



NVM Express®

Key Value Command Set Specification

Revision 1.1

August 5th, 2024

Please send comments to info@nvmexpress.org

NVM Express® Key Value Command Set Specification, Revision 1.1 is available for download at <https://nvmexpress.org>. The NVM Express Key Value Command Set Specification, Revision 1.1 incorporates NVM Express® Key Value Command Set Specification, Revision 1.0, ratified on June 3, 2021, ECN 001, ECN102, ECN106, ECN108, ECN109, ECN110, ECN111, ECN116, ECN118, ECN119, ECN120, TP 4015b, TP4055, TP 4074a, TP4095a, TP 4097a, TP4135, TP4138, TP4152, TP4162a, and TP8012 (refer to <https://nvmexpress.org/changes-in-nvme-revision-2-1> for details).

SPECIFICATION DISCLAIMER

LEGAL NOTICE:

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

This NVM Express Key Value Command Set, Revision 1.1 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, Revision 1.1 subject, however, to the Member’s continued compliance with the Company’s Intellectual Property Policy and Bylaws and the Member’s Participation Agreement.

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

LEGAL DISCLAIMER:

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

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

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

PCIe® is a registered trademark of PCI-SIG.

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

Table of Contents

1	INTRODUCTION.....	6
1.1	Overview.....	6
1.2	Scope.....	6
1.3	Conventions.....	7
1.4	Definitions.....	7
1.5	References.....	8
2	KEY VALUE COMMAND SET MODEL.....	9
2.1	Theory of operation.....	9
2.2	I/O Controller Requirements.....	11
3	I/O COMMANDS FOR THE KEY VALUE COMMAND SET.....	13
3.1	Submission Queue Entry and Completion Queue Entry.....	13
3.2	Key Value Command Set Commands.....	14
4	ADMIN COMMANDS FOR THE KEY VALUE COMMAND SET.....	21
4.1	Admin Command behavior for the Key Value Command Set.....	21
5	EXTENDED CAPABILITIES.....	29
5.1	Namespace Management.....	29
5.2	Reservations.....	29
5.3	Sanitize Operations.....	29

Table of Figures

Figure 1: NVMe Family of Specifications	6
Figure 2: I/O Controller – Key Value Command Set Support	11
Figure 3: I/O Controller – Feature Support	12
Figure 4: Status Code – Generic Command Status Values, Key Value Command Set	13
Figure 5: Opcodes for Key Value Command Set Commands	14
Figure 6: Delete – Command Dword 11	14
Figure 7: Delete – Command Dword 2 and Command Dword 3	14
Figure 8: Delete – Command Dword 14 and Command Dword 15	15
Figure 9: Delete – Generic Command Status Values	15
Figure 10: List – Command Dword 10	15
Figure 11: List – Command Dword 11	15
Figure 12: List – Command Dword 2 and Command Dword 3	15
Figure 13: List – Command Dword 14 and Command Dword 15	16
Figure 14: List – Generic Command Status Values	16
Figure 15: List – Return data structure	16
Figure 16: Key data structure	16
Figure 17: Retrieve – Data Pointer	17
Figure 18: Retrieve – Command Dword 10	17
Figure 19: Retrieve – Command Dword 11	17
Figure 20: Retrieve – Command Dword 13	17
Figure 21: Retrieve – Command Dword 2 and Command Dword 3	17
Figure 22: Retrieve – Command Dword 14 and Command Dword 15	17
Figure 23: Retrieve – Generic Command Status Values	18
Figure 24: Exist – Command Dword 11	18
Figure 25: Exist – Command Dword 2 and Command Dword 3	18
Figure 26: Exist – Command Dword 14 and Command Dword 15	18
Figure 27: Exist – Generic Command Status Values	18
Figure 28: Store – Data Pointer	19
Figure 29: Store – Command Dword 10	19
Figure 30: Store – Command Dword 11	19
Figure 31: Store – Command Dword 13	19
Figure 32: Store – Command Dword 2 and Command Dword 3	19
Figure 33: Store – Command Dword 14 and Command Dword 15	20
Figure 34: Store – Generic Command Status Values	20
Figure 35: Feature Identifiers – Key Value Command Set	21
Figure 36: Key Value Config – Command Dword 11	21
Figure 37: Get Log Page – Log Page Identifiers	22
Figure 38: Error Information Log Entry Data Structure	22
Figure 39: Self-test Results Data Structure	22
Figure 40: Identify – CNS Values	23
Figure 41: Identify – I/O Command Set Specific Identify Namespace Data Structure, Key Value Type Specific	23
Figure 42: KV Format Data Structure	25
Figure 43: I/O Command Set Specific Identify Controller Data Structure for the Key Value Command Set	26
Figure 44: Key Value Command Set Specification Version Descriptor Field Values	26
Figure 45: Command Dword 11 - CNS Specific Identifier	27
Figure 46: Namespace Management – Host Software Specified Fields	28
Figure 47: Command Behavior in the Presence of a Reservation	29

