Command Set Specification

Software-Enabled Flash™

Version: 1.14

SEF-CMD-01-14

©2023 Software-Enabled Flash Project. All Rights Reserved.

LEGAL DISCLAIMER

THIS DOCUMENT AND THE INFORMATION CONTAINED HEREIN IS PROVIDED ON AN "AS IS" BASIS. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, THE SOFTWARE-ENABLED FLASH PROJECT, THE LINUX FOUNDATION, AND 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.

DEFINITIONS AND CLARIFICATIONS

Definition of capacity: we define a megabyte (MB) as 1,000,000 bytes, a gigabyte (GB) as 1,000,000,000 bytes and a terabyte (TB) as 1,000,000,000,000 bytes. A computer operating system, however, reports storage capacity using powers of 2 for the definition of 1GB $= 2^{30} =$ 1,073,741,824 bytes and therefore shows less storage capacity. Available storage capacity (including examples of various media files) will vary based on file size, formatting, settings, software and operating system, such as Microsoft Operating System and/or pre-installed software applications, or media content. Actual formatted capacity may vary.

KiB: A kibibyte (KiB) means 2^{10}, or 1,024 bytes, a mebibyte (MiB) means 2^{20}, or 1,048,576 bytes, and a gibibyte (GiB) means 2^{30}, or 1,073,741,824 bytes.

Read and write speed may vary depending on the host device, read and write conditions, and file size.

TRADEMARKS

NVM Express and NVMe are registered or unregistered marks of NVM Express, Inc. in the United States and other countries.

PCI Express, PCIe and PCI are trademarks or registered trademarks of PCI-SIG.

Other company names, product names and service names may be trademarks of third-party companies.

©2023 Software-Enabled Flash Project. All Rights Reserved

Contents

©2023 Software-Enabled Flash Project. All Rights Reserved

©2023 Software-Enabled Flash Project. All Rights Reserved

1 Revision History

Version	Date	Description of change(s)
1.14	2023.08.30	Initial version of the document

©2023 Software-Enabled Flash Project. All Rights Reserved

2 Abstract

2.1 Introduction

NVM Express™ (NVMe™) Base specification defines an interface for host software to communicate with nonvolatile memory subsystems over a variety of memory-based transports and message-based transports.

This document defines a specific NVMe I/O Command Set, the Software-Enabled Flash (SEF) Command Set, which extends the NVMe Base Specification.

2.2 Scope

Figure 1 of NVMe 2.0 shows the relationship of the NVM Express specifications. This document belongs to "Command Set Specification" like NVM or Key Value.

This specification supplements the NVMe 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 SEF 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.

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 Base Specification
3. NVMe transport specifications
4. NVMe I/O command set specifications
5. NVMe-MI specification

©2023 Software-Enabled Flash Project. All Rights Reserved

2.3 References

- NVM Express Base Specification Revision 2.0 (NVMe 2.0)
- NVM Express NVM Command Set Specification Revision 1.0 (NVM 1.0)

2.4 Features

- Commands are added based on NVMe 2.0
  - Additional OPC, fields and others are subject to change
- Read/Write in units of ADU
- Supports Nameless Write / Copy
- Performance control on a per QoS Domain basis
- Supports Patrol

2.5 Class Code

In order to distinguish SEF devices from standard NVMe, ZNS, or other devices, an undefined value of ABh is used for the field of CC.PI. This is tentative, and will be changed to the standard 02h after the SEF Command Set is submitted to the standards organization.

Table 2.1: PCI Configuration Space - Class Code (CC: Offset 09h)

Bits	Type	Reset	Description
07:00	RO	02h or ABh	**Programming Interface (PI):** This field specifies the controller uses the NVM Express (02h) or the SEF programming interface (ABh).

2.6 Command Sets Supported

SEF devices shall set `CAP.CSS bit 6` to `1` to indicate the device supports one or more I/O Command Sets. Refer to Figure 36 in NVMe 2.0.

©2023 Software-Enabled Flash Project. All Rights Reserved

3 Geometry

3.1 Physical Geometry

Figure 3.1 depicts physical geometry of typical flash devices. Abstracted geometry described in the next section is used in Software-Enabled Flash (SEF) commands which specify flash addresses for the maximum performance.

Main differences from the abstracted geometry are the location of plane and the existence of string. The number of pages per string for SLC / MLC / TLC / QLC is 1 / 2 / 3/ 4 respectively.

There are two modes to record data, general and pSLC. The general mode writes all blocks with the same ID in each plane simultaneously, and uses all pages in every string. pSLC mode works like SLC, writing planes independently, and using only one page per string.

