



LEGAL NOTICE:

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

This NVM Express revision 1.4 technical proposal is proprietary to the NVM Express, Inc. (also referred to as “Company”) and/or its successors and assigns.

NOTICE TO USERS WHO ARE NVM EXPRESS, INC. MEMBERS: Members of NVM Express, Inc. have the right to use and implement this NVM Express revision 1.4 technical proposal subject, however, to the Member’s continued compliance with the Company’s Intellectual Property Policy and Bylaws and the Member’s Participation Agreement.

NOTICE TO NON-MEMBERS OF NVM EXPRESS, INC.: If you are not a Member of NVM Express, Inc. and you have obtained a copy of this document, you only have a right to review this document or make reference to or cite this document. Any such references or citations to this document must acknowledge NVM Express, Inc. copyright ownership of this document. The proper copyright citation or reference is as follows: “© 2007 – 2021 NVM Express, Inc. ALL RIGHTS RESERVED.” When making any such citations or references to this document you are not permitted to revise, alter, modify, make any derivatives of, or otherwise amend the referenced portion of this document in any way without the prior express written permission of NVM Express, Inc. Nothing contained in this document shall be deemed as granting you any kind of license to implement or use this document or the specification described therein, or any of its contents, either expressly or impliedly, or to any intellectual property owned or controlled by NVM Express, Inc., including, without limitation, any trademarks of NVM Express, Inc.

LEGAL DISCLAIMER:

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

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

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

NVM Express Technical Proposal for New Feature

Technical Proposal ID	4015b – NVMe KV
Change Date	4/8/2021
Builds on Specification	NVM Express 1.4

Technical Proposal Author(s)

Name	Company
Bill Martin	Samsung
Fred Knight	NetApp

Change History

08/19/2018	Initial revision
09/06/2018	Changed to use SGL instead of PRP for long keys
09/13/2018	<p>Changed to allow SGL or PRP for long key</p> <p>Removed Fused operation</p> <p>Modified Keyspace Attach and Keyspace Management commands</p> <p>Added Identify Controller Data structure with modifications for NVM Express Base specification</p> <p>Added Identify Keyspace Data Structure for inclusion into base specification</p>
09/27/2018	<p>Deleted sections that do not need modifications from the base specification</p> <p>Added sections at the bottom where the base specification needs modifications</p>
10/26/2018	Remove sections no longer needed as a result of NS Type TP
11/29/2018	Final revision with changes from base specification highlighted incorporated change request from conference calls
01/10/2018	<p>Removed change tracking of changes from base specification. Following this is only changes made during the process of updating this proposal.</p> <p>Incorporated changes that had been previously agreed to in comments from the 11/29/2018 version.</p>
1/31/2019	Added clause 6 with information on the operating model for NVMe KV
2/7/2019	<p>Updates from 1/31 conference call</p> <p>Moved KV formats back into Identify Namespace data structure</p>

2/13/2019	Added the data structure for the list command and comment resolution from the meeting on 2/7/2019. Modified the KV Format data structure in Identify.
2/27/2019	Added Fused operation description and FUSES support for the fused operations description. Modified the description of the key in the commands based on the value of the Fused Operation field in the command.
3/6/2019	Removed comments that were previously resolved, Renumbered the Identify Namespace data structure. Updated the Key Value Format descriptor.
3/28/2019	Added references to the following commands in base specification: <ul style="list-style-type: none"> a) Flush command b) Reservation Acquire command c) Reservation Register command d) Reservation Release command e) Reservation Report command. Modified the submission queue entry to remove fields not used by NVMe KV Updated “key” to “KV key” and “value” to “KV value” as appropriate.
4/18/2019	Updated comments to isolate the remaining Phase 2 issues and resolve agreed to comments
4/19/2019	Updates from 4/18/2019 meeting
4/23/2019	Removed change notation on agreed to changes Removed comments that have been resolved from previous meetings
5/2/2019	Limited keys to 16 bytes Removed Set Key command Removed fused operation Moved Key bytes [7:0] to CDW2 and CDW3 – this allows for future use of the fields in CDW12 and CDW13 in the base specification if need be. Both of these changes had ripple effects oin the document.
5/8/2019	Removed all resolved comments and change tracking from previous revision Added definition of read command and write command
5/9/2019	Commented from Technical work group and sub-team
5/10/2019	Removed change markups from 5/9/2019 revision Resolved remaining open issues as discussed in the sub-team meeting.
5/29/2019	Numerous editorial changes fixed based on input from Seagate Seagate additional comments attached with responses Work done to bring in line with the Namespace types proposal
5/31/2019	Comments from Seagate with resolution from working group meeting
6/19/2019	Removed resolved issues from 5/31 version Accepted all changes from previous revisions Pulled in ALL referenced definitions from 1.4
6/27/2019	Accepted previously reviewed changes

	Made consistent NVM Express Base Specification consistent
7/11/2019	Comments incorporated from 7/11 meeting
8/1/2019	Removed changes previously accepted Removed resolved comments Added proposed structure for KV key descriptors in Delete Command for review before incorporating elsewhere
8/15/2019	Cleaned up all accepted changes Made documentation of Command Dword 2 and Command Dword 3; and Command Dword 14 and Command Dword 15 consistent for commands.
3/11/2020	Updated format to the I/O Command Set template
3/20/2020	Moved some sections around No technical changes from previous revision
3/25/2020	Added Get Log Page section to specification (moved from Base specification modifications section) Added Set features section to specification (moved from Base specification modifications section) Deleted read command and write command definition Pulled in necessary log page definitions related to Host Read Commands and Host Write Commands Added reservations section
3/26/2020	Added new section "I/O Command Set specific definitions used in the NVMe Base specification" Moved changes back into Base specification changes for the log page modifications. Added new section for Base specification of "I/O Command Set specific definitions used in the NVMe Base specification"
4/3/2020	Moved submission queue section up in document Removed I/O from name Removed reference to 6.1 for the specification of the KV key
4/3/2020a	Added Copy command to User Data In Command and User Data Out Command for modification in the base specification. Added note in the Data Units Read and Data Units Written fields in the SMART Health log page Updated definitions to be full sentences.
4/8/2020	Removed all answered comments and change markup in preparation for member review.
5/14/2020	Member Review comment resolution
5/14/2020a	Clean Integration Ready document
06/07/2020	Integrated into the NVM Express Base Specification
11/12/2020	TP 4015a – removed one incorrect reference in 4.2.2 Replaced two occurrences of Retrieve with Retrieve

	<p>Duplicate table number 9 fixed</p> <p>Table X1 needs to be numbered in final specification</p> <p>Changes are highlighted for the member review version</p>
2/3/2021	TP4051b - Added necessary changes to get CSI value assigned in both the Key Value Command Set and in the NVMe Base specification
4/6/2021	Integrated into the NVMe Base Specification.
4/7/2021	Updated the CNS title and number in the header.
4/8/2021	Removed comments, accepted all changes, and converted references/cross-references to text.