1 Introduction

1.1 Overview

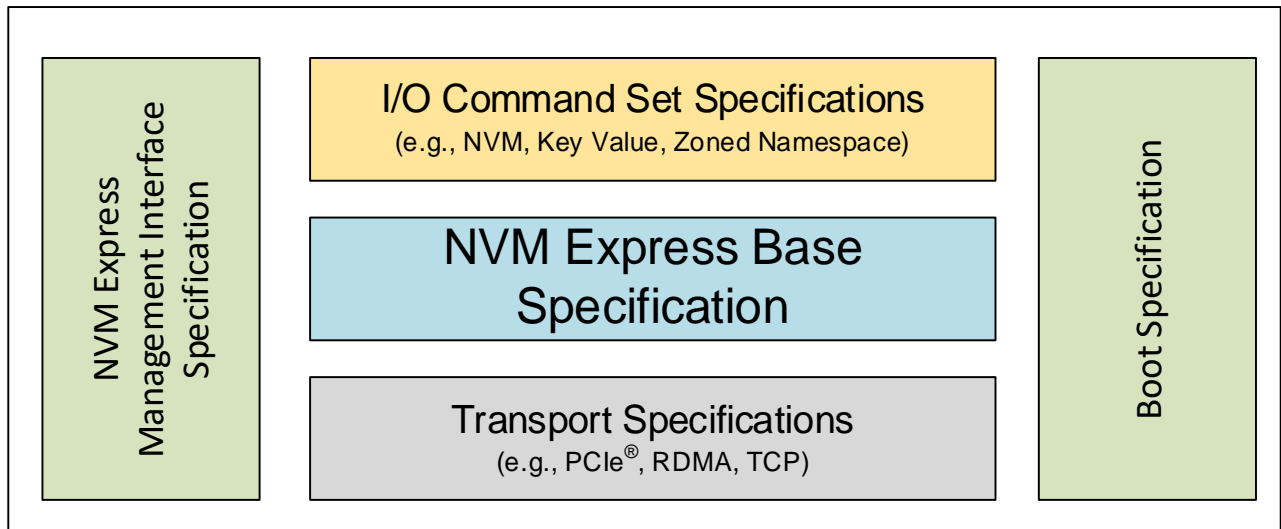
NVM Express® (NVMe®) Base Specification defines an interface for host software to communicate with a non-volatile memory subsystem (NVM subsystem) over a variety of memory-based transports and message-based transports.

This document defines a specific NVMe I/O Command Set, the Key Value Command Set, which extends the NVM Express Base Specification.

1.2 Scope

Figure 1 shows the relationship of the NVM Express® Key Value Command Set Specification to other specifications within the NVMe Family of Specifications.

Figure 1: NVMe Family of Specifications



This specification supplements the NVMe Express Base Specification. 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 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 Express Base Specification.

If a conflict arises among requirements defined in different specifications, then a lower-numbered specification in the following list shall take precedence over a higher-numbered specification:

1. Non-NVMe specifications
2. NVMe Express Base Specification
3. NVMe transport specifications
4. NVMe I/O command set specifications
5. NVMe Express Management Interface Specification
6. NVMe Express® Boot Specification

1.3 Conventions

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

1.4 Definitions

1.4.1 Definitions from the NVM Express Base Specifications

This specification uses the definitions in the NVM Express Base Specification.

1.4.2 Definitions in the NVM Express Base Specification specified in the Key Value Command Set

The following terms used in this specification and the NVM Express Base Specification are as defined here.

1.4.2.1 Endurance Group Host Read Command

A Retrieve command.

1.4.2.2 Format Index

A value used to index into the KV Formats list (refer to Figure 41).

1.4.2.3 Identify Controller data structures

All controller data structures that are able to be retrieved via the Identify command for the Key Value Command Set:

- the Identify Controller data structure (refer to the NVM Express Base Specification); and
- the I/O Command Set specific Identify Controller data structure for the Key Value Command Set (refer to section 4.1.5.2).

1.4.2.4 Identify Namespace data structures

All namespace data structures that are able to be retrieved via the Identify command for the Key Value Command Set:

- the I/O Command Set Independent Identify Namespace data structure (refer to the NVM Express Base Specification); and
- the I/O Command Set specific Identify Namespace data structure for the Key Value Command Set (refer to section 4.1.5.1).

1.4.2.5 SMART Data Units Read Command

A Retrieve command.

1.4.2.6 SMART Host Read Command

A Retrieve command.

1.4.2.7 User Data Format

The layout of the data on the NVM media as described by the Key Value Format of the namespace.

1.4.2.8 User Data Out Command

A Store command.

1.4.3 Definitions specific to the Key Value Command Set

This section defines terms that are specific to this specification.

1.4.3.1 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.

1.4.3.2 KV key

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

1.4.3.3 KV value

The value that is associated with a key value pair.

1.5 References

NVM Express® Base Specification, Revision 2.1. Available from <https://www.nvmexpress.org>.

2 Key Value Command Set Model

The NVM Express Base Specification defines a register level interface for host software to communicate with a non-volatile memory subsystem (NVM 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 NVM Express Base Specification. The Key Value Command Set is assigned a CSI value of 01h.

2.1 Theory of operation

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.

2.1.1 Namespaces

A namespace is a set of resources that may be accessed by a host and is as defined in the NVM Express Base Specification. A namespace has an associated namespace identifier that a host uses to access that namespace.

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.

The Key Value Command Set specific Identify Namespace data structure (refer to section 4.1.5.1) contains related fields reporting the namespace size, capacity and utilization:

- The Namespace Size (NSZE) field defines the total size of the namespace in bytes.
- The Namespace Utilization (NUSE) field defines the number of bytes of namespace capacity that are currently in use to store KV keys and KV values.

