe-MI commands supported. No longer using the Identify Controller data structure (reference NVMe base spec 2.0 section 15.16.1.19).
    - Required to set bit 5 in the Log Page Attributes field in the Identify Controller data structure (reference NVMe base spec 2.0 section 5.17.2).
  - References
    - Technical Proposal 4071b
- **Persistent Event Log Context Counter**
  - Required Changes
    - A Generation Number field is added to the Persistent Event Log Header to be incremented whenever the persistent event reporting contact is established reference NVMe base spec 2.0 section 5.16.1.4).
  - References
    - Technical Proposal 4083
- **Time To Ready Enhancements**

- Required Changes
  - This TP adds the ability for a controller to become ready before or after media is ready based on host request. This section assume controller is ready when media is ready is the baseline. CAP.TO is allowed to be modal in what is reported however the default is the time to ready to include namespaces being ready which I assume meets our definition. CAP.TO Controller Ready Modes supported bits are required to be supported. Bit 0 of that field is required to be set to '1'. the other bit setting is optional (reference NVMe base spec 2.0 section 3.1.3.1)
  - Required to add the Controller Ready Timeouts register at offset 68h (reference NVMe base spec 2.0 section 3.1.3.21)
  - Required to log the event if media not ready in specified reported time (reference NVMe base spec 2.0 section 5.16.1.14.1.5).
  - Required to support the Namespace Status filed in the I/O Command Set Independent Identify Namespace data structure (see TP 4105a) (reference NVMe base spec 2.0 section 5.17.2.8).
- References
  - Technical Proposal 4084
- **Persistent Event Log (PEL) State and Source Reporting**
  - Required Changes
    - Added a new Action value in CDW10 for the Persistent Event log page to establish a reporting contact and read the header information. Required to Add the Reporting Contact Information field to the Persistent Event Log Header (reference NVMe base spec 2.0 section 5.16.1.14).
    - Indicate if the new Action is supported in the Supported Log Pages for the Log Identifier 0Dh (reference NVMe base spec 2.0 section 15.16.1.1).
  - References
    - Technical Proposal 4086
- **IO Command Set Independent Data Structure**
  - Required Changes
    - The Namespace Attribute Changed notice occurs on changes to the new I/O Command Set Independent Identify Namespace data structure defined by TP 4105a (reference NVMe base spec 2.0 section 5.2.1).
    - New CNS 08h defined to access the new I/O Command Set Independent Identify Namespace data structure defined by TP 4105a. This namespace data structure is for fields not associated with an I/O command set. Some fields are replicated from the Identify Namespace data structure defined by the NVM Command Set Specification 1.0 (reference NVMe base spec 2.0 section 15.17.1).
    - The Changed Namespace List log page is updated if the new I/O Command Set Independent Identify Namespace data structure changes (reference NVMe base spec 2.0 section 5.16.1.5).
    - “If the Command Slot was paused and a More Processing Required Response has not yet been transmitted for the Command Message being processed, then the request-to-response timer shall be reset and restarted (refer to section 4.2.1.2 for details on the request-to-response time).”
  - References
    - Technical Proposal 6031



- **Transport Layer Security (TLS) 1.3 Support**
  - Required Changes
    - NVMe/TCP implementations that support TLS are required to support TLS 1.3 as specified in the NVMe TCP Transport Specification sections 3.6.1.1-3.6.1.6.
    - TLS 1.2 support is also allowed as specified in NVMe-oF 1.1, but TLS 1.2 support provisions have not been carried forward to the NVMe 2.0 family of specifications.
  - References
    - Technical Proposal 8011