NVM Express™

Key Value Command Set Specification

Revision 0.30

April 8, 2021

Please send comments to info@nvmexpress.org

SPECIFICATION DISCLAIMER

LEGAL NOTICE:

© 2021 NVM Express, Inc. ALL RIGHTS RESERVED.

This NVM Express Key Value Command Set specification revision 0.30 is proprietary to the NVM Express, Inc. (also referred to as “Company”) and/or its successors and assigns.

NOTICE TO USERS WHO ARE NVM EXPRESS, INC. MEMBERS: Members of NVM Express, Inc. have the right to use and implement this NVM Express Key Value Command Set specification 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: “© 2021 NVM Express, Inc. ALL RIGHTS RESERVED.” When making any such citations or references to this document you are not permitted to revise, alter, modify, make any derivatives of, or otherwise amend the referenced portion of this document in any way without the prior express written permission of NVM Express, Inc. Nothing contained in this document shall be deemed as granting you any kind of license to implement or use this document or the specification described therein, or any of its contents, either expressly or impliedly, or to any intellectual property owned or controlled by NVM Express, Inc., including, without limitation, any trademarks of NVM Express, Inc.

LEGAL DISCLAIMER:

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

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

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

Table of Contents

1	INTRODUCTION	10
1.1	Overview.....	10
1.2	Scope.....	10
1.3	Conventions.....	10
1.4	Definitions.....	11
1.5	References	11
2	KEY VALUE COMMAND SET OVERVIEW.....	12
2.1	Key Value Namespace Size implications	12
2.2	Key Size implications.....	12
2.3	Submission Queue Entry – Command Format.....	12
2.4	Completion Queue Entry	13
2.5	Key Value Command Set Commands.....	15
2.6	Delete command	16
2.7	List command	17
2.8	Exist command	17
2.9	Store command	17
2.10	Retrieve command.....	17
3	KEY VALUE COMMAND SET SPECIFIC DATA STRUCTURES.....	18
3.1	Identify Data Structures	18
4	KEY VALUE COMMAND SET SPECIFIC BEHAVIOR FOR ADMIN COMMANDS	20
4.1	Asynchronous Event Request command	20
4.2	Format NVM command – Key Value Command Set Specific	20
4.3	Get Log Page command.....	21
4.4	Reservations.....	22
4.5	Sanitize command – Key Value Command Set Specific.....	23
4.6	Set Features command	23
5	KEY VALUE COMMAND SET	25
5.1	Namespaces.....	25
5.2	Command Ordering Requirements	25
5.3	Atomic Operation.....	25
5.4	Delete command	25
5.5	List command	26
5.6	Retrieve command	27
5.7	Exist command	29
5.8	Store command	29
5.9	Key Value Configuration (Feature Identifier 20h), (Mandatory for Key Value Command Set).....	32

1 Introduction

1.1 Overview

NVM Express™ (NVMe™) Base specification revision 2.0 defines an interface for host software to communicate with a non-volatile memory subsystems over a variety of memory based transports and message based transports. The NVM Express Base specification revision 2.0 is referred to as the NVMe Base specification.

This document defines a specific I/O Command Set: Key Value Command Set which extends the NVMe Base specification.

The Key Value Command Set host controller interface is designed to address the needs of Enterprise and Client systems that utilize key value drives. The interface utilizes the command submission and completion paths defined in the NVMe Base specification.

The interface has the following attributes:

- Provides a KV key associated with a KV value that is stored, retrieved, or deleted from the media;

This specification along with the NVMe Base specification and transport specific specifications (e.g., NVMe over PCIe, NVMe over Fabrics) defines a streamlined set of registers whose functionality includes:

- Indication of controller capabilities;
- Status for controller failures (command status is processed via CQ directly);
- Admin Queue configuration (I/O Queue configuration processed via Admin commands); and
- Doorbell registers for scalable number of Submission and Completion Queues.

This specification defines the Key Value Command Set.

The host selects the I/O Command Set that is used for a namespace.

1.2 Scope

This specification defines additional Data Structures, Features, log pages, commands, and status values. This specification also defines extensions to existing Data Structures, Features, log pages, commands, and status values. This specification supplements the NVMe Base specification.

Unless otherwise specified, this specification conforms to the System Bus (PCI Express) Registers section, Control Registers section, Admin Command Set section, Controller Architecture section, and Namespaces section of the NVMe Base specification.

This specification defines requirements and behaviors that are specific to the Key Value Command Set. Functionality that is applicable generally to NVMe or that is applicable across multiple I/O Command Sets is defined in the NVMe Base specification.

Other published specifications referred to in this document, even if required for compliance, are outside the scope of this specification; this includes published specifications for NVMe transports, NVMe Management Interface, and other technologies referred to by this specification.

1.3 Conventions

This specification conforms to the Conventions section, Keywords section, and Byte, Word, and Dword Relationships section of the NVMe Base specification.

1.4 Definitions

1.4.1 Definitions from the NVMe Base specification

This specification uses the definitions in the NVMe Base specification.

1.4.2 Definitions specific to the Key Value Command Set

This section defines terms that are specific to this specification.

1.4.2.1 key value pair

An associated KV key and KV value that may be stored on media where the KV key identifies the associated KV value.

1.4.2.2 KV key

The part of a key value pair that is used to identify that key value pair.

1.4.2.3 KV value

The value that is associated with a key value pair.

1.4.3 I/O Command Set specific definitions used in the NVMe Base specification

1.4.3.1 User Data In Command

A Retrieve command.

1.4.3.2 User Data Out Command

A Store command.

1.5 References

NVMe Base specification, Revision 1.4. 10 June 2019 Available from <http://www.nvmexpress.org>.

2 Key Value Command Set Overview

The NVMe Base specification defines a register level interface for host software to communicate with a non-volatile memory subsystem. This specification defines additional functionality for the Key Value Command Set.

Each I/O Command Set is assigned a specific Command Set Identifier (CSI) value by the NVMe Base Specification. The Key Value Command Set is assigned a CSI value of 01h.

An NVM subsystem may contain controllers that implement the Key Value Command Set. Key Value storage is measured in bytes. The amount of storage required to store a key value pair is the sum of the KV key size and the KV value size. A KV value is allowed to have a length of zero bytes. Supported KV key and KV value sizes are reported in the I/O Command Set specific Identify Namespace data structure for the Key Value Command Set.

A device that implements the Key Value Command Set provides access to data identified by a KV key. The KV key may be variable length and the length of the KV key is specified in the command. Two KV keys that have different lengths are not the same. The KV value that is associated with a KV key has a length in bytes that is specified in the command that stores that KV value. The length in bytes of a KV value is indicated in the response to a query about that KV value (e.g., Retrieve command, Exist command). The length in bytes of a KV key is indicated in the response to a List command that returns that KV key.

While a controller may perform operations (e.g., compression) on data before the data is stored on the media and perform the reverse of that operation (e.g., decompression) when retrieving the data from the media, this functionality is outside of the scope of this specification.

The maximum size of any KV key and the maximum size of any KV value in a namespace is specified when the namespace is formatted and is selected from the matrix of KV formats in the I/O Command Set specific Identify Namespace data structure.