The Namespace Utilization (NUSE) field shall be less than or equal to the Namespace Size (NSZE) field.

2.1.2 Command Ordering Requirements

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, the host enforces those requirements above the level of the controller.

2.1.3 Fused Operation

The Key Value Command Set does not support any Fused Operations.

2.1.4 Atomic Operation

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

2.1.5 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 a status code of Invalid Field in Command.

2.1.6 Command Operation

2.1.6.1 Delete command

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

2.1.6.2 List command

The host may request a list of the KV keys in the 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 Sanitize, Format NVM, Store, and Delete commands the order of the KV keys in the list shall be constant. The KV key that is specified in a List 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 data structure. 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 Sanitize, Format NVM, Store, and Delete commands the first KV key returned shall not change.

2.1.6.3 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 completion queue entry (CQE) for that command.

2.1.6.4 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 (e.g., 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).

2.1.6.5 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 portion of the KV value and the full 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.

2.2 I/O Controller Requirements

2.2.1 Command Support

This specification implements the command support requirements for I/O Controllers defined in the NVM Express Base Specification. Figure 2 defines Key Value Command Set specific definitions for I/O commands that are mandatory, optional, and prohibited for an I/O controller that supports the Key Value Command Set.

Figure 2: I/O Controller – Key Value Command Set Support

Command	Command Support Requirements ¹
Store	M
Retrieve	M
Delete	M
Exist	M
List	M
Notes: 1. O = Optional, M = Mandatory, P = Prohibited	

2.2.2 Log Page Support

This specification implements the log page support requirements for I/O Controllers defined in the NVM Express Base Specification. There are no additional Key Value Command Set specific definitions for log pages that are mandatory, optional, and prohibited for an I/O controller that supports the Key Value Command Set.

2.2.3 Features Support

This specification implements the feature support requirements for I/O Controllers defined in the NVM Express Base Specification. Figure 3 defines Key Value Command Set specific definitions for features that are mandatory, optional, prohibited, and not recommended for an I/O Controller that supports the Key Value Command Set.

Figure 3: I/O Controller – Feature Support

Feature Name	Feature Support Requirements ¹	Logged in Persistent Event Log
Key Value Configuration	M	Yes
Notes: 1. O = Optional, M = Mandatory, P = Prohibited, NR = Not Recommended		

3 I/O Commands for the Key Value Command Set

This section specifies the Key Value Command Set I/O commands.

3.1 Submission Queue Entry and Completion Queue Entry

The Submission Queue Entry (SQE) and the fields that are common to all NVMe I/O Command Sets are defined in the Submission Queue Entry section in the NVM Express Base Specification. The completion queue entry (CQE) and the fields that are common to all NVMe I/O Command Sets are defined in the Completion Queue Entry section in the NVM Express Base Specification.

The Key Value Command Set uses the Common Command Format as defined in the NVM Express Base Specification.

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

The command specific fields in the SQE (i.e., SQE Command Dword 2, Command Dword 3, Command Dwords 10-15) and the CQE (i.e., CQE Dword 0, and Dword 1) for the Key Value Command Set are defined in this section.

3.1.1 Common Command Format

The Common Command Format is as defined in the NVM Express Base Specification.

SQE Command Dword 2 and Command Dword 3 contain KV key bytes [7:0]. SQE Command Dword 14 and Command Dword 15 contain KV key bytes [15:8].

3.1.2 Key Value Command Set Specific Status Values

No command specific status values are defined in this specification.

This specification supports the Generic Command status values defined in the NVM Express Base Specification. Generic Command status values that are reported by commands defined in this specification are described in Figure 4.

Figure 4: Status Code – Generic Command Status Values, Key Value Command Set

Value	Definition	Commands Affected
81h	Capacity Exceeded	Store
82h	Namespace Not Ready	Delete, Exist, Retrieve, Store
83h	Reservation Conflict	Delete, Store, Retrieve
84h	Format In Progress	Delete, Exist, List, Retrieve, Store
85h	Invalid Value Size	Store
86h	Invalid Key Size	List, Retrieve, Store
87h	KV Key Does Not Exist	Delete, Exist, Retrieve, Store
88h	Unrecovered Error	Retrieve
89h	Key Exists	Store

3.2 Key Value Command Set Commands

The Key Value Command Set includes the commands listed in Figure 5. Section 3.2 describes the definition for each of the commands defined by this specification. Commands are submitted as described in the NVM Express Base Specification.

Figure 5: Opcodes for Key Value Command Set Commands

Opcode by Field		Combined Opcode ¹	Command ²	Reference
(07:02)	(01:00)			
Function	Data Transfer ³			
Refer to the NVM Express Base Specification			Flush	NVM Express Base Specification
Refer to the NVM Express Base Specification			Reservation Register	NVM Express Base Specification
Refer to the NVM Express Base Specification			Reservation Report	NVM Express Base Specification
Refer to the NVM Express Base Specification			Reservation Acquire	NVM Express Base Specification
Refer to the NVM Express Base Specification			Reservation Release	NVM Express Base Specification
Refer to the NVM Express Base Specification			Cancel ⁴	NVM Express Base Specification
0000 00b	01b	01h	Store	3.2.5
0000 00b	10b	02h	Retrieve	3.2.3
0001 00b	00b	10h	Delete	3.2.1
0001 01b	00b	14h	Exist	3.2.4
0000 01b	10b	06h	List	3.2.2
Notes:				
1. Opcodes not listed are defined in the NVM Express 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.				