©2023 Software-Enabled Flash Project. All Rights Reserved

Figure 3.1: Physical Geometry

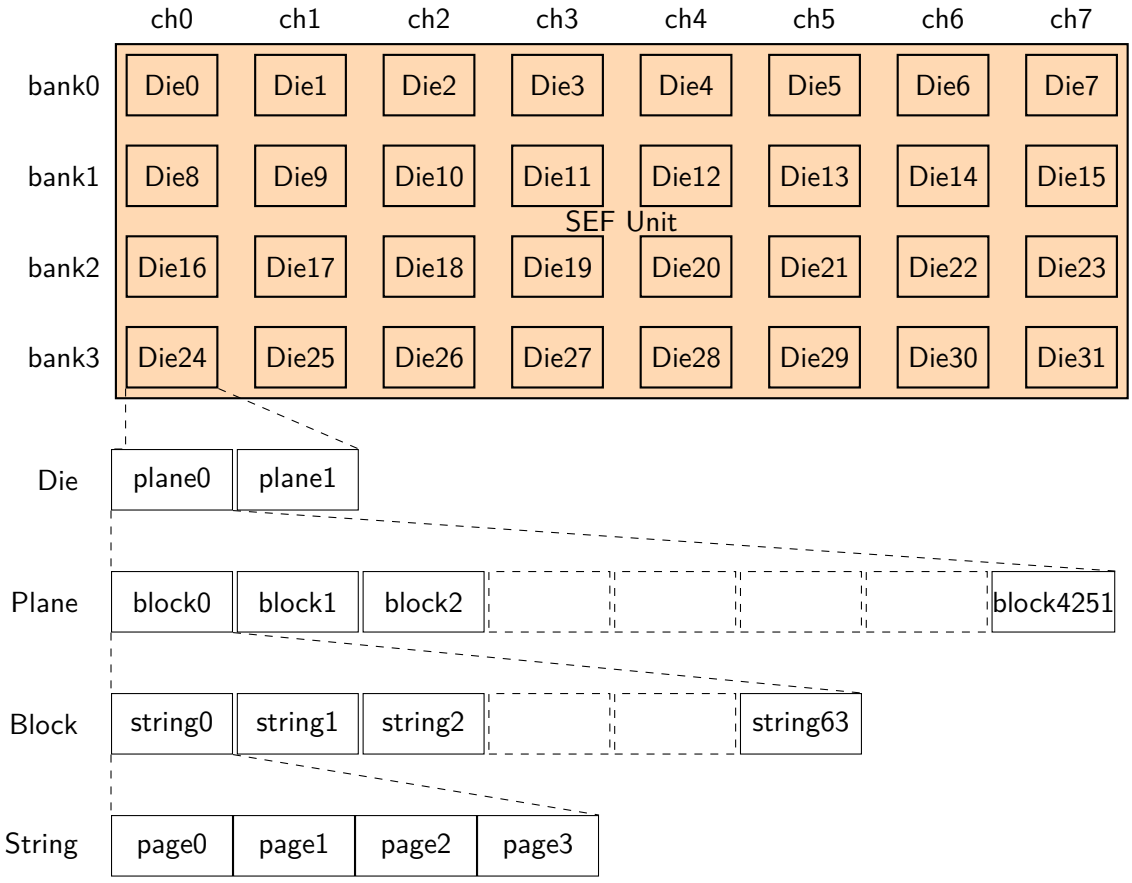

©2023 Software-Enabled Flash Project. All Rights Reserved

3.2 Abstracted Geometry

Figure 3.2: Abstracted Geometry

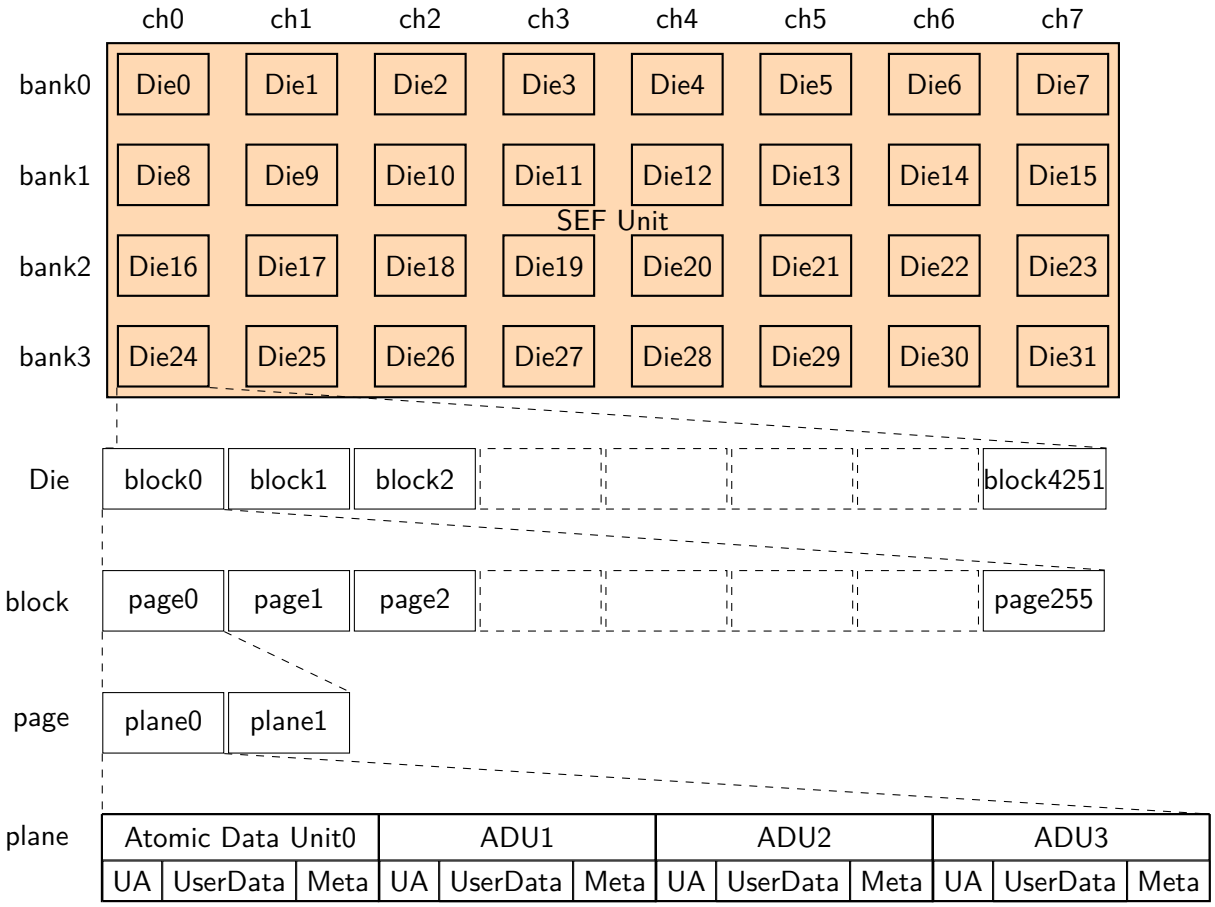

Table 3.1: Size of Elements

Element	Min	Max	Example	Remarks
channels(ch) / SEF	1	64	8	Fixed size by device
banks / SEF	1	32	4	Fixed size by device
blocks / Die	1	16,384	4,252	Fixed size by device
pages / block	128	8,192	256	Fixed size by device
planes / page	1	64	2	Fixed size by device
plane	16KiB	1MiB	16KiB	Fixed size by device
User Address (UA)	8 bytes	8 bytes	8 bytes	Fixed size
User Data	4KiB	1MiB	4KiB	Variable size by QoS Domain Should be the plane size of one over power of 2, or of times of power of 2

©2023 Software-Enabled Flash Project. All Rights Reserved

metadata	0	4KiB	16 bytes	The maximum depends on the User Data size

The fixed sizes above are given in the SEF Command Set specific Identify Controller data. Available sets of User Data and metadata size are listed in the LBA Format Data Structure in the SEF Command Set Identify Namespace data.

In commands and data structures in this document, capacity-related fields are mostly expressed in the number of planes as it is the smallest fixed-size element.

3.3 Die

Refers to a Flash memory chip. In NVMe 2.0, the term "Channel Media Unit" is used.

Each Die has a unique ID of `N x m + n`, where the number of channels and banks are N and M respectively, and the Die is located at the channel number and bank number of `n (0 <= n < N)` and `m (0 <= m < M)` respectively.

3.4 Virtual Device

In NVMe 2.0, the term "Endurance Group" is used.

Virtual Device (VD) is an execution unit which consists of an arbitrary number of dies.

Virtual Devices are created by issuing a Capacity Management command. Before creating Virtual Devices, Capacity Configuration data should be registered using Set Features command with Capacity Configuration Feature ID, if a pre-defined configuration is not used.

All Virtual Devices on a SEF device are created at deployment and remain unchanged throughout device lifetime.

Each Virtual Device has a unique ID between 1 and the number of dies.

Virtual Devices are internally managed and controlled with Super Blocks, Super Pages in a Super Block, and Atomic Data Units (ADUs) in a Super Page.

©2023 Software-Enabled Flash Project. All Rights Reserved

Figure 3.3: Virtual Device

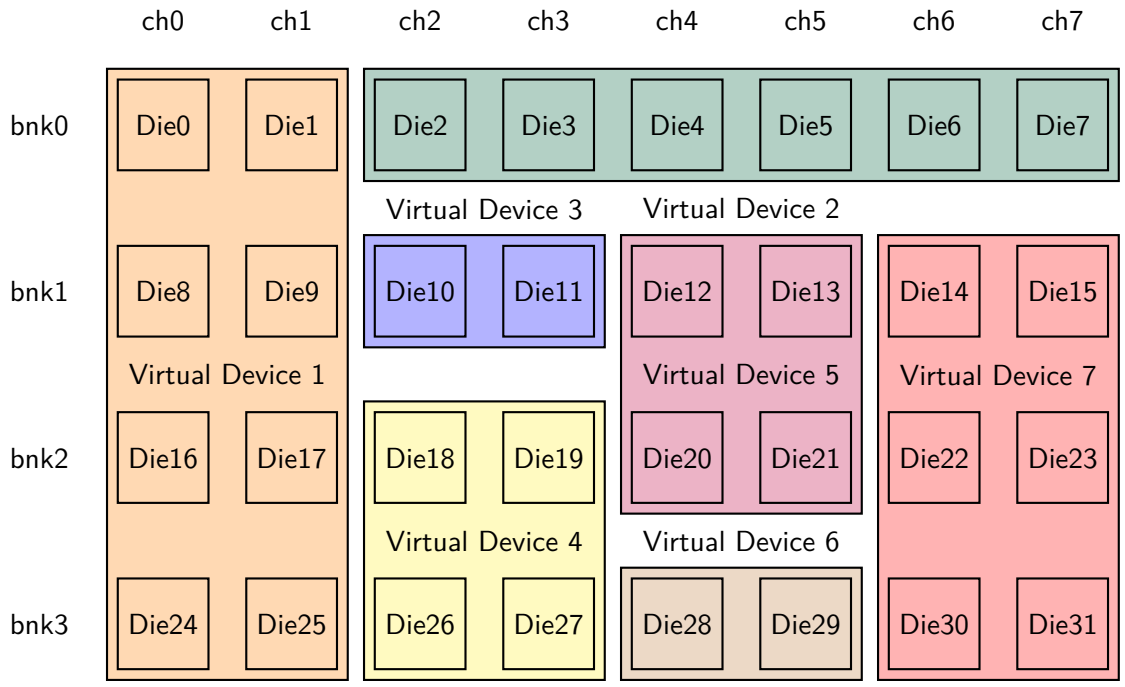

3.5 Super Block

There are two types of Super Block, a general Super Block, and a pSLC Super Block.