2.1 Key Value Namespace Size implications

The number of bytes required to store a key value pair is related to the KV key size and the KV value size. Supported KV key sizes and KV value sizes are reported in the KV Format data structures in the Identify Namespace data structure.

The number of bytes required to store a given key value pair is greater than or equal to the sum of the size of the KV key and the size of the KV value. Namespace Size and Namespace Utilization reflect the number of bytes required to store the KV value and KV key.

2.2 Key Size implications

The maximum KV key size is 16 bytes.

The KV key is specified in Command Dword 2, Command Dword 3, Command Dword 14, and Command Dword 15.

If a command specifies a KV key size greater than 16 bytes, that command is aborted with Invalid Field in Command.

2.3 Submission Queue Entry – Command Format

The Submission Queue Entry (SQE) structure and the fields that are common to all I/O Command Sets are defined in the NVMe Base specification.

Each command is 64 bytes in size.

Command Dword 0, Namespace Identifier, Metadata Pointer, PRP Entry 1, PRP Entry 2, SGL Entry 1, and Metadata SGL Segment Pointer have common definitions for all Admin commands and I/O commands and are described in the Submission Queue Entry – Command Format section in the NVMe Base specification.

The Key Value Command Set uses the Common Command Format as defined in the NVMe Base specification.

The command format for the Key Value Command set is defined in Figure 1.

Figure 1: Command Format – Key Value Command Set

Bytes	Description
03:00	Command Dword 0 (CDW0): This field is described in the NVMe Base specification.
07:04	Namespace Identifier (NSID): This field is described in the NVMe Base specification.
11:08	Command Dword 2 (CDW2): This field contains KV KEY [bytes 3:0].
15:12	Command Dword 3 (CDW3): This field contains KV KEY [bytes 7:4].
23:16	Reserved for MPTR
39:24	Data Pointer (DPTR): This field is described in the NVMe Base specification.
43:40	Command Dword 10 (CDW10): This field is command specific Dword 10.
47:44	Command Dword 11 (CDW11): This field is command specific Dword 11.
51:48	Command Dword 12 (CDW12): This field is command specific Dword 12.
55:52	Command Dword 13 (CDW13): This field is command specific Dword 13..
59:56	Command Dword 14 (CDW14): This field contains KV KEY [bytes 11:8].
63:60	Command Dword 15 (CDW15): This field contains KV KEY [bytes 15:12].

2.4 Completion Queue Entry

The Completion Queue Entry (CQE) structure and the fields that are common to all I/O Command Sets are defined in the NVMe Base specification. For the Key Value Command Set, additional status codes are defined in this subclause. Status codes not defined in this specification are defined in the NVMe Base specification.

2.4.1.1 Key Value Command Set - command Specific Status Definitions

Completion queue entries with a Status Code Type of Command Specific Errors indicate an error that is specific to a particular command opcode. Status codes of 00h to 7Fh are for Admin command errors. Refer to NVMe Base specification. Status codes of 80h to BFh are specific to a specific I/O Command Set. The status codes that are specific to the Key Value Command Set are defined in Figure 8.

Figure 8: Key Value Generic Command Status

Value	Description	Commands Affected
81h	Capacity Exceeded: Refer to NVMe Base specification	Store
82h	Namespace Not Ready: Refer to NVMe Base specification	Delete, Exist, Retrieve, Store
83h	Reservation Conflict: Refer to NVMe Base specification	Delete, Store, Retrieve
84h	Format In Progress: Refer to NVMe Base specification	Delete, Exist, List, Retrieve, Store
85h	Invalid Value Size: The KV value size is not valid	Store

Figure 9: Key Value Command Set Command Specific status codes

Value	Description	Commands Affected
85h	Invalid Value Size: The KV value size is not valid	Store
86h	Invalid Key Size: The KV key size is not valid	List, Retrieve, Store
87h	KV Key Does Not Exist: The KV key does not exist	Delete, Exist, Retrieve, Store
88h	Unrecovered Error: There was an unrecovered error when reading from the medium	Retrieve
89h	Key Exists: Store Option bit 9 is set to '1' and the KV key exists	Store

2.5 Key Value Command Set Commands

The Submission Queue Entry (SQE) structure and the fields that are common to all NVMe commands are defined in the NVMe Base specification and in section 2.3. The Completion Queue Entry (CQE) structure and the fields that are common to all NVMe commands are defined in the NVMe Base specification. The command specific fields in the SQE and CQE structures (i.e., SQE Command Dwords 10-15 and CQE Dword 0) for the Key Value Command Set are defined in this section.

The Key Value Command Set includes the commands listed in [Figure x1](#). Section 3 describes the definition for each of the commands defined by this specification. Commands shall only be submitted by the host when the controller is ready as indicated in the Controller Status register (CSTS.RDY) and after appropriate I/O Submission Queue(s) and I/O Completion Queue(s) have been created.

Figure X1: Opcodes for Key Value Command Set Commands

Opcode by Field			Combined Opcode ¹		Command ²	Reference
(07)	(06:02)	(01:00)				
Standard Command	Function	Data Transfer ³				
0b	000 00b	00b	00h		Flush	NVM Command Set
0b	000 11b	01b	0Dh		Reservation Register	NVM Command Set
0b	000 11b	10b	0Eh		Reservation Report	NVM Command Set
0b	001 00b	01b	11h		Reservation Acquire	NVM Command Set
0b	001 01b	01b	15h		Reservation Release	NVM Command Set
0b	000 00b	01b	01h		Store	5.8
0b	000 00b	10b	02h		Retrieve	5.6
0b	001 00b	00b	10h		Delete	5.4
0b	001 01b	00b	14h		Exist	5.7
0b	000 01b	10b	06h		List	5.5

NOTES:

1. Opcodes not listed are defined in the NVMe Base specification.
2. All Key Value Command Set commands use the Namespace Identifier (NSID) field. The value FFFFFFFFh is not supported in this field unless footnote 4 in this figure indicates that a specific command does support that value.
3. Indicates the data transfer direction of the command. All options to the command shall transfer data as specified or transfer no data. All commands, including vendor specific commands, shall follow this convention: 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional.
4. This command may support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh.

2.6 Delete command

The Delete command requests the controller to delete the specified key value pair from the namespace.

2.7 List command

The host may request a list of the KV keys in that namespace. This is accomplished using the List command. The KV keys in the data structure returned from the List command are not in any specified order, but in the absence of Store and Delete commands the order of the KV keys in the list shall be constant. The KV key that is sent in that command specifies the starting point in the list of KV keys. If that KV key exists, then that KV key is the first key returned in the datastructure. If that KV key does not exist, then the device returns KV keys where the first KV key returned is vendor specific, but in the absence of Store and Delete commands the first KV key returned shall not change.

2.8 Exist command

The Exist command is used to determine if a specified KV key exists in the namespace. The existence of the KV key is indicated by the value returned in the CQE for that command.

2.9 Store command

The Store command is used to store a key value pair to the namespace. The length of the KV value is specified in the Store command and the location of the KV value to be stored is specified by either the SGL or the PRP in the command. The Store command is an atomic command and following a Store command of a key value pair that existed prior to that command, a Retrieve command returns either all of the previous KV value or all of the KV value in that Store command but shall not return a combination of previous data and data from that Store command.