3.2.1 Delete command

The Delete command deletes the KV key and the associated KV value for the specified 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 6: Delete – Command Dword 11

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

Figure 7: Delete – Command Dword 2 and Command Dword 3

Bits	Description
63:0	KV Key[63:00]: This field specifies the least-significant 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 8: Delete – Command Dword 14 and Command Dword 15

Bits	Description
63:0	KV Key[127:64]: This field specifies the most-significant 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.

3.2.1.1 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 generic status values are defined in Figure 9.

Figure 9: Delete – Generic Command 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 Key Value Command Set.

3.2.2 List command

The List command retrieves a list of KV keys that exist for the specified namespace starting at the KV key specified. The number of keys returned are the minimum of:

- a) the number of keys in the controller; or
- b) the number of complete keys that fit in the buffer provided by the host.

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 10: List – Command Dword 10

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

Figure 11: List – Command Dword 11

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

Figure 12: List – Command Dword 2 and Command Dword 3

Bits	Description
63:0	KV Key[63:00]: This field specifies least-significant 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 13: List – Command Dword 14 and Command Dword 15

Bits	Description
63:0	KV Key[127:64]: This field specifies the most-significant 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.

3.2.2.1 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 3.2.2.2.

List command generic status values are defined in Figure 14.

Figure 14: List – Generic Command Status Values

Value	Definition
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.

3.2.2.2 List command return data structure

The data structure returned for the list command is as defined in Figure 15.

Figure 15: List – Return data structure

Bytes	Description
03:00	Number of Returned Keys (NRK): This value reflects how many KV keys are returned in this data structure.
Key Data Structure List	
Variable:04	Key Data Structure 1: The first KV Key (refer to Figure 16), if any.
Variable	Key Data Structure 2: The second KV Key (refer to Figure 16), if any.
...	...
Variable	Key Data Structure NRK: The last KV Key (refer to Figure 16), if any.

Figure 16: Key data structure

Bytes	Description
01:00	Key Length (KL): indicates the length of the KV key in bytes that this data structure represents.
n:02	KV Key (KVK): KV key that this entry describes.
m:n+1	Pad: Pad necessary, if any to end the data structure on a 4 byte boundary.

3.2.3 Retrieve command

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

The command uses Command Dword 2, Command Dword 3, Command Dword 10, Command Dword 11, Command Dword 13, 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

Bits	Description
127:00	Data Pointer (DPTR): This field specifies where data is transferred to. Refer to the NVM Express Base Specification for the definition of this field.

Figure 18: Retrieve – Command Dword 10

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

Figure 19: Retrieve – Command Dword 11

Bits	Description						
31:16	Reserved						
15:8	Retrieve Option (RO): This field specifies the retrieve option.						
	<table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>15:9</td> <td>Reserved</td> </tr> <tr> <td>8</td> <td>Return Raw Data (RRD): If this bit is set to '1', then the controller shall return raw data (i.e., no decompression is performed on the data). If this bit cleared to '0', then 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. If the controller does not compress data then this bit is ignored.</td> </tr> </tbody> </table>	Bits	Description	15:9	Reserved	8	Return Raw Data (RRD): If this bit is set to '1', then the controller shall return raw data (i.e., no decompression is performed on the data). If this bit cleared to '0', then 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. If the controller does not compress data then this bit is ignored.
	Bits	Description					
15:9	Reserved						
8	Return Raw Data (RRD): If this bit is set to '1', then the controller shall return raw data (i.e., no decompression is performed on the data). If this bit cleared to '0', then 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. If the controller does not compress data then this bit is ignored.						
7:0	Key Length (KL): Specifies the length of the KV key in bytes.						

Figure 20: Retrieve – Command Dword 13

Bits	Description
31:20	Reserved
19:16	Command Extension Type (CETYPE): Specifies the Command Extension Type that applies to the command (refer to the Key Per I/O section in the NVM Express Base Specification).
15:00	Command Extension Value (CEV): The definition of this field is dependent on the value of the CETYPE field. Refer to the Key Per I/O section in the NVM Express Base Specification.

Figure 21: Retrieve – Command Dword 2 and Command Dword 3

Bits	Description
63:0	KV Key[63:00]: This field specifies the least-significant 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 22: Retrieve –Command Dword 14 and Command Dword 15

Bits	Description
63:0	KV Key[127:64]: This field specifies the most-significant 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.

3.2.3.1 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. On successful completion of 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 size of the KV value then the portion of the KV value that fits in the host buffer shall be returned starting at the beginning of the KV value. If the host requires the entire KV

value, then the host should issue a subsequent Retrieve command with a buffer large enough to retrieve the KV value length returned in the I/O Completion Queue.

Retrieve command generic status values are defined in Figure 23.

Figure 23: Retrieve – Generic Command Status Values

Value	Definition
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.

3.2.4 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 24: Exist – Command Dword 11

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

Figure 25: Exist – Command Dword 2 and Command Dword 3

Bits	Description
63:0	KV Key[63:00]: This field specifies the least-significant 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 26: Exist – Command Dword 14 and Command Dword 15

Bits	Description
63:0	KV Key[127:64]: This field specifies most-significant 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.