A general Super Block consists of one or more physical blocks in general mode, each from different physical planes of dies in a Virtual Device, while a pSLC Super Block consists of one or more physical blocks in pSLC mode, each from different physical plane per die in a Virtual Device. Therefore, the size of a healthy pSLC Super Block is `1 / n` of a healthy general Super Block, where `n` is the number of bits per physical cell (e.g. 3 for TLC, 4 for QLC, . . .).

The number of dies involved in a Super Block is set to the number of dies in the Virtual Device by default, and it can be modified using Set Features command with Virtual Device Feature ID before any QoS Domains are created in the Virtual Device. Once a QoS Domain has been created in a virtual device, the number cannot be changed.

Individual Super Block capacity may gradually decrease over time due to media defects depending on device mode.

3.6 Super Page

A Super Page consists of physical pages with the same page number in each physical block of a Super Block. Thus, all Super Blocks in a device have the same number of Super Pages regardless the number

©2023 Software-Enabled Flash Project. All Rights Reserved

of physical blocks a Super Block occupies. Depending on device mode, there is a case where the size of Super Pages vary when some planes in a Super Page are defective.

3.7 Atomic Data Unit (ADU)

A unit of user data and associated metadata that is written or read atomically. The size is given in the FLBAS field in the Identify Namespace Data.

3.8 User Address

Eight-bytes of data associated with read and write commands. The User Address given in a write command is written along with an ADU. The one given with a read command is checked with the one read from flash memory to make sure that the data being returned is the data that was requested.

User Address may be used as a key for a Look Up Table (LUT). Stored User Address data may also be used to rebuild the LUT if it is lost after an unsafe shutdown.

3.9 QoS Domain

An access entity within a Virtual Device. The term "Namespace" in NVMe is treated as a QoS Domain. QoS Domains are created by using the Namespace Management command.

All data access specifies the QoS Domain ID for the operation. In order to isolate QoS Domains, when multiple QoS Domains are within the same Virtual Device, data for different QoS Domains is never mixed within a Super Block. Super Blocks are individually allocated to QoS Domains, however over the lifetime of the device a Super Block may be released from one QoS Domain and later allocated to another QoS Domain.

3.10 Placement ID

Placement IDs are used in each QoS Domain for Flash Address Request and Nameless Write commands. Data with a different Placement ID is stored in a different Super Block.

IDs are valid within a QoS Domain and the number of Placement IDs is specified when creating a QoS Domain. Possible Placement ID for a command is between 0 and one smaller than the number of Placement IDs.

©2023 Software-Enabled Flash Project. All Rights Reserved

3.11 Flash Address (FLA)

Figure 3.4: Flash Address

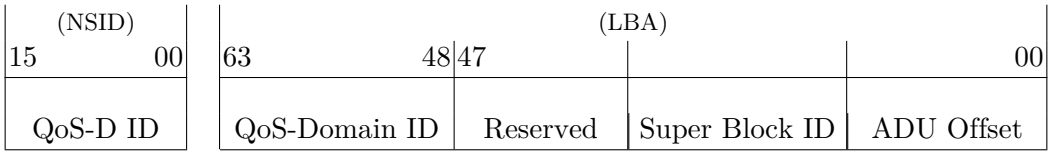

Figure 3.4 shows the Flash Address format. A 16-bit QoS Domain ID uses lower Namespace ID (NSID) field specified by NVMe. For certain commands, the QoS Domain field may be specified as `FFFFh` to denote all QoS Domains.

The lower 48-bit field of LBA field consists of Super Block ID in the upper and ADU Offset in the lower portions, and the remaining part is reserved. Bit lengths may be different Virtual Device by Virtual Device, and are provided in the Virtual Device Information Log page. The bit length of ADU Offset are up to 32 bits.

ADU Offset is constructed with ADU number, Plane number, Die number and Page number in low-to-high order. Note that each element is not always a power of two.

Figure 3.5: Elements constructing ADU Offset

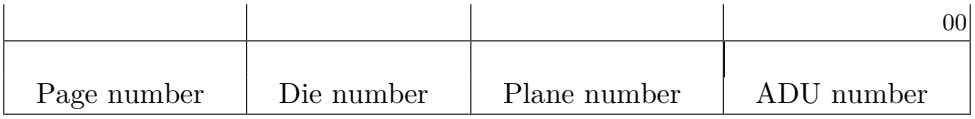

There are three methods, Perfect, Fragmented, and Packed, to express ADU Offset when one or more blocks in a Super Block are unavailable due to bad blocks or other reasons. Methods supported are given with SEF Command Set Identify Controller data. The method may vary by QoS Domain, and is specified on creation of a QoS Domain.

3.11.1 Perfect Method

A method to keep all Super Blocks in a QoS Domain the same size using alternate blocks when bad blocks are encountered.

3.11.2 Fragmented Method

A method in which ADU Offset is non-contiguous and the missing plane is provided with a Defective Bitmap.

The bit length of the Defective Bitmap is (the number of Planes per Block) x (the number of Dies), and the map provides missing Plane number with '1'. The corresponding Plane numbers are not used. It is

©2023 Software-Enabled Flash Project. All Rights Reserved

provided by the device to the host and is included in ARR in Flash Address Request command, ACO in Nameless Write command and return data of Erase in Super Block Management command.

Figure 3.6: Defective Bitmap layout of n Dies and m Planes

nm-1 (n-1)m		2m-1 m	m-1 0
Die n-1 DBM	--	Die 1 DBM	Die 0 DBM

3.11.3 Packed Method

A method where ADU Offset is always contiguous even when bad blocks exist. Unlike Fragmented method, the Defective Bitmap is not required for normal operations.

3.12 Super Block Assignment

Assigning a Super Block involves the following parameters specified when creating a QoS Domain.

- GCap: Guaranteed number of Planes.
- Quota: Maximum number of Planes.

A new QoS Domain can be created if both GCaps, general and pSLC Super Blocks, can be guaranteed.

A Super Block beyond GCap can be allocated if the remainder of the Virtual Device is available and in-use number of Planes is less than its Quota. The remainder is the difference between the Virtual Device's current capacity and the total tolerated GCap. The tolerated GCap is GCap plus one less than the number of healthy (w/o defect) Super Block's Planes.

This allows for the implementation of thinly provisioned QoS Domains.