2.10 Retrieve command

The Retrieve command is used to retrieve a key value pair from the namespace. The length to be retrieved of the KV value is specified in the Retrieve command and the location to transfer the KV value to is specified by either the SGL or the PRP in the command. If the length specified in the command is less than the length of the KV value that is being retrieved, then the device returns the requested amount and the length of the KV value is returned in the CQE. If the length specified in the command is greater than the length of the KV value that is being retrieved, then the device returns the data from the media and the length of that KV value is returned in the CQE.

3 Key Value Command Set Specific data structures

3.1 Identify Data Structures

This specification implements the Identify Data structures defined in the NVMe Base specification. Additionally, the Key Value Command Set specifies the data structures defined in this section.

3.1.1 I/O Command Set specific Identify Namespace data structure (CNS 05h, CSI 01h)

The I/O Command Set specific Identify Namespace data structure (i.e., CNS 05h) for the Key Value Command Set is defined in Figure 2.

Figure 2: Identify – Identify Namespace Data Structure, Key Value Type Specific

Bytes	O/M ¹	Description
07:00	M	Namespace Size (NSZE): This field indicates the total size of the namespace in bytes. This is the space to store KV keys and KV values. This field is undefined prior to the namespace being formatted.
15:08		Reserved
23:16	M	<p>Namespace Utilization (NUSE): This field indicates the current number of bytes allocated in the namespace. This is the space to store KV keys and KV values. This field is smaller than or equal to the Namespace Size.</p> <p>A key value pair is allocated when it is written with a Store command. A key value pair is deallocated using the Delete command.</p> <p>If the controller supports Asymmetric Namespace Access Reporting (refer to the CMIC field), and the relationship between the controller and the namespace is in the ANA Inaccessible state (refer to the NVMe Base specification) or the ANA Persistent Loss state (refer to the NVMe Base specification), then this field shall be cleared to 0h.</p>
24	M	<p>Namespace Features (NSFEAT): This field defines features of the namespace.</p> <p>Bits 7:4 are reserved.</p> <p>Bit 3 if set to '1' indicates that the non-zero NGUID and non-zero EUI64 fields for this namespace are never reused by the controller. If cleared to '0', then the NGUID and EUI64 values may be reused by the controller for a new namespace created after this namespace is deleted. This bit shall be cleared to '0' if both NGUID and EUI64 fields are cleared to 0h. Refer to the NVMe Base specification.</p> <p>Bits 2:0 are reserved.</p>
25	M	Number of KV Formats (NKVF): This field defines the number of KV format descriptors supported by the namespace. KV formats shall be allocated in order (starting with 0) and packed sequentially. This is a 0's based value. The maximum number of KV formats that may be indicated as supported is 16. The supported KV formats are indicated in bytes 72 to 327 in this data structure. The KV Format fields with an index beyond the value set in this field are invalid and not supported. KV Formats that are valid, but not currently available may be indicated by setting the KV Key Max and KV Value Max both to 0000h for that KV Format.
26	O	Namespace Multi-path I/O and Namespace Sharing Capabilities (NMIC): Refer to the NMIC field in the Identify Namespace data structure in the NVMe Base specification.
27	O	Reservation Capabilities (RESCAP): Refer to NVMe Base specification.
28	O	Format Progress Indicator (FPI): Refer to NVMe Base specification.
31:29		Reserved
35:32	O	Namespace Optimal Value Granularity (NOVG): This field indicates the optimal value granularity for this namespace. This field is specified in bytes. The host should construct Store commands that store multiples of NOVG bytes to achieve optimal performance. A value of 0h indicates that no optimal value granularity is reported.

Figure 2: Identify – Identify Namespace Data Structure, Key Value Type Specific

Bytes	O/M ¹	Description
39:36	O	ANA Group Identifier (ANAGRPID): Refer to NVMe Base specification.
42:40		Reserved
43	O	Namespace Attributes (NSATTR): Refer to NVMe Base specification.
45:44	O	NVM Set Identifier (NVMSETID): Refer to NVMe Base specification.
47:46	O	Endurance Group Identifier (ENDGID): Refer to NVMe Base specification.
63:48	O	Namespace Globally Unique Identifier (NGUID): Refer to NVMe Base specification.
71:64	O	IEEE Extended Unique Identifier (EUI64): Refer to NVMe Base specification.
87:72	M	KV Format 0 Support (KVF0): This field indicates the KV format 0 that is supported by the controller. The KV format field is defined in Figure 1.
103:88	O	KV Format 1 Support (KVF1): This field indicates the KV format 1 that is supported by the controller. The KV format field is defined in Figure 1.
119:104	O	KV Format 2 Support (KVF2): This field indicates the KV format 2 that is supported by the controller. The KV format field is defined in Figure 1.
135:120	O	KV Format 3 Support (KVF3): This field indicates the KV format 3 that is supported by the controller. The KV format field is defined in Figure 1.
151:136	O	KV Format 4 Support (KVF4): This field indicates the KV format 4 that is supported by the controller. The KV format field is defined in Figure 1.
167:152	O	KV Format 5 Support (KVF5): This field indicates the KV format 5 that is supported by the controller. The KV format field is defined in Figure 1.
183:168	O	KV Format 6 Support (KVF6): This field indicates the KV format 6 that is supported by the controller. The KV format field is defined in Figure 1.
199:184	O	KV Format 7 Support (KVF7): This field indicates the KV format 7 that is supported by the controller. The KV format field is defined in Figure 1.
215:200	O	KV Format 8 Support (KVF8): This field indicates the KV format 8 that is supported by the controller. The KV format field is defined in Figure 1.
231:216	O	KV Format 9 Support (KVF9): This field indicates the KV format 9 that is supported by the controller. The KV format field is defined in Figure 1.
247:232	O	KV Format 10 Support (KVF10): This field indicates the KV format 10 that is supported by the controller. The KV format field is defined in Figure 1.
263:248	O	KV Format 11 Support (KVF11): This field indicates the KV format 11 that is supported by the controller. The KV format field is defined in Figure 1.
279:264	O	KV Format 12 Support (KVF12): This field indicates the KV format 12 that is supported by the controller. The KV format field is defined in Figure 1.
295:280	O	KV Format 13 Support (KVF13): This field indicates the KV format 13 that is supported by the controller. The KV format field is defined in Figure 1.
311:296	O	KV Format 14 Support (KVF14): This field indicates the KV format 14 that is supported by the controller. The KV format field is defined in Figure 1.
327:312	O	KV Format 15 Support (KVF15): This field indicates the KV format 15 that is supported by the controller. The KV format field is defined in Figure 1.
3839:328		Reserved
4095:3840	O	Vendor Specific

NOTES:

1. O/M definition: O = Optional, M = Mandatory.

Figure 3: Identify – KV Format Data Structure

Bytes	Description										
01:00	KV Key Max Length: Maximum length of a KV key in a key value pair in bytes. The maximum value that is supported by the commands in this version of the Key Value Command Set specification is 16.										
02	Reserved										
03	<p>Additional format options:</p> <p>Bits 7:2 Reserved</p> <p>Bits 1:0 Relative Performance (RP): This field indicates the relative performance of the KV format indicated, relative to other KV formats supported by the controller. Depending on the characteristics of the format, there may be performance implications. The performance analysis is based on better performance on a queue depth of 32 with(<i>need words here</i>). The meanings of the values indicated are included in the following table.</p> <table> <tr> <th>Value</th><th>Definition</th></tr> <tr> <td>00b</td><td>Best performance</td></tr> <tr> <td>01b</td><td>Better performance</td></tr> <tr> <td>10b</td><td>Good performance</td></tr> <tr> <td>11b</td><td>Degraded performance</td></tr> </table>	Value	Definition	00b	Best performance	01b	Better performance	10b	Good performance	11b	Degraded performance
Value	Definition										
00b	Best performance										
01b	Better performance										
10b	Good performance										
11b	Degraded performance										
07:04	KV Value Max Length: Maximum length in bytes of a KV value in a key value pair.										
11:08	Max Num Keys: Maximum number of KV keys allowed in the namespace. A value of zero indicates that no maximum number is indicated.										
15:12	Reserved										

3.1.2 Identify I/O Command Set specific Controller data structure (CNS 06h, CSI 01h)

The Key Value Command Set does not have an Identify I/O Command Set specific Controller data structure (i.e., CNS 06h). The controller shall return a zero filled data structure for this CNS value.

4 Key Value Command Set specific behavior for Admin Commands

The Admin Commands are as defined in the NVMe Base specification. The Key Value Command Set specific behavior for Admin Commands is described in this section.

4.1 Asynchronous Event Request command

The Asynchronous Event Request command operates as defined in the NVMe Base specification. In addition to the Asynchronous Events defined in the NVMe Base specification, the Key Value Command Set does not define any additional Asynchronous Events.

4.2 Format NVM command – Key Value Command Set Specific

The Format NVM command operates as defined in the NVMe Base specification. The Format Index points to a valid KV Format in the I/O Command Set Specific, Identify Namespace data structure.

4.2.1 Command Completion

Command Completion is as defined in the NVMe Base specification.

4.3 Get Log Page command

The Get Log Page command operates as described in the NVMe Base specification. In addition to the requirements in the NVMe Base specification, Mandatory, optional, and prohibited Log Identifiers are defined in Figure 4. If a Get Log Page command is processed that specifies a Log Identifier that is not supported, then the controller should abort the command with status Invalid Field in Command.

4.3.1 Log Specific Information

Figure 4 defines the Log pages that may be retrieved with the Get Log Page command and the scope of the information that is returned in those Log pages for the Key Value Command Set.

Log page scope is as defined in the NVMe Base specification.

The rules for namespace identifier usage are specified in the NVMe Base specification.

Figure 4: Get Log Page – Log Page Identifiers

Log Identifier	Scope	Log Page Name	Reference	Support Requirement ⁸
00h	Reserved			
01h	Controller	Error Information	NVMe Base specification	M
02h	NVM subsystem ¹	SMART / Health Information	NVMe Base specification	O
	Namespace ²			M
03h	Domain/NVM subsystem i.a.i.6	Firmware Slot Information	NVMe Base specification	M
04h	Controller	Changed Namespace List	NVMe Base specification	O
05h	Controller	Commands Supported and Effects	NVMe Base specification	O

Figure 4: Get Log Page – Log Page Identifiers

Log Identifier	Scope	Log Page Name	Reference	Support Requirement ⁸
06h	Controller ³	Device Self-test ⁵	NVMe Base specification	O
	Domain/NVM subsystem ⁴ , i.a.i.6			O
07h	Vendor Specific	Telemetry Host-Initiated ⁵	NVMe Base specification	O
08h	Vendor Specific	Telemetry Controller-Initiated ⁵	NVMe Base specification	O
09h	Domain/NVM subsystem i.a.i.6	Endurance Group Information	NVMe Base specification	O
0Ah	Domain/NVM subsystem i.a.i.6	Predictable Latency Per NVM Set	NVMe Base specification	O
0Bh	Domain/NVM subsystem i.a.i.6	Predictable Latency Event Aggregate	NVMe Base specification	O
0Ch	Controller	Asymmetric Namespace Access	NVMe Base specification	O
0Dh	NVM subsystem	Persistent Event Log ⁵	NVMe Base specification	O
0Eh	Controller	LBA Status Information	NVMe Base specification	P
0Fh	Domain/NVM subsystem i.a.i.6	Endurance Group Event Aggregate	NVMe Base specification	O
10h	Domain/NVM subsystem i.a.i.5, i.a.i.6	Media Unit Status	NVMe Base specification	M ⁷
11h	Domain/NVM subsystem i.a.i.6	Supported Capacity Configuration List	NVMe Base specification	M ⁷
12h to 6Fh	Reserved			
70h	Discovery (refer to the NVMe over Fabrics specification)			
71h to 7Fh	Reserved for NVMe over Fabrics implementations			
80h	Controller	Reservation Notification	NVMe Base specification	O
81h	NVM subsystem	Sanitize Status	NVMe Base specification	O
82h to BFh	I/O Command Set Specific			

Figure 4: Get Log Page – Log Page Identifiers

Log Identifier	Scope	Log Page Name	Reference	Support Requirement ⁸
C0h to FFh	Vendor specific ⁵			
KEY: Namespace = The log page contains information about a specific namespace. Controller = The log page contains information about the controller that is processing the command. NVM subsystem = The log page contains information about the NVM subsystem. Vendor Specific = The log page contains information that is vendor specific.				
NOTES: 1. For namespace identifiers of 0h or FFFFFFFFh. 2. For namespace identifiers other than 0h or FFFFFFFFh. 3. Bit 0 is cleared to '0' in the DSTO field in the Identify Controller data structure (refer to the NVMe Base specification). 4. Bit 0 is set to '1' in the DSTO field in the Identify Controller data structure. 5. Selection of a UUID may be supported. Refer to the NVMe Base specification. 6. For NVM subsystems that support multiple domains (refer to the MDS bit in the Identify Controller data structure), Domain scope information is returned. 7. Optional for controllers that do not support Fixed Capacity Management. 8. O = Optional, M = Mandatory, P = Prohibited				

4.4 Reservations

Reservations operate as defined in 8.8 of the NVMe Base specification with the additional Command Behavior in the Presence of a Reservation defined in Figure 487

Figure 5: Command Behavior in the Presence of a Reservation

NVMe Command	Write Exclusive Reservation		Exclusive Access Reservation		Write Exclusive Registrants Only or Write Exclusive All Registrants Reservation		Exclusive Access Registrants Only or Exclusive Access All Registrants Reservation	
	Non-Registrant	Registrant	Non-Registrant	Registrant	Non-Registrant	Registrant	Non-Registrant	Registrant
Key Value Read Command Group								
Retrieve	A	A	C	C	A	A	C	A
Key Value Write Command Group								

Figure 5: Command Behavior in the Presence of a Reservation

NVMe Command	Write Exclusive Reservation		Exclusive Access Reservation		Write Exclusive Registrants Only or Write Exclusive All Registrants Reservation		Exclusive Access Registrants Only or Exclusive Access All Registrants Reservation	
	Non-Registrant	Registrant	Non-Registrant	Registrant	Non-Registrant	Registrant	Non-Registrant	Registrant
Delete Flush Format NVM (Admin) Namespace Attachment (Admin) Namespace Management (Admin) Sanitize (Admin) Security Send (Admin) Store	C	C	C	C	C	A	C	A
Key: A definition: A=Allowed, command processed normally by the controller C definition: C=Conflict, command aborted by the controller with status Reservation Conflict								

4.5 Sanitize command – Key Value Command Set Specific

The Sanitize command operates as described in the NVMe Base specification. In addition to the requirements in the NVMe Base specification, the following Key Value Command Set commands shall be aborted if a sanitize operation is in progress:

- Delete;
- List;
- Retrieve;
- Exist; and
- Store.