3.2.4.1 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 generic status values are defined in Figure 27.

Figure 27: Exist – Generic Command Status Values

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

3.2.5 Store command

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

The command uses Command Dword 2, Command Dword 3, Command Dword 10, Command Dword 11, Command Dword 13, 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 28: Store – Data Pointer

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

Figure 29: Store – Command Dword 10

Bits	Description
31:00	Value Size (VS): This field indicates the KV value size in bytes. A KV value of 0h specifies that there is no value associated with this KV key but that the KV key exists.

Figure 30: Store – Command Dword 11

Bits	Description										
31:16	Reserved										
15:8	Store Option (SO): Specifies the store option.										
	<table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>15:11</td> <td>Reserved</td> </tr> <tr> <td>10</td> <td>No Compression (NOCOMP): If this bit is set to '1', then the controller shall not compress the KV value. If this bit is cleared to '0', then the controller shall compress the KV value if compression is supported.</td> </tr> <tr> <td>9</td> <td>Store If No Key Exists (SINKE): If this bit is set to '1', then the controller shall not store the KV value if the KV key exists. If this bit is cleared to '0', then the controller shall store the KV value if other Store Options are met.</td> </tr> <tr> <td>8</td> <td>Store If Key Exist (SIKE): If this bit is set to '1', then the controller shall not store the KV value if the KV key does not exist. If this bit is cleared to '0', then the controller shall store the KV value if other Store Options are met.</td> </tr> </tbody> </table>	Bits	Description	15:11	Reserved	10	No Compression (NOCOMP): If this bit is set to '1', then the controller shall not compress the KV value. If this bit is cleared to '0', then the controller shall compress the KV value if compression is supported.	9	Store If No Key Exists (SINKE): If this bit is set to '1', then the controller shall not store the KV value if the KV key exists. If this bit is cleared to '0', then the controller shall store the KV value if other Store Options are met.	8	Store If Key Exist (SIKE): If this bit is set to '1', then the controller shall not store the KV value if the KV key does not exist. If this bit is cleared to '0', then the controller shall store the KV value if other Store Options are met.
	Bits	Description									
	15:11	Reserved									
10	No Compression (NOCOMP): If this bit is set to '1', then the controller shall not compress the KV value. If this bit is cleared to '0', then the controller shall compress the KV value if compression is supported.										
9	Store If No Key Exists (SINKE): If this bit is set to '1', then the controller shall not store the KV value if the KV key exists. If this bit is cleared to '0', then the controller shall store the KV value if other Store Options are met.										
8	Store If Key Exist (SIKE): If this bit is set to '1', then the controller shall not store the KV value if the KV key does not exist. If this bit is cleared to '0', then the controller shall store the KV value if other Store Options are met.										
7:0	Key Length (KL): Specifies the length of the KV key in bytes.										

Figure 31: Store – Command Dword 13

Bits	Description
31:20	Reserved
19:16	Command Extension Type (CETYPE): Specifies the Command Extension Type that applies to the command (refer to the Key Per I/O section in the NVM Express Base Specification).
15:00	Command Extension Value (CEV): The definition of this field is dependent on the value of the CETYPE field. Refer to the Key Per I/O section in the NVM Express Base Specification.

Figure 32: Store – Command Dword 2 and Command Dword 3

Bits	Description
63:0	KV Key[63:00]: This field specifies the least-significant 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 33: Store –Command Dword 14 and Command Dword 15

Bits	Description
63:0	KV Key[127:64]: This field specifies the most-significant 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.

3.2.5.1 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 generic errors are defined in Figure 34.

Figure 34: Store – Generic Command Status Values

Value	Definition
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: The Store If No Key Exists (SINKE) bit is set to '1' in the Store Option field and the KV key exists.
87h	KV Key Does Not Exist: The Store If Key Exists (SIKE) bit is set to '1' in the Store Option field-and the KV key does not exist.

4 Admin Commands for the Key Value Command Set

4.1 Admin Command behavior for the Key Value Command Set

The Admin commands are as defined in the NVM Express Base Specification. The Key Value Command Set specific behavior for Admin commands is described in this section.

4.1.1 Asynchronous Event Request command

The Asynchronous Event Request command operates as defined in the NVM Express Base Specification. The Key Value Command Set does not define any additional Asynchronous Events.

4.1.2 Format NVM command – Key Value Command Set Specific

The Format NVM command operates as defined in the NVM Express Base Specification. The Format Index indicates a valid KV Format from the KV Format field in the Key Value Command Set specific Identify Namespace data structure.

4.1.3 Get Features & Set Features commands

Figure 35 defines the Features support requirements for I/O Controllers supporting the Key Value Command Set.

Figure 35: Feature Identifiers – Key Value Command Set

Feature Identifier	Persistent Across Power Cycle and Reset ¹	Uses Memory Buffer for Attributes	Description	Scope
20h	Yes	No	Key Value Configuration	Namespace
Notes:				
1. This column is only valid if the feature is not saveable (refer to the NVM Express 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.				

4.1.3.1 Key Value Configuration (Feature Identifier 20h)

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 36 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 NVM Express Base Specification), then the host is able to configure this Feature when initially provisioning a device.