©2023 Software-Enabled Flash Project. All Rights Reserved

4 | Super Block State

Figure 4.1: Super Block State Diagram

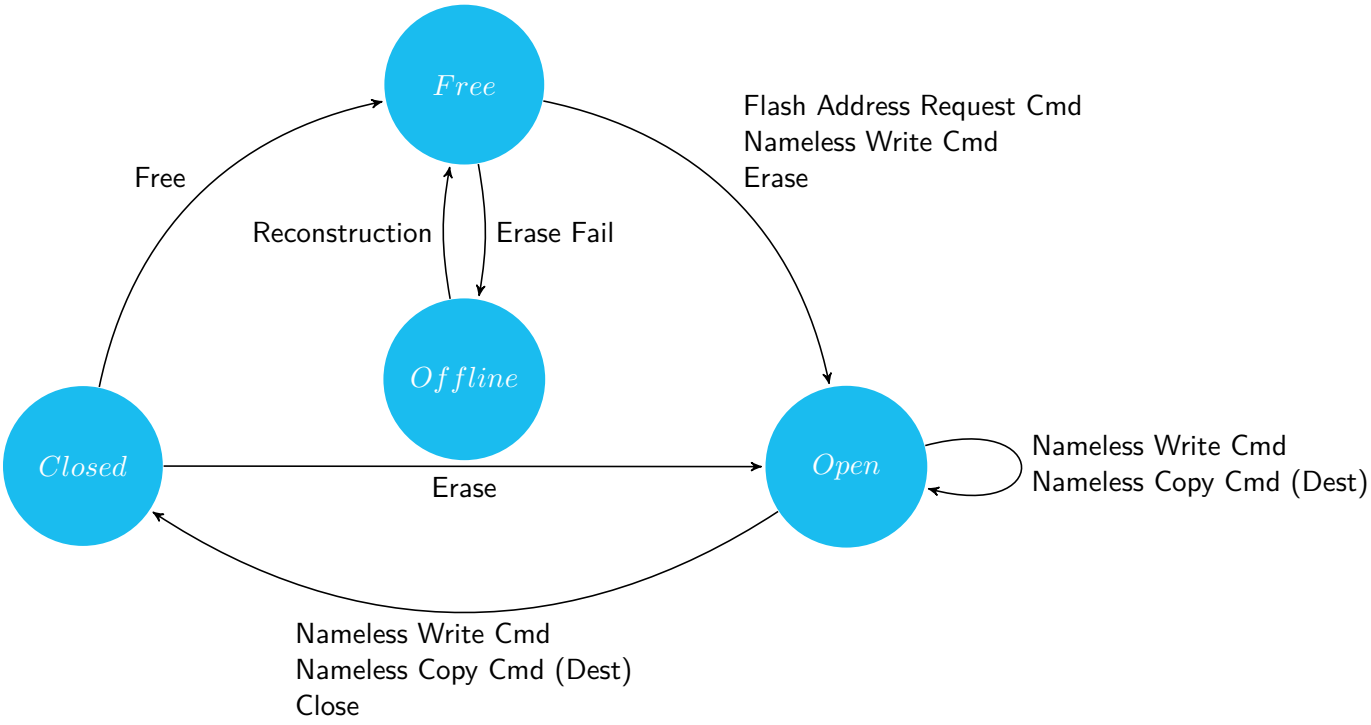

4.1 Free

This is the initial state for Super Blocks.

'Closed' Super Block transits to 'Free' by Free operation of a Super Block Management command.

4.2 Open

This is the partially programmed Super Block state.

©2023 Software-Enabled Flash Project. All Rights Reserved

'Free' Super Block transits to 'Open' by either Flash Address Request / Nameless Write command or Erase operation of a Super Block Management command.

Internally, there are two 'Open' states.

Open for Nameless Write: A Super Block dedicated to Nameless Write (NLW). The Super Block transits to this state by a Flash Address Request / Nameless Write command without explicit Super Block ID.

Open by Erase: A Super Block opened by a Super Block Management Command. This Super Block can be used for Nameless Write and Nameless Copy by specifying explicit Super Block ID.

The maximum number of simultaneous open blocks is determined per the device, or per each Virtual Device. The maximum number for a QoS Domain is specified at creation and can be changed anytime through Set Features command on condition that the total maximum number of QoS Domains doesn't exceed the maximum number of the device or the Virtual Device. When a QoS Domain tries to open a Super Block exceeding its limit, the Super Block least recently opened is closed or the command aborts.

4.3 Closed

This is the state of Super Blocks which retain effective data after all Super Pages are programmed, the user explicitly issues a close operation, or the device forces the Super Block to close.

Currently the device may force a Super Block to the "closed" state due to either internal hardware timing constraints on how long the Super Block has been in the "open" state or due to internal resource constraints such as exceeding the maximum number of open Super Blocks. In the case of resource constraints, the device may select the Super Block to close by any algorithm it desires such as the oldest Super Block currently open or the least recently written Super Block.

4.4 Offline

This is the state of Super Blocks which are out of service.

©2023 Software-Enabled Flash Project. All Rights Reserved

5 | Command Scheduling

The Software-Enabled Flash (SEF) Device controls priorities of host command processing through three schedulers.

Figure 5.1: Command Scheduling

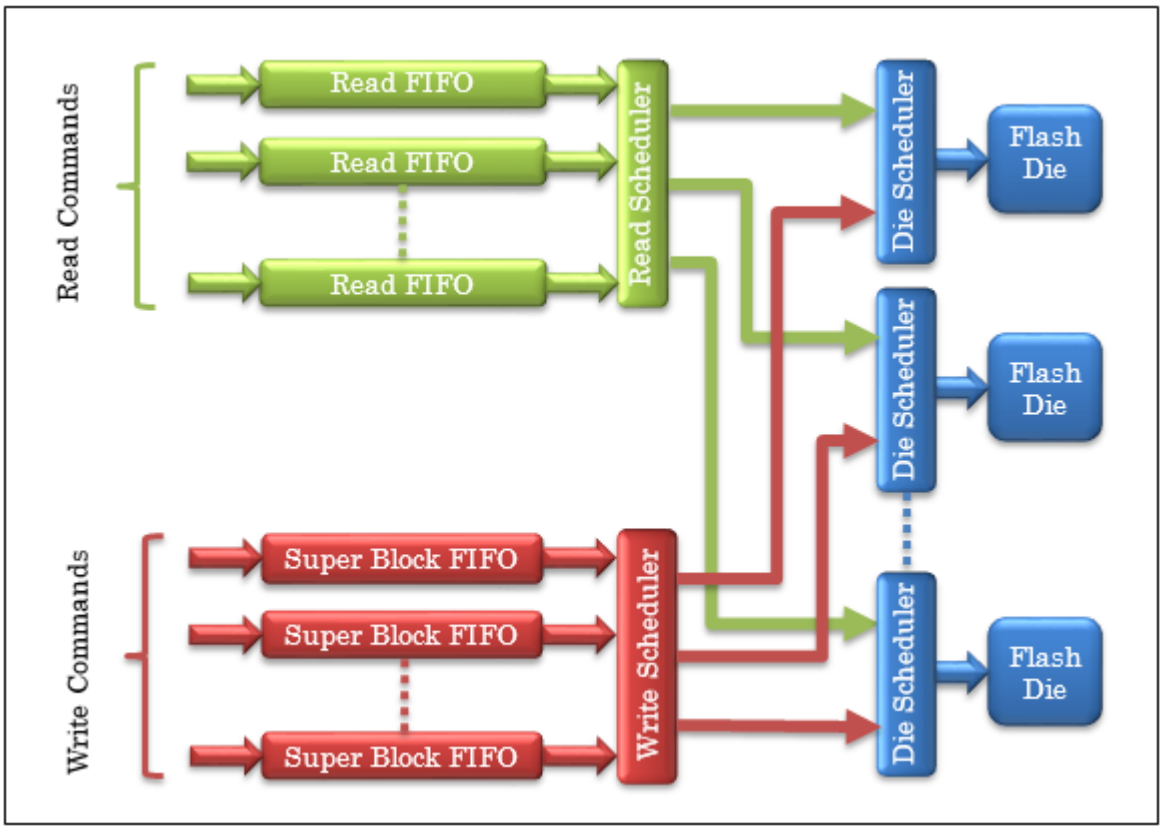

1. The Read Scheduler can operate as a strict priority based scheduler, a round robin scheduler, or a WFQ scheduler depending on the assigned weights, and send its parsed per-die read commands to the Die Scheduler. When Virtual Devices are created by the Capacity Management command, a default Read FIFO is assigned to each Virtual Device. Read FIFOs can be attached, detached, or changed by the Set Features command with Virtual Device Management FID. A default Read FIFO can also be assigned to a QoS Domain.