4.6 Set Features command

Figure 12 defines the Features support requirements for the Key Value Command Set.

Figure 6: Set Features – Feature Identifiers

Feature Identifier	O/M/P ⁶	Persistent Across Power Cycle and Reset ²	Uses Memory Buffer for Attributes	Description
00h				Reserved
01h	M	No	No	Arbitration
02h	M	No	No	Power Management
03h	P		Yes	LBA Range Type

Figure 6: Set Features – Feature Identifiers

Feature Identifier	O/M/P ⁶	Persistent Across Power Cycle and Reset ²	Uses Memory Buffer for Attributes	Description
04h	M	No	No	Temperature Threshold
05h	P		No	Error Recovery
06h	O	No	No	Volatile Write Cache
07h	M	No	No	Number of Queues
08h	NOTE 5	No	No	Interrupt Coalescing
09h	NOTE 5	No	No	Interrupt Vector Configuration
0Ah	M	No	No	Prohibited
0Bh	M	No	No	Asynchronous Event Configuration
0Ch	O	No	Yes	Autonomous Power State Transition
0Dh	O	No ³	No ⁴	Host Memory Buffer
0Eh	O	No	Yes	Timestamp
0Fh	O	No	No	Keep Alive Timer
10h	O	Yes	No	Host Controlled Thermal Management
11h	O	No	No	Non-Operational Power State Config
12h	O	Yes	No	Read Recovery Level Config
13h	O	Yes	Yes	Predictable Latency Mode Config
14h	O	Yes	No	Predictable Latency Mode Window
15h	P	No	No	LBA Status Information Report Interval
16h	M	No	Yes	Host Behavior Support
17h	O	Yes	No	Sanitize Config
18h	O	No	No	Endurance Group Event Configuration
20h	M	Yes	No	Key Value Configuration
21h to 77h				Reserved
78h to 7Fh		Refer to the NVMe Management Interface Specification for definition.		
80h to BFh				Command Set Specific (Reserved)
C0h to FFh				Vendor Specific ¹

NOTES:

1. The behavior of a controller in response to an inactive namespace ID to a vendor specific Feature Identifier is vendor specific.
2. This column is only valid if the feature is not saveable (refer to the NVMe Base specification). If the feature is saveable, then this column is not used and any feature may be configured to be saved across power cycles and reset.
3. The controller does not save settings for the Host Memory Buffer feature across power states and reset events, however, host software may restore the previous values. Refer to the NVMe Base specification.
4. The feature does not use a memory buffer for Set Features, but it does use a memory buffer for Get Features. Refer to the NVMe Base specification.
5. The feature is mandatory for NVMe over PCIe implementations. This feature is not supported for NVMe over Fabrics implementations.
6. O/M/P: O = Optional, M = Mandatory, P=Prohibited.

5 Key Value Command Set

5.1 Namespaces

A namespace is a collection of NVM and is as defined in the NVMe Base specification.

5.2 Command Ordering Requirements

For all commands, each command is processed as an independent entity without reference to other commands submitted to the same I/O Submission Queue or to commands submitted to other I/O Submission Queues. Specifically, the controller is not responsible for checking the KV key of a Retrieve or Store command to ensure any type of ordering between commands. For example, if a Retrieve command is submitted for KV key *x* and there is a Store command also submitted for KV key *x*, then there is no guarantee of the order of completion for those commands (the Retrieve command may finish first or the Store command may finish first). If there are ordering requirements between these commands, host software or the associated application is required to enforce that ordering above the level of the controller.

5.3 Atomic Operation

All Store commands and Delete commands are atomic with respect to the associated key value pair.

5.4 Delete command

The Delete command deletes the KV key and the associated KV value for the specified KV namespace.

The command uses Command Dword 2, Command Dword 3, Command Dword 11, Command Dword 14, and Command Dword 15 fields. All other command specific fields are reserved.

If the value in the Key Length field is greater than 16, then the controller shall abort the command with Invalid Field in Command.

Figure 7: Delete – Command Dword 11

Bit	Description
31:8	Reserved
7:0	Key Length (KL): Specifies the length of the KV key in bytes.

Figure 8: Delete – Command Dword 2 and Command Dword 3

Bit	Description
63:0	KV key[63:00]: This field specifies 64-bits of the KV Key to be used for the command. Command Dword 2 contains bits 31:00; Command Dword 3 contains bits 63: 32.

Figure 9: Delete – Command Dword 14 and Command Dword 15

Bit	Description
63:0	KV key[127:64]: This field specifies 64-bits of the KV Key to be used for the command. Command Dword 14 contains bits 95:64; Command Dword 15 contains bits 127: 96.

5.4.1 Command Completion

Upon completion of the Delete command, the controller posts a completion queue entry (CQE) to the associated I/O Completion Queue. If the status code returned is 00h, then the KV key and its associated KV value have been deleted.

Delete command specific status values are defined in Figure 15

Figure 10: Delete – Command Specific Status Values

Value	Description
87h	KV Key Does Not Exist: The KV key does not exist
0Bh	Invalid Namespace or Format: The namespace or the format of that namespace is invalid or the namespace is not associated with the KV Command Set..

5.5 List command

The List command retrieves a list of KV keys that exist for the specified KV namespace starting at the KV key specified.

The command uses Command Dword 2, Command Dword 3, Command Dword 10, Command Dword 11, Command Dword 14, and Command Dword 15 fields. If the command uses PRPs for the data transfer, then the PRP Entry 1, and PRP Entry 2 fields are used. If the command uses SGLs for the data transfer, then the SGL Entry 1 field is used

If the value in the Key Length field is greater than 16, then the controller shall abort the command with Invalid Field in Command.

Figure 11: List – Command Dword 10

Bit	Description
31:00	Host Buffer Size (HBS): This field indicates the host buffer size in bytes.

Figure 12: List – Command Dword 11

Bit	Description
31:8	Reserved
7:0	Key Length (KL): Specifies the length of the KV key in bytes.

Figure 13: List – Command Dword 2 and Command Dword 3

Bit	Description
63:0	KV key[63:00]: This field specifies 64-bits of the KV Key to be used for the command. Command Dword 2 contains bits 31:00; Command Dword 3 contains bits 63: 32.