Figure 36: Key Value Config – Command Dword 11

Bits	Description
31:01	Reserved

Figure 36: Key Value Config – Command Dword 11

Bits	Description
00	<p>Error on Delete of Non-Existent KV Key (EDNEK): This bit defines the response of the controller to a Delete command processed for a KV key that does not exist.</p> <p>If this bit is set to '1' and the controller processes a Delete command that specifies a KV key that does not exist, then the controller shall abort that command with a status code of KV Key Does Not Exist.</p> <p>If this bit is cleared to '0' and the controller processes a Delete command that specifies a KV key that does not exist, then the controller shall not abort that 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>

4.1.4 Get Log Page command

The Get Log Page command operates as defined in the NVM Express Base Specification. In addition to the requirements in the NVM Express Base Specification, mandatory, optional, and prohibited Log Page Identifiers are defined in Figure 37. If a Get Log Page command is processed that specifies a Log Page Identifier that is not supported, then the controller should abort the command with a status code of Invalid Field in Command.

Log page scope is as defined in the NVM Express Base Specification.

The rules for namespace identifier usage are specified in the NVM Express Base Specification.

Figure 37: Get Log Page – Log Page Identifiers

Log Page Identifier	Scope and Support	Log Page Name	Reference
01h	Refer to the NVM Express Base Specification	Error Information	4.1.4.1
06h	Refer to the NVM Express Base Specification	Device Self-test	4.1.4.2

4.1.4.1 Error Information (Log Page Identifier 01h)

The Error Information log page is as defined in the NVM Express Base Specification. Figure 38 describes the Key Value Command Set specific definition of the LBA field.

Figure 38: Error Information Log Entry Data Structure

Bytes	Description
23:16	Logical Block Address (LBA): This field is reserved.

4.1.4.2 Device Self-test (Log Page Identifier 06h)

The Device Self-test log page is as defined in the NVM Express Base Specification. Figure 39 describes the Key Value Command Set specific definition of the Failing LBA field.

Figure 39: Self-test Results Data Structure

Bytes	Description
23:16	Failing LBA (FLBA): This field is reserved.

4.1.5 Identify Command

This specification implements the Identify command and associated Identify data structures defined in the NVM Express Base Specification. Additionally, the Key Value Command Set specifies the data structures defined in this section.

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

Figure 40: Identify – CNS Values

CNS Value	O/M ¹	Definition	NSID ²	CNTID ³	CSI ⁴	Reference Section
Active Namespace Management						
05h	M ⁵	Identify I/O Command Set specific Namespace data structure for the specified NSID for the I/O Command Set specified in the CSI field.	Y	N	Y	4.1.5.1
06h	M	Identify I/O Command Set Specific Controller data structure for the controller processing the command.	Y	N	Y	4.1.5.2
0Ah	O	I/O Command Set specific Identify Namespace data structure for the specified Format Index for the I/O Command Set specified in the CSI field. ⁶	N	N	Y	4.1.5.3

Notes:

- O/M definition: O = Optional, M = Mandatory.
- The NSID field is used: Y = Yes, N = No.
- The CDW10.CNTID field is used: Y = Yes, N = No.
- The CDW11.CSI field is used: Y = Yes, N = No.
- Mandatory for controllers that support the Namespace Management capability (refer to the NVM Express Base Specification).
- Selection of a UUID may be supported. Refer to the Universally Unique Identifiers (UUIDs) for Vendor Specific Information section in the NVM Express Base Specification.

4.1.5.1 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 41.

Figure 41: Identify – I/O Command Set Specific Identify Namespace Data Structure, Key Value Type Specific

Bytes	O/M ¹	Description	Reported ²
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.	No
15:08		Reserved	

Figure 41: Identify – I/O Command Set Specific Identify Namespace Data Structure, Key Value Type Specific

Bytes	O/M 1	Description	Reported ²								
23:16	M	<p>Namespace Utilization (NUSE): This field indicates the current number of bytes of namespace capacity that are in use to store KV keys and KV values. This field is less than or equal to the Namespace Size field.</p> <p>A key value pair begins to use namespace capacity when the key value pair is written with a Store command. A key value pair ceases to use namespace capacity when the key value pair is deleted 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 NVM Express Base Specification) or the ANA Persistent Loss state (refer to the NVM Express Base Specification), then this field shall be cleared to 0h.</p>	No								
24	M	<p>Namespace Features (NSFEAT): This field defines features of the namespace.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:4</td> <td>Reserved</td> </tr> <tr> <td>3</td> <td> <p>Unique ID Reuse (UIDR): If this bit is set to '1', then the non-zero NGUID and non-zero EUI64 fields for this namespace are never reused by the controller.</p> <p>If this bit is 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.</p> <p>This bit shall be cleared to '0' if both NGUID and EUI64 fields are cleared to 0h. Refer to the NVM Express Base Specification.</p> </td> </tr> <tr> <td>2:0</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	7:4	Reserved	3	<p>Unique ID Reuse (UIDR): If this bit is set to '1', then the non-zero NGUID and non-zero EUI64 fields for this namespace are never reused by the controller.</p> <p>If this bit is 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.</p> <p>This bit shall be cleared to '0' if both NGUID and EUI64 fields are cleared to 0h. Refer to the NVM Express Base Specification.</p>	2:0	Reserved	No
Bits	Description										
7:4	Reserved										
3	<p>Unique ID Reuse (UIDR): If this bit is set to '1', then the non-zero NGUID and non-zero EUI64 fields for this namespace are never reused by the controller.</p> <p>If this bit is 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.</p> <p>This bit shall be cleared to '0' if both NGUID and EUI64 fields are cleared to 0h. Refer to the NVM Express Base Specification.</p>										
2:0	Reserved										
25	M	<p>Number of KV Formats (NKVF): This field defines the number of KV format descriptors supported by the namespace. KV formats shall be packed sequentially starting at the KV Format 0 Support (KVF0) field. This is a 0's based value.</p> <p>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.</p> <p>The KV Format fields with a Format 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 clearing the KV Key Max Length field to 0h and clearing the KV Value Max Length field to 0h for that KV Format.</p>	Yes								
26	O	<p>Namespace Multi-path I/O and Namespace Sharing Capabilities (NMIC): Refer to the NMIC field in the I/O Command Set Independent Identify Namespace data structure in the NVM Express Base Specification.</p>	No								
27	O	<p>Reservation Capabilities (RESCAP): Refer to the NVM Express Base Specification.</p>	No								
28	O	<p>Format Progress Indicator (FPI): Refer to the NVM Express Base Specification.</p>	No								

Figure 41: Identify – I/O Command Set Specific Identify Namespace Data Structure, Key Value Type Specific

Bytes	O/M 1	Description	Reported ²						
29	M	<p>KV Format Capabilities (KVFC): This field defines capabilities associated with the KV format.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:4</td> <td>Reserved</td> </tr> <tr> <td>3:0</td> <td>KV Format Index (KVI): This field indicates the Format Index that was used to format the namespace.</td> </tr> </tbody> </table> <p>This field is contained in the Formatted LBA Size (FLBAS) field in the Changed Namespace Event data structure of the Changed Namespace Event in the Persistent Event Log (refer to the NVM Express Base Specification).</p>	Bits	Description	7:4	Reserved	3:0	KV Format Index (KVI): This field indicates the Format Index that was used to format the namespace.	No
		Bits	Description						
7:4	Reserved								
3:0	KV Format Index (KVI): This field indicates the Format Index that was used to format the namespace.								
31:30	Reserved	No							
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.	No						
39:36	O	ANA Group Identifier (ANAGRPID): Refer to the NVM Express Base Specification.	No						
42:40		Reserved	No						
43	O	Namespace Attributes (NSATTR): Refer to the NVM Express Base Specification.	No						
45:44	O	NVM Set Identifier (NVMSETID): Refer to the NVM Express Base Specification.	No						
47:46	O	Endurance Group Identifier (ENDGID): Refer to the NVM Express Base Specification.	No						
63:48	O	Namespace Globally Unique Identifier (NGUID): Refer to the NVM Express Base Specification.	No						
71:64	O	IEEE Extended Unique Identifier (EUI64): Refer to the NVM Express Base Specification.	No						
KV Formats									
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 42.	Yes						
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 42.	Yes						
...									
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 42.	Yes						
3839:328		Reserved							
4095:3840	O	Vendor Specific (VS)	No						
Notes:									
1. O/M definition: O = Optional, M = Mandatory.									
2. Identifies fields that report information for the Identify command when querying the capabilities of LBA formats.									

Figure 42: KV Format Data Structure

Bytes	Description
01:00	KV Key Max Length (KVKML): Maximum length of a KV key in a key value pair in bytes. The value of this field shall be less than or equal to 16.
02	Reserved

Figure 42: KV Format Data Structure

Bytes	Description																	
03	Additional Format Options (AFO):																	
	<table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:2</td> <td>Reserved</td> </tr> <tr> <td rowspan="4">1:0</td> <td>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 4KiB KV value reads. The meanings of the values indicated are included in the following table.</td> </tr> <tr> <td> <table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <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> </tbody> </table> </td> </tr> </tbody> </table>	Bits	Description	7:2	Reserved	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 4KiB KV value reads. The meanings of the values indicated are included in the following table.	<table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <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> </tbody> </table>	Value	Definition	00b	Best performance	01b	Better performance	10b	Good performance	11b	Degraded performance
	Bits	Description																
	7:2	Reserved																
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 4KiB KV value reads. The meanings of the values indicated are included in the following table.																	
	<table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <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> </tbody> </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 (KVVML): Maximum length in bytes of a KV value in a key value pair.																	
11:08	Max Num Keys (MNKS): Maximum number of KV keys allowed in the namespace. A value of 0h indicates that no maximum number is indicated.																	
15:12	Reserved																	

4.1.5.2 I/O Command Set Specific Identify Controller Data Structure (CNS 06h, CSI 01h)

Figure 43 defines the I/O Command Set Specific Identify Controller data structure for the Key Value Command Set.

Figure 43: I/O Command Set Specific Identify Controller Data Structure for the Key Value Command Set

Bytes	O/M ¹	Description
03:00	M	Version (VER): This field contains a Specification Version Descriptor (refer to the NVM Express Base Specification) indicating the version of this specification supported by the controller, as defined in Figure 44.
4095:04		Reserved
Notes:		
1. O/M definition: O = Optional, M = Mandatory.		

Published versions of this specification and the values that shall be reported by compliant controllers are defined in Figure 44.

Figure 44: Key Value Command Set Specification Version Descriptor Field Values

Specification Versions ¹	MJR Field	MNR Field	TER Field
1.0	1h	0h	0h
1.1	1h	1h	0h
Notes:			
1. The specification version listed includes lettered versions (e.g., 1.0 includes 1.0, 1.0a, 1.0b, etc.).			

4.1.5.3 I/O Command Set specific Identify Namespace Data Structure (CNS 0Ah, CSI 01h)

An I/O Command Set specific Identify Namespace data structure for the Key Value Command Set (refer to Figure 41) is returned to the host for the Format Index specified by the CNS Specific Identifier field as defined in Figure 45. The returned I/O Command Set specific Identify Namespace data structure for the Key Value Command Set specifies fields that define capabilities used by a host to format or create a namespace. If the Format Index specified is valid, then the controller shall return an I/O Command Set specific Identify Namespace data structure for the Key Value Command Set that contains:

- fields in Figure 41 where the Reported column indicates “Yes” are set to a value that is the same for all namespaces using the specified Format Index; and
- fields in Figure 41 where the Reported column indicates “No” are cleared to 0h.

Figure 45: Command Dword 11 - CNS Specific Identifier

Bits	Description
15:0	Format Index Identifier (FIDXI): This field specifies the Format Index (refer to Figure 41) identifying the KV Format capabilities that are to be returned.

4.1.5.4 Command Set Index Usage for the Key Value Command Set

The following sections provide an example on how a host uses the CSI value of 01h for accessing Identify Namespace data structures for a namespace associated with the Key Value Command Set.

4.1.5.4.1 Determining the Identify Command Information Associated with a Namespace

For a host to determine the Identify Namespace data structures (refer to section 1.4.2.4) for a namespace associated with the Key Value Command Set, the host is required to issue the following Identify commands in any order:

- a) An Identify command with:
 - a. the CNS field set to 08h; and
 - b. the NSID field set to the NSID of the namespace,

to access the I/O Command Set Independent Identify Namespace data structure (refer to the NVM Express Base Specification);
- b) An Identify command with:
 - a. the CNS field set to 05h;
 - b. the CSI field set to 01h; and
 - c. the NSID field set to the NSID of the namespace,

to access the I/O Command Set specific Identify Namespace data structure for the Key Value Command Set (refer to section 4.1.5.1).

4.1.5.4.2 Determining the Identify Command Information Associated with a Format Index

For a host to determine the Identify Namespace Data Structures associated with a specific Format Index (i.e., determining information about a namespace associated with the Key Value Command Set prior to creating that namespace), the host is required to issue the following Identify commands in any order:

- a) An Identify command with:
 - a. the CNS field set to 08h; and

- b. the NSID field set to FFFFFFFFh,
to access the I/O Command Set Independent Identify Namespace data structure (refer to the NVM Express Base Specification);
- b) An Identify command with:
 - a. the CNS set to 0Ah;
 - b. the CSI set to 01h;
 - c. the NSID field set to 0h; and
 - d. the CNS Specific Identifier field set to the Format Index,
to access the I/O Command Set specific Identify Namespace data structure for the Key Value Command Set (refer to 4.1.5.3).

4.1.6 Namespace Management command

The Namespace Management command operates as defined in the NVM Express Base Specification.

The host specified namespace management fields are specific to the I/O Command Set. The data structure passed to the create operation for the Key Value Command Set (CSI 01h) is defined in Figure 46. Fields that are reserved should be cleared to 0h by host software. After successful completion of a Namespace Management command with the create operation, the namespace is formatted with the specified attributes.

Figure 46: Namespace Management – Host Software Specified Fields

Bytes	Description	Host Specified
These fields are the same fields as defined in the I/O Command Set specific Identify Namespace data structure (refer to Figure 41).		
07:00	Namespace Size (NSZE)	Yes
29:08	Reserved	
30	Namespace Multi-path I/O and Namespace Sharing Capabilities (NMIC)	Yes
31	KV Format Capabilities (KVFC)	Yes
91:32	Reserved	
95:92	ANA Group Identifier (ANAGRPID) ¹	Yes
99:96	Reserved	
101:100	NVM Set Identifier (NVMSETID) ¹	Yes
103:102	Endurance Group Identifier (ENDGID)	Yes
511:104	Reserved	
Notes:		
1. A value of 0h specifies that the controller determines the value to use (refer to the Namespace Management section in the NVM Express Base Specification). If the associated feature is not supported, then this field is ignored by the controller.		

4.1.7 Sanitize command

The Sanitize command operates as defined in the NVM Express Base Specification. There are no Key Value Command Set specific requirements on the Sanitize command.

5 Extended Capabilities

5.1 Namespace Management

Namespace Management operates as defined in the NVM Express Base Specification.

5.2 Reservations

Reservations operate as defined the NVM Express Base Specification with the additional Command Behavior in the Presence of a Reservation defined in Figure 47.

Figure 47: 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 Command Set Read Command Group								
Retrieve	A	A	C	C	A	A	C	A
Key Value Command Set Write Command Group								
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								

5.3 Sanitize Operations

A sanitize operation is performed as defined in the NVM Express Base Specification, with the following exceptions:

- Namespaces associated with the Key Value command set are not accessible in the Media Verification state or the Post-Verification Deallocation state (refer to the Sanitize Operations section in the NVM Express Base Specification).
- Namespaces associated with the Key Value command set do not support additional media modification (refer to the NVM Express Base Specification).