©2023 Software-Enabled Flash Project. All Rights Reserved

2. The Write Scheduler can operate as a strict priority based scheduler, a round robin scheduler, or a WFQ scheduler depending on the assigned weights, and send its parsed per-die write commands to the Die Scheduler.
3. The Die Scheduler determines a read or write command, examining their weights, and activates the die. When a new higher priority read command comes while a lower priority write or erase command is in progress, the read command may be processed by suspending the current operation.
4. Smaller weight is higher priority, and the number of ADUs in each FIFO in a certain period of time is, on the whole, processed in proportion to the reciprocal of the weight.
5. Write/Erase weights are on a QoS Domain basis, and specified when QoS Domains are created.
6. The weights can be overridden in each individual I/O command if necessary.

©2023 Software-Enabled Flash Project. All Rights Reserved

6 | Admin Commands for the Software Enabled Flash (SEF) Command Set

Table 6.1 shows a list of Admin commands the Software-Enabled Flash (SEF) devices support.

Table 6.1: Opcodes for Admin Commands

Opcode	Command	Remarks
00h	Delete I/O Submission Queue	NVMe 2.0 compatible
01h	Create I/O Submission Queue	NVMe 2.0 compatible
02h	Get Log Page	Added SEF specific pages
04h	Delete I/O Completion Queue	NVMe 2.0 compatible
05h	Create I/O Completion Queue	NVMe 2.0 compatible
06h	Identify	Added SEF specific fields
08h	Abort	NVMe 2.0 compatible
09h	Set Features	Added SEF specific features
0Ah	Get Features	Added SEF specific features
0Ch	Asynchronous Event Request	Added SEF specific events
0Dh	Namespace Management	NVMe 2.0 compatible
10h	Firmware Commit	NVMe 2.0 compatible
11h	Firmware Image Download	NVMe 2.0 compatible
15h	Namespace Attachment	NVMe 2.0 compatible
20h	Capacity Management	NVMe 2.0 compatible

Green highlighted fields in this chapter are NVMe 2.0 compatible.

©2023 Software-Enabled Flash Project. All Rights Reserved

6.1 Asynchronous Event Notification Request command

Super Block state change events and Capacity events are added to the I/O Command specific status, and they are identified with the Log Page Identifier in the CQE DW0. The capacity events include capacity reduction event and out of Super Block event, and they are identified with the Virtual Device Information log page.

- Super Block State Change Event: Indicates the state of a Super Block has changed. The closure event is notified through this command only when it happens due to the circumstances of the device. The closures due to explicit NLW/NLC/Flush/Close operations by the host are not notified with this command, but with the status of the NLW/NLC commands.
- Capacity Reduction Event: Indicates the capacity of a Virtual Device was decreased due to a physical block failure and has become less than the total guaranteed capacity of its QoS Domains. The Virtual Device Information log pages shows '1' for Capacity Reduction bit of Critical Warning field, and VDTCAP $<$ VDGCAP.
- Out of Super Block Event: Indicates the number of free Super Blocks in a Virtual Device has become the number of its QoS Domains or less. The Virtual Device Information log page shows '1' for either or both of Out of Super Block bit of Critical Warning field, and SBSTA.NFRSB $<$ ASS_QOSD and/or PSBSTA.NFRSB $<$ ASS_QOSD.

Example of host handling of Super Block State Change event

1. Identify Virtual Device ID, Super Block ID and the transition from the Super Block Information log page with CDW11.LSI=FFFFh.
2. When reliability index of a Super Block changes to Unknown, the host issues Patrol command on it.
3. When a Super Block in Open State for Nameless Write transits to Closed state, the host determines if Garbage Collection is needed.
4. When a Super Block in Open state for Nameless Copy transits to Closed state, the host reflects LUT with the accumulated ACRs, and determines if Garbage Collection is needed.
5. Issue an Asynchronous Event Notification Request command, if necessary.
6. If CDW10.RAE=1 at step (1), reset the event by reading the Super Block Information log page with CDW10.RAE=0. If another Super Block state change exists, another event is raised at this timing.

Example of host handling of Capacity Reduction event

1. Identify Virtual Device ID and the excess number of Planes from the Virtual Device Information log page with CDW11.LSI=FFFFh.

©2023 Software-Enabled Flash Project. All Rights Reserved

2. Determine QoS Domain(s) whose guaranteed capacity should be reduced, and reduce its (their) guaranteed capacity by issuing Set Features command with Feature ID = QoS Domain Management.
3. Issue an Asynchronous Event Notification Request command, if necessary.
4. If CDW10.RAE=1 at step (1), reset the event by reading the Virtual Device Information log page with CDW10.RAE=0. If the capacity is still short, another event is raised at this timing.

Example of host handling of Out of Super Block event

1. Identify Virtual Device ID and the excess number of Super Blocks from the Virtual Device Information log page with CDW11.LSI=FFFFh.
2. Determine QoS Domain(s), Super Blocks of which should be freed, and do garbage collection to free Super Blocks.
3. Issue an Asynchronous Event Notification Request command, if necessary.
4. If CDW10.RAE=1 at step (1), reset the event by reading the Virtual Device Information log page with CDW10.RAE=0. If the number of free Super Block is still less, another event is raised at this timing.

6.1.1 Command Completion

The information about this event is returned in Completion Queue Entry Dword 0.

Table 6.2: SEF Capacity Event - Completion Queue Entry Dword 0

Bits	Description
31:24	Reserved
23:16	**Log Page Identifier:** Indicates the log page associated with the asynchronous event. This log page needs to be read by the host to clear the event. D0h: SEF Capacity event D3h: Super Block State Change event
15:08	**Asynchronous Event Information (AEI):** Returns one of the following value for IO Command Set Specific Status. 40h is for SEF Command Set Status. Value / Description: 00h-3Fh: Defined in standard NVMe family specifications 40h: SEF Command Set Specific Status (Tentative) 41h-FFh: Reserved
07:03	Reserved

©2023 Software-Enabled Flash Project. All Rights Reserved

02:00	**Asynchronous Event Type:** Indicates the type of the asynchronous event. More specific information on the event is provided in the Asynchronous Event Information field. 110b is used for SEF.

Value	Description
000b	Error status
001b	SMART / Health status
010b	Notice
011b to 101b	Reserved
110b	I/O Command Set specific status
111b	Vendor Specific

6.2 Capacity Management command

Host software uses the Capacity Management command to configure Endurance Groups (Virtual Devices) by selecting one of a set of supported configurations (i.e. Fixed Capacity Management). SEF Command Set doesn't support Variable Capacity Management. The Capacity Management command specifies the operations defined in Table 6.3.

The Capacity Management command uses the Command Dword 10.

Table 6.3: Capacity Management - Command Dword 10

Bits	Description
31:16	**Element Identifier:** This field contains a value specific to the value of the Operation field. Specify Capacity Configuration ID (0, 1, 2, or 3) for Select Capacity Configuration. ID 0 is for deletion, 1 is built-in, and 2 and 3 are configurable.
15:04	Reserved
03:00	**Operation:** Specifies the operation to be performed by the controller:

Value	Description
0h	**Select Capacity Configuration:** Endurance Groups and NVM Sets are configured as indicated by the Capacity Configuration Descriptor specified by this command.
1h - 4h	Not supported by the SEF Command Set
5h - Fh	Reserved

If the Element Identifier field is cleared to '0', then the controller shall clear the configuration. It can be done only when Delete Endurance Group (Bit13) in CTRATT (Byte99-96) in the Identify Controller Data is '1'. If a QoS Domain exists, the controller aborts the command. If the Element Identifier field specifies the current configuration, then the controller shall complete the command without error and

©2023 Software-Enabled Flash Project. All Rights Reserved

shall make no change to the capacity configuration. If the Element Identifier field is not 0, and specifies other than the current configuration, then the controller aborts the command.

6.3 Get/Set Features command

Table 6.4 defines Features that are added to NVMe 2.0.

Table 6.4: Feature Identifiers

Feature Identifier	Persistent Across Cycle and Reset	Uses Memory Buffer for Attributes	Description
05h	No	No	Error Recovery (From NVM1.0 specification)
D0h	Yes	Yes	Virtual Device Management
D1h	Yes	No	QoS Domain Management
D2h	Yes	Yes	Capacity Configuration Registration

6.3.1 Error Recovery (Feature Identifier 05h)

NVM 1.0 defines that the default value of DULBE is '0' and host software shall only enable this error if it is supported for this namespace as indicated in the Namespace Features field of the Identify Namespace data structure. The default value for this field shall be '0'. In SEF, the Namespace Features field of the Identify Namespace data structure indicates it is supported, however, DULBE is always '1' and unchangeable.

Table 6.5: Error Recovery - Command Dword 11

Bits	Description
31:17	Reserved
16	**Deallocated or Unwritten Logical Block Error Enable (DULBE):** If set to '1', then the Deallocated or Unwritten Logical Block error is enabled for the namespace specified in CDW1.NSID. '1' is always returned for Get Features and unchangeable against Set Features.
15:00	**Time Limited Error Recovery (TLER):** See Figure 84 in NVM 1.0

6.3.2 Virtual Device Management (Feature Identifier D0h)

Set Features sets Virtual Device settings, and Get Features returns information of a Virtual Device. CDW10.SV of Set Features is ignored and settings are always saved.

©2023 Software-Enabled Flash Project. All Rights Reserved

Table 6.6: Virtual Device Management - Command Dword 11

Bits	Description
31:04	This field depends on CDW11.SEL. If not indicated, the field is reserved.
03:00	**Select (SEL):** This field selects the type of operation to perform.

Value	Set Features	Get Features
0h	Set Number of Dies	Reserved
1h	Attach Read FIFO	Reserved
2h	Change Read FIFO	Reserved
3h	Detach Read FIFO	Reserved
4h	Set Number of pSLC Blocks	Reserved
5h	Control Asynchronous Event	Reserved
Fh	Reset Scheduler	Reserved
6h-Eh	Reserved	

Set Number of Dies

Defines the number of Dies to compose a Super Block for the Virtual Device specified with CDW1.NSID. This operation is executable only before a QoS Domain is created.

Set Number of Dies uses CDW11.

Table 6.7: Set Number of Dies - Command Dword 11

Bits	Description
31:16	Reserved
15:08	**Number of Dies (ND)**: Indicates the number of Dies constructing a Super Block. This is a 0'sbased value. The number must be a divisor of the number of Dies in the Virtual Device. If it isn't, then the controller aborts the command.
07:04	Reserved
03:00	**Select (SEL):** '0' for Set Number of Dies.

Attach Read FIFO

Attaches a Read FIFO specified with CDW12.RFID to the Virtual Device specified with CDW1.NSID.

When SEFCAP.bit13 in the Identify SEF Command Set Controller data is cleared to '0', and the Read FIFO is already attached to another Virtual Device, the command aborts.

Attach Read FIFO uses CDW11 and 12.

©2023 Software-Enabled Flash Project. All Rights Reserved

Table 6.8: Attach Read FIFO - Command Dword 11

Bits	Description
31:16	**Weight (WT):** Read Weight
15:04	Reserved
03:00	**Select (SEL):** '1' for Attach Read FIFO

Table 6.9: Attach Read FIFO - Command Dword 12

Bits	Description
31	**Default (D):** If set to '1' the Read FIFO is set as a default. If cleared to '0', the default setting stays.
30:00	**Read FIFO Identifier (RFID):** Specifies Read FIFO ID to attach or change.

Change Read FIFO

Changes the weight and/or default setting for the Read FIFO specified with CDW12.RFID. If the Read FIFO is not attached to the Virtual Device, the command aborts. Change Read FIFO uses CDW11 and 12.

Table 6.10: Change Read FIFO - Command Dword 11

Bits	Description
31:16	**Weight (WT):** Read Weight
15:04	Reserved
03:00	**Select (SEL):** '2' for Change Read FIFO

Table 6.11: Change Read FIFO - Command Dword 12

Bits	Description
31	**Default (D):** If set to '1' the Read FIFO is set as a default. If cleared to '0', the default setting stays.
30:00	**Read FIFO Identifier (RFID):** Specifies Read FIFO ID to attach or change.

Detach Read FIFO

Detaches a Read FIFO specified with CDW12.RFID from the Virtual Device specified with CDW1.NSID. If the Read FIFO is set as a default of a Virtual Device or a QoS Domain, the command is aborted. Detach Read FIFO uses CDW11 and 12.

©2023 Software-Enabled Flash Project. All Rights Reserved

Table 6.12: Detach Read FIFO - Command Dword 11

Bits	Description
31:16	**Weight (WT):** Read Weight
15:04	Reserved
03:00	**Select (SEL):** '3' for Detach Read FIFO

Table 6.13: Detach Read FIFO - Command Dword 12

Bits	Description
31:00	**Read FIFO Identifier (RFID):** Specifies Read FIFO ID to detach. If the specified Read FIFO is not attached to the Virtual Device, or is used as a default Read FIFO by the Virtual Device or any QoS Domains, the controller aborts the command.

Set Number of pSLC Blocks

Defines the number of regular Blocks per Die to transform to pSLC on the Virtual Device specified with CDW1.NSID. The number of pSLC blocks per die will be the-number-of-Plane times of the specified number.

If either of GCaps for general and pSLC Super Blocks is not satisfied, the command is aborted.

Set Number of pSLC Blocks uses CDW11 and 12.

Table 6.14: Set Number of pSLC Blocks - Command Dword 11

Bits	Description
31:04	Reserved
03:00	**Select (SEL):** '4' for Set Number of pSLC Blocks

Table 6.15: Set Number of pSLC Blocks - Command Dword 12

Bits	Description
31:00	**Number of pSLC Blocks (NPBLK):** Specifies the number of pSLC physical blocks per die. This is not a 0's based value. If the value exceeds the number of available pSLC physical blocks, all Super Blocks are configured in pSLC mode.

Control Asynchronous Event

Enables or disables asynchronous events of the Virtual Device specified with CDW1.NSID.

When disabled, Asynchronous Event Notification Request command doesn't return for the event, but the

©2023 Software-Enabled Flash Project. All Rights Reserved

behavior of the corresponding Get Log Page command is the same as when enabled.

Control Asynchronous Event uses CDW11.

Table 6.16: Control Asynchronous Event - Command Dword 11

Bits	Description
31:24	Reserved
23:16	**Critical Warning Control:** Controls the settings of the critical warnings. The value of '0' disables the event, while '1' enables. The default is '0'.

Bits	Definition
7:3	Reserved
2	Capacity Reduction Event
1	Out of pSLC Super Block Event
0	Out of general Super Block Event

15:14	Reserved
15:14	**Select (SEL):** '5' for Control Asynchronous Event

Reset Scheduler

Resets schedulers specified with CDW11.SCH of the Virtual Device specified with CDW1.NSID. If FFFFh is specified for CDW1.NSID, all Virtual Devices are the targets.

This is for maintenance and will be deleted from formal publication.

Reset Scheduler uses CDW11.

Table 6.17: Reset Scheduler - Command Dword 11

Bits	Description
31:08	Reserved
07:04	**Scheduler (SCH):** Specify schedulers to reset.

Bits	Description
7	Internal Read Scheduler
6	Read Scheduler
5	Write Scheduler
4	Die Scheduler

03:00	**Select (SEL):** 'Fh' for Reset Scheduler

Command Completion

©2023 Software-Enabled Flash Project. All Rights Reserved

Table 6.18: Virtual Device Management - Completion Queue Entry Dword 0

Bits	Description
31:16	Reserved
15:00	**Virtual Devices ID (VDID):** Virtual Device ID processed.

6.3.3 QoS Domain Management (Feature Identifier D1h)

Set Features manages a QoS Domain and Get Features returns a Root Pointer of a QoS Domain. The type of operation is specified in CDW11.SEL. CDW10.SV of Set Features is ignored and settings are always saved.

Table 6.19: QoS Domain Management - Command Dword 11

Bits	Description
31:04	This field depends on CDW11.SEL. If not indicated, the field is reserved.
03:00	**Select (SEL):** This field selects the type of operation to perform.

Value	Set Features	Get Features
0h	Set a Root Pointer	Get a Root Pointer
1h	Change Capacity	Reserved
2h	Change Weights	Reserved
3h	Change Read Deadline	Reserved
4h	Change Maximum Number of Open Super Blocks	Reserved
5h	Change Default Read FIFO	Reserved
6h-Fh	Reserved	

Set/Get a Root Pointer

Set Features sets and Get Features gets the Root Pointer specified with CDW11.RPID of the QoS Domain specified in CDW1.NSID.

64-bit Root Pointers are prepared for each QoS Domain. Their purposes are up to the host. Most likely usage is pointers to locations where important data such as management data is stored.

Table 6.20: Set/Get a Root Pointer - Command Dword 11

Bits	Description
31:08	Reserved
07:04	**Root Pointer ID (RPID):** Specified the Root Pointer ID to manage.
03:00	**Select (SEL):** '0' for Set/Get a Root Pointer

©2023 Software-Enabled Flash Project. All Rights Reserved

Table 6.21: Set a Root Pointer - Command Dword 12 and 13

Bits	Description
63:00	**Root Pointer (RPTR):** Specifies the Root Pointer to be set.

Change Capacity

This operation changes guaranteed capacity and/or quota of the QoS Domain ID specified in CDW1.NSID.

If CDW12.GCAP doesn't satisfy that of other QoS Domains in the Virtual Device, the controller aborts the command.

Change Capacity uses CDW11, CDW12 and CDW13.

Table 6.22: Change Capacity - Command Dword 11

Bits	Description
31:17	Reserved
16	**Type of Super Block:**'0' specifies general Super Block, '1' specifies pSLC Super Block
15:04	Reserved
03:00	**Select (SEL):** '1' for Change Capacity

Table 6.23: Change Capacity - Command Dword 12 and 13

Bits	Description
63:00	**Guaranteed Capacity (GCAP):** Specifies the number of guaranteed Planes to be used. This is not a 0's based value. When 0 is specified, the current GCAP is retained.

Table 6.24: Change Capacity - Command Dword 14 and 15

Bits	Description
63:00	**Quota (QUOTA):** Indicates the maximum number of Planes the QoS Domain may possibly use. This is not a 0's based value. When 0 is specified, the current Quota is retained. If QUOTA $<$ max(GCAP, UCAP), QUOTA is rounded up to max(GCAP, UCAP), where UCAP is the utilization (see Table 6.56).

Change Weights

This operation changes weights of the QoS Domain specified with CDW1.NSID. Change Weights uses CDW11 and 12.

©2023 Software-Enabled Flash Project. All Rights Reserved

Table 6.25: Change Weights - Command Dword 11

Bits	Description
31:04	Reserved
03:00	**Select (SEL):** '2' for Change Weights

Table 6.26: Change Weights - Command Dword 12

Bits	Description
31:16	**Erase Weight (EWT)**
15:0	**Write Weight (WWT)**

Change Read Deadline

This operation changes Read Deadline of the QoS Domain specified in CDW1.NSID. Change Read Deadline uses CDW11.

Table 6.27: Change Read Deadline - Command Dword 11

Bits	Description
31:10	Reserved
09:08	**Read Deadline (RDL):** Value: 0 = Fastest; 1 = Typical; 2 = Long; 3 = Heroic
07:04	Reserved
03:00	**Select (SEL):** '3' for Change Read Deadline

Change Maximum Number of Open Super Blocks

This operation changes the maximum number of Super Blocks the QoS Domain specified in CDW1.NSID is able to open simultaneously. Change Maximum Number of Open Super Blocks uses CDW11.

Table 6.28: Change Maximum Number of Open Super Blocks - Command Dword 11

Bits	Description
31:16	**Maximum Number of Open Super Blocks (MAXOSB):** This is not a 0's based value. If the number is less than or equal to the number of NPLID+2, NPLID+2 is set, where NPLID is in the SEF Command Set Identify Namespace data.

©2023 Software-Enabled Flash Project. All Rights Reserved

15:04	Reserved
03:00	**Select (SEL):** '4' for Change Maximum Number of Open Super Blocks

Change Default Read FIFO

This operation changes the default Read FIFO of the QoS Domain specified with CDW1.NSID. Change Default Read FIFO uses CDW11 and 12.

Table 6.29: Change Default Read FIFO - Command Dword 11

Bits	Description
31:04	Reserved
03:00	**Select (SEL):** '5' for Change Default Read FIFO

Table 6.30: Change Default Read FIFO - Command Dword 12

Bits	Description
31:00	**Read FIFO ID (RFID):** If 0 is specified, the default Read FIFO for the Virtual Device is used. If the specified Read FIFO is not 0, nor attached to the Virtual Device, the controller aborts the command.

Command Completion

Table 6.31: QoS Domain Management - Completion Queue Entry 0 and 1 for Get a Root Pointer

Bits	Description
63:00	**Root Pointer (RPTR):** Root Pointer

Table 6.32: QoS Domain Management - Completion Queue Entry 0 for other than Get a Root Pointer

Bits	Description
31:16	Reserved
15:00	**QoS Domain ID (QOSDID):** Indicates QoS Domain ID processed.

6.3.4 Capacity Configuration Registration (Feature Identifier D2h)

Set Features registers the Capacity Configuration specified with the Capacity Configuration Descriptor. The SEF unit has three sets of Capacity Configurations, ID=1 and 2, and ID=2 can be registered. When the ID currently selected is specified, the controller aborts the command. This feature additionally uses DPTR, CDW10 and CDW11.

©2023 Software-Enabled Flash Project. All Rights Reserved

Table 6.33: Capacity Configuration Registration - Command Dword 11

Bits	Description
31:16	**Number of Endurance Group Configuration Descriptors (EGCN):** This field indicates the number of Endurance Groups in the Capacity Configuration Descriptor.
15:00	**Total Number of Channel Media Units (TCHMUS):** This field indicates the number of Channel Media Units that appear in the Capacity Configuration Descriptor.

The Size of the Capacity Configuration Descriptor is as follows:
20h + EGCN x 58h + TCHMUS x 8h In the Capacity Configuration Descriptor, Endurance Group Configuration Descriptors are listed in ascending order by Endurance Group Identifier, and each Endurance Group Identifier shall appear only once. Get Features is not supported for this identifier. Use Get Log Page command with Supported Capacity Configuration List (Log Identifier 11h) instead.

Table 6.34: Capacity Configuration Registration - Capacity Configuration Descriptor

Bytes	Description
1:0	**Capacity Configuration Identifier:** This field indicates the identifier for this Capacity Configuration. Shall be '2'.
3:2	**Domain Identifier:** Domain is not supported, and must be '0'.
5:4	**Number of Endurance Group Configuration Descriptors (EGCN):** This field indicates the number of Endurance Group (Virtual Device) Configuration Descriptors in the list.
31:6	Reserved
NOTE 1	**Endurance Group Configuration 1 Descriptor:** This field indicates the first Endurance Group Configuration Descriptor in the list.
—	
NOTE 1	**Endurance Group Configuration EGCN Descriptor:** This field indicates the last Endurance Group Configuration Descriptor in the list, if any.
NOTES: 1. Endurance Group Configuration Descriptor may be different lengths.	

In the Endurance Group Configuration Descriptor, just one Channel Configuration Descriptor must be listed. It means all Dies of the SEF Unit are regarded to be connected to a single virtual channel.

Table 6.35: Endurance Group Configuration Descriptor

Bytes	Description

©2023 Software-Enabled Flash Project. All Rights Reserved

1:0	**Endurance Group Identifier (ENDGID):** This field indicates the identifier of the Endurance Group (Virtual Device) described by this Endurance Group Configuration Descriptor. This field shall indicate a value greater than or equal to 1h and less than or equal to the number of Dies in the SEF device.
3:2	**Capacity Adjustment Factor:** This field is not supported and must be '0'.
79:4	Reserved
NVM Set Identifiers	
81:80	**Number of NVM Sets (EGSETS):** NVM Sets are not supported and must be '0'.
Channel Configuration Descriptors	
83:82	**Number of Channels (EGCHANS):** In the Capacity Configuration Descriptor, all Media Units (Dies) are regarded to be connected to a single virtual channel, and must be '1'.
NOTE 1	**Channel 0 Configuration Descriptor:** This field contains the Channel Configuration Descriptor for the first and the only Channel in this Endurance Group.
NOTES: 1. Channel Configuration Descriptor may be different lengths.	

The Channel Configuration Descriptor lists the Media Units (Dies) attached to a Channel. Media Unit Configuration Descriptors are listed in ascending order by Media Unit Identifier, and each Media Unit Identifier shall appear only once throughout all Endurance Group Descriptors.

Table 6.36: Channel Configuration Descriptor

Bytes	Description
1:0	**Channel Identifier:** Not supported, and must be FFFFh.
3:2	**Number of Channel Media Units (CHMUS):** This field indicates the number of Media Units (Dies) that are attached to this Channel.
11:4	**Media Unit 0 Configuration Descriptor:** This field contains the Media Unit Configuration Descriptor for the first Media Unit attached to this Channel.
—	
CHMUS*8+3:CHMUS*8-4	**Media Unit CHMUS-1 Configuration Descriptor:** This field contains the Media Unit Configuration Descriptor for the last Media Unit attached to this Channel.

©2023 Software-Enabled Flash Project. All Rights Reserved

Table 6.37: Media Unit Configuration Descriptor

Bytes	Description
1:0	**Media Unit Identifier:** This field indicates the identifier of this Media Unit.
5:2	Reserved
7:6	**Media Unit Descriptor Length (MUDL):** This field contains the length in bytes of the descriptor information that follows. The total length of the Media Unit Configuration Descriptor in bytes is the value in this field plus 8. In controllers compliant to this revision of this specification, this field shall be cleared to 0h.

Figure 6.1: Example of a Capacity Configuration Descriptor

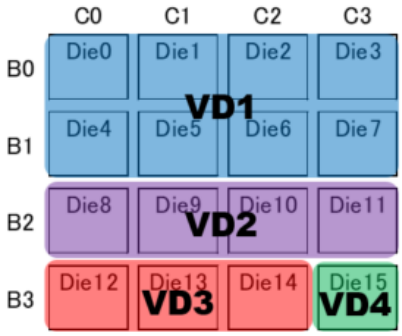

	00	01	02	03	04	05	06	07	08	09	0A	0B	0C	0D	0E	0F
0000	02	00	00	00	04	00	00	00	00	00	00	00	00	00	00	00
0010	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
0020	01	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
0030	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
0040	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
0050	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
0060	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
0070	00	00	01	00	FF	FF	08	00	00	00	00	00	00	00	00	00
0080	01	00	00	00	00	00	00	00	02	00	00	00	00	00	00	00
0090	03	00	00	00	00	00	00	00	04	00	00	00	00	00	00	00
00A0	05	00	00	00	00	00	00	00	06	00	00	00	00	00	00	00
00B0	07	00	00	00	00	00	00	00	02	00	00	00	00	00	00	00
00C0	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
00D0	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
00E0	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
00F0	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
0100	00	00	00	00	00	00	00	00	00	00	01	00	FF	FF	04	00
0110	08	00	00	00	00	00	00	00	09	00	00	00	00	00	00	00
0120	0A	00	00	00	00	00	00	00	0B	00	00	00	00	00	00	00
0130	03	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
0140	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
0150	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
0160	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
0170	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
0180	00	00	01	00	FF	FF	03	00	0C	00	00	00	00	00	00	00
0190	0D	00	00	00	00	00	00	00	0E	00	00	00	00	00	00	00
01A0	04	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
01B0	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
01C0	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
01D0	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
01E0	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
01F0	00	00	01	00	FF	FF	01	00	0F	00	00	00	00	00	00	00

©2023 Software-Enabled Flash Project. All Rights Reserved

6.4 Get Log Page command

The Get Log Page command returns a data buffer containing the log page requested.

The Get Log Page command uses the Data Pointer, Command Dword 10, Command Dword 11, Command Dword 12, Command Dword 13, and Command Dword 14 fields. All other command specific fields are reserved.

Table 6.38 illustrates Command Dword 11, which includes an additional Log Specific Identifier filed for Virtual Device Information Log Page.

Table 6.38: Get Log Page - Command Dword 11

<table>
<tr><th>Bits</th><th>Description</th></tr>
<tr><td>31:16</td><td>Log Specific Identifier (LSI): This field specifies an identifier that is required for a particular log page. The log pages that require a log specific identifier are indicated in the table below.
<table>
<tr><th>Log Page</th><th>Description</th></tr>
<tr><td>Virtual Device Information
Super Block Information
Read FIFO List</td><td>Endurance Group (Virtual Device) Identifier</td></tr>
</table></td></tr>
<tr><td>15:00</td><td>Number of Dwords (NUMDU)</td></tr>
</table>

Table 6.39 defines the Log pages that are added to NVMe 2.0.

Table 6.39: Get Log Page - Additional Log Page Identifiers

Log Identifier	Description
11h	Supported Capacity Configuration List
D0h	Virtual Device Information
D2h	Super Block List
D3h	Super Block Information
D4h	User Address List
D5h	Address Change Order
D6h	Read FIFO List

6.4.1 Supported Capacity Configuration List (Log Identifier 11h)

This log page is used to provide a list of Supported Capacity Configuration Descriptors. Each entry in the list defines a different configuration of Endurance Groups (Virtual Devices) supported by the device. SEF Devices support two configurations, ID=1 and 2. ID 1 is a predefined fixed configuration of Endurance Groups (Virtual Devices) with one die each. ID 2 can be registered through Set Features command with

©2023 Software-Enabled Flash Project. All Rights Reserved

Feature ID at D2h. Its initial configuration is up to the device.

6.4.2 Virtual Device Information (Log Identifier D0h)

This log page provides the information of the Virtual Device specified with Endurance Group Identifier in CDW11. If the Identifier is FFFFh, it returns the information of a Virtual Device which raised a critical warning. If no Virtual Device raised a warning, the controller aborts the command.

Table 6.40: Get Log Page - Virtual Device Information Structure

<table>
<tr><th>Bytes</th><th>Description</th></tr>
<tr><td>1:0</td><td>Virtual Device ID (VDID): This field indicates Virtual Device Identifier.</td></tr>
<tr><td>2</td><td>Critical Warning: This field indicates critical warnings for the state of the Virtual Device. Each bit corresponds to a critical warning type; multiple bits may be set to '1'. If a bit is cleared to '0', then that critical warning does not apply. Critical warnings may result in an asynchronous event notification to the host. Bits in this field represent the current associated state and are not persistent.
<table>
<tr><th>Bits</th><th>Definition</th></tr>
<tr><td>7:3</td><td>Reserved</td></tr>
<tr><td>2</td><td>If set to '1', then the total capacity (VDTCAP) has become greater than the total number of guaranteed capacity (VDGCAP).</td></tr>
<tr><td>1</td><td>If set to '1', then the number of free pSLC Super Blocks (PSBSTA.NFRSB) has become less than the number of its QoS Domains (ASS_QOSD).</td></tr>
<tr><td>0</td><td>If set to '1', then the number of free general Super Blocks (SBSTA.NFRSB) has become less than the number of its QoS Domains (