Figure 14: List – Command Dword 14 and Command Dword 15

Bit	Description
63:0	KV key[127:64]: This field specifies 64-bits of the KV Key to be used for the command. Command Dword 14 contains bits 95:64; Command Dword 15 contains bits 127: 96.

5.5.1 Command Completion

Upon completion of the List command, the controller shall post a completion queue entry to the associated I/O Completion Queue indicating the status for the command.

The command returns a list of KV keys that exist as described in 5.5.2

List command specific status values are defined in Figure 15

Figure 15: List – Command Specific Status Values

Value	Description
86h	Invalid Key Size: The KV key size is not valid
0Bh	Invalid Namespace or Format: The namespace or the format of that namespace is invalid.

5.5.2 List command return data structure

The data structure returned for the list command is as defined in Figure 23

Figure 16: List – Return data structure

Byte	Description
03:00	Number of Returned Keys (NRK): This value reflects how many KV keys are returned in this data structure
	Key data structure 1 (refer to Figure)
	Key data structure 2 (refer to Figure)
	...
	Key data structure n (refer to Figure)

Figure 24 Key data structure

Byte	Description
01:00	Key Length (KL): indicates the length of the KV key in bytes that this data structure represents
n:02	Key: KV key that this entry describes
m:n	Pad: Pad necessary, if any to end the data structure on a 4 byte boundary

5.6 Retrieve command

The Retrieve command retrieves a KV value from the NVM KV controller for the KV key specified.

The command uses Command Dword 2, Command Dword 3, Command Dword 10, Command Dword 11, Command Dword 14, and Command Dword 15 fields. All other command specific fields are reserved. If the command uses PRPs for the data transfer, then the PRP Entry 1, and PRP Entry 2 fields are used. If the command uses SGLs for the data transfer, then the SGL Entry 1 field is used.

If the value in the Key Length field is greater than 16, then the controller shall abort the command with Invalid Field in Command.

Figure 17: Retrieve – Data Pointer

Bit	Description
127:00	Data Pointer (DPTR): This field specifies where data is transferred to. Refer to the NVMe Base specification for the definition of this field.

Figure 18: Retrieve – Command Dword 10

Bit	Description
31:00	Host Buffer Size (HBS): This field indicates the host buffer size in bytes.

Figure 19: Retrieve – Command Dword 11

Bit	Description
31:16	Reserved
15:8	<p>Retrieve Option (RO): This field specifies the retrieve option. Bits 15:9 are reserved.</p> <p>Bit 8 if set to '1' specifies that the controller shall return raw data (i.e., no decompression is performed on the data). Bit 8 if cleared to '0' specifies that the controller shall return decompressed data if compression is supported. Control of compression algorithms, if any, and their use by the controller is outside the scope of this specification.</p> <p>If the controller does not compresses data then this bit is ignored.</p>
7:0	Key Length (KL): Specifies the length of the KV key in bytes.

Figure 20: Retrieve – Command Dword 2 and Command Dword 3

Bit	Description
63:0	KV key[63:00]: This field specifies 64-bits of the KV Key to be used for the command. Command Dword 2 contains bits 31:00; Command Dword 3 contains bits 63: 32.

Figure 21: Retrieve –Command Dword 14 and Command Dword 15

Bit	Description
63:0	KV key[127:64]: This field specifies 64-bits of the KV Key to be used for the command. Command Dword 14 contains bits 95:64; Command Dword 15 contains bits 127: 96.

5.6.1 Command Completion

Upon completion of the Retrieve command, the controller shall post a completion queue entry to the associated I/O Completion Queue indicating the status for the command. Dword 0 of the completion queue entry contains the KV value size in bytes.

If the host buffer size is less than the value size then only the portion of the value that fits in the host buffer is returned starting at the beginning of the KV value. If the host requires the entire value, then the host should issue a subsequent Retrieve command with a buffer large enough to retrieve the value length returned in the I/O Completion Queue.

If the size of the value is greater than the buffer size, then as much of the value as fits in the buffer allocated shall be returned starting at the beginning of the KV value.

Retrieve command specific status values are defined in Figure 22.

Figure 22: Retrieve – Command Specific Status Values

Value	Description
86h	Invalid Key Size: The KV key size is not valid
0Bh	Invalid Namespace or Format: The namespace or the format of that namespace is invalid.
87h	KV Key Does Not Exist: The KV key does not exist
88h	Unrecovered Error: There was an unrecovered error when reading from the medium

5.7 Exist command

The Exist command returns a status indicating if the specified KV key exists.

The command uses Command Dword 2, Command Dword 3, Command Dword 11, Command Dword 14, and Command Dword 15 fields. All other command specific fields are reserved.

If the value in the Key Length field is greater than 16, then the controller shall abort the command with Invalid Field in Command.

Figure 23: Exist – Command Dword 11

Bit	Description
31:8	Reserved
7:0	Key Length (KL): Specifies the length of the KV key in bytes.

Figure 24: Exist – Command Dword 2 and Command Dword 3

Bit	Description
63:0	KV key[63:00]: This field specifies 64-bits of the KV Key to be used for the command. Command Dword 2 contains bits 31:00; Command Dword 3 contains bits 63: 32.

Figure 25: Exist – Command Dword 14 and Command Dword 15

Bit	Description
63:0	KV key[127:64]: This field specifies 64-bits of the KV Key to be used for the command. Command Dword 14 contains bits 95:64; Command Dword 15 contains bits 127: 96.

5.7.1 Command Completion

Upon completion of the Exist command, the controller posts a completion queue entry (CQE) to the associated I/O Completion Queue. If the status code returned is 00h, then the KV key exists. The Exist command specific status values are defined in Figure 26.

Figure 26: Exist – Command Specific Status Values

Value	Description
87h	KV Key Does Not Exist: The KV key does not exist.

5.8 Store command

The Store command stores a value to the NVM KV controller for the KV key specified.

The command uses Command Dword 2, Command Dword 3, Command Dword 10, Command Dword 11, Command Dword 14, and Command Dword 15 fields. If the command uses PRPs for the data transfer, then the PRP Entry 1, and PRP Entry 2 fields are used. If the command uses SGLs for the data transfer, then the SGL Entry 1 field is used.

Figure 27: Store – Data Pointer

Bit	Description
127:00	Data Pointer (DPTR): This field specifies the location of a data buffer from which data is transferred. Refer to the NVMe Base specification for the definition of this field.

Figure 28: Store – Command Dword 10

Bit	Description
31:00	Value Size (VS): This field indicates the KV value size in bytes. A KV value of zero specifies that there is no value associated with this KV key but that the KV key exists.

Figure 29: Store – Command Dword 11

Bit	Description
31:16	Reserved
15:8	Store Option (SO): Specifies the store option
	Bits 15:11 are reserved
	Bit 10 if set to '1' specifies that the controller shall not compress the KV value. Bit 10 if cleared to '0' specifies that the controller shall compress the KV value if compression is supported.
	Bit 9 if set to '1' specifies that the controller shall not store the KV value if the KV key exists. Bit 9 if cleared to '0' specifies that the controller shall store the KV value if other Store Options are met.
7:0	Bit 8 if set to '1' specifies that the controller shall not store the KV value if the KV key does not exist. Bit 8 if cleared to '0' specifies that the controller shall store the KV value if other Store Options are met.
	Key Length (KL): Specifies the length of the KV key in bytes.

Figure 30: Store – Command Dword 2 and Command Dword 3

Bit	Description
63:0	KV key[63:00]: This field specifies 64-bits of the KV Key to be used for the command. Command Dword 2 contains bits 31:00; Command Dword 3 contains bits 63: 32.

Figure 31: Store –Command Dword 14 and Command Dword 15

Bit	Description
63:0	KV key[127:64]: This field specifies 64-bits of the KV Key to be used for the command. Command Dword 14 contains bits 95:64; Command Dword 15 contains bits 127: 96.

5.8.1 Command Completion

Upon completion of the Store command, the controller shall post a completion queue entry to the associated I/O Completion Queue indicating the status for the command.

Store command specific errors are defined in Figure 32.

Figure 32: Store – Command Specific Status Values

Value	Description
85h	Invalid Value Size: The value size is not valid
86h	Invalid Key Size: The KV key size is not valid
0Bh	Invalid Namespace or Format: The namespace or the format of that namespace is invalid.
81h	Capacity Exceeded: The capacity of the device was exceeded
89h	Key Exists: Store Option bit 9 is set to '1' and the KV key exists
87h	KV Key Does Not Exist: Store Option bit 8 is set to '1' and the KV key does not exist

5.9 Key Value Configuration (Feature Identifier 20h), (Mandatory for Key Value Command Set)

This Feature controls behavior of the Key Value Command Set. The scope of this Feature is the namespace.

The attributes are indicated in Command Dword 11.

If a Get Features command is submitted for this Feature, the attributes specified in **Figure 33** are returned in Dword 0 of the completion queue entry for that command.

If the capabilities of the Key Value Config Feature Identifier are both changeable and saveable (refer to the NVMe Base specification), then the host is able to configure this Feature when initially provisioning a device.

Figure 33: Key Value Config – Command Dword 11

Bits	Description
31:01	Reserved
00	<p>Error on Delete of Non-Existent KV Key (EDNEK): The EDNEK bit defines the response of the controller to a Delete command processed for a KV Key that does not exist.</p> <p>If the EDNEK bit is set to '1' and the controller process a Delete command that specifies a KV key that does not exist, then the controller shall abort the command with a status code of KV Key Does Not Exist.</p> <p>If the EDNEK bit is cleared to '0' and the controller process a Delete command that specifies a KV key that does not exist, then the controller shall not abort the command with a status code of KV Key Does Not Exist. (i.e., complete the command as if the KV key existed and was deleted).</p>

BASE SPEC CHANGES

Add Section 1.6a as shown below:

1.6a I/O Command Set specific definitions used in the NVMe Base specification

The following terms used in this specification are defined in each I/O Command Set specification

1.6a.1 User Data In Command

An I/O Command Set specific command that transfers user data from the Controller to the Host.

For the NVM Command Set these commands are Compare command, Copy command, and Read command.

1.6a.2 User Data Out Command

An I/O Command Set specific command that transfers user data from the Host to the Controller.

For the NVM Command Set these commands are Copy command, and Write command.

Modify figure TBDb I/O Command Set Vector as shown below:

5.15.2. TBDb Identify I/O Command Set data structure (CNS 1Ch)

...

Modify Figure TBDb: Command Set Vector <Refer to the integration of TP 4056b>:

Figure TBDb: I/O Command Set Vector

Bit	Description
63:2	Reserved
1	Key Value Command Set: This bit is set '1' if the Key Value Command Set is selected. This bit is cleared to '0' if the Key Value Command Set is not selected.
0	NVM Command Set: This bit is set to '1' if the NVM Command Set is selected. This bit is cleared to '0' if the NVM Command Set is not selected.

...

Modify figure 202 Get Log Page – SMART / Health Information Log as shown below:

Figure 34: Get Log Page – SMART / Health Information Log

Bytes	Description
...	...
31:07	Reserved

Figure 34: Get Log Page – SMART / Health Information Log

Bytes	Description
47:32	<p>Data Units Read: Contains the number of 512 byte data units the host has read from the controller <u>as part of processing a User Data In Command</u>; this value does not include metadata. This value is reported in thousands (i.e., a value of 1 corresponds to 1,000 units of 512 bytes read) and is rounded up (e.g., one indicates that the number of 512 byte data units read is from 1 to 1,000, three indicates that the number of 512 byte data units read is from 2,001 to 3,000).</p> <p>For the NVM command set, logical blocks read as part of Compare, Read, and Verify commands are also included in this value. operations shall be included in this value. The host does not read any data from the controller for Copy commands.</p> <p>A value of 0h in this field indicates that the number of Data Units Read is not reported.</p>
63:48	<p>Data Units Written: Contains the number of 512 byte data units the host has written to the controller <u>as part of processing a User Data Out Command</u>; this value does not include metadata. This value is reported in thousands (i.e., a value of 1 corresponds to 1,000 units of 512 bytes written) and is rounded up (e.g., one indicates that the number of 512 byte data units written is from 1 to 1,000, three indicates that the number of 512 byte data units written is from 2,001 to 3,000).</p> <p>For the NVM command set, the host does not write any user data to the controller for Copy commands. logical blocks written as part of Write operations shall be included in this value. Write Uncorrectable commands shall not impact this value. A value of 0h in this field indicates that the number of Data Units Written is not reported.</p>
79:64	<p>Host Read Commands: Contains the number of read commands <u>User Data In Commands</u> completed by the controller.</p> <p>For the NVM command set, this value is the sum of the number of Compare commands and the number of Read commands.</p>
95:80	<p>Host Write Commands: Contains the number of <u>User Data Out Commands</u> completed by the controller.</p> <p>For the NVM command set, this is the number of Write commands.</p>
...	...

Modify figure 212 Get Log Page – Endurance Group Log (Log Identifier 09h) as shown below:

Figure 35: Get Log Page – Endurance Group Log (Log Identifier 09h)

Bytes	Description
...	...
111:96	<p>Host Read Commands: Contains the number of read commands <u>User Data In Commands</u> completed by the controller.</p> <p>For the NVM command set, this value is the sum of the number of Compare commands and the number of Read commands.</p>
127:112	<p>Host Write Commands: Contains the number of <u>User Data Out Commands</u> completed by the controller.</p> <p>For the NVM command set, this is the number of Write commands.</p>
...	...

Modify figure X1: Command Set Identifiers <Refer to the integration of TP 4056b>:

Figure X1: Command Set Identifiers

Namespace Type Value	Description	Reference Section
00h	NVM Command Set	5.15.2. TBDa .1.2
01h	Key Value Command Set	Refer to the Key Value Command Set specification
02h to 2Fh	Reserved	
30h to 3Fh	Vendor specific	
40h to FFh	Reserved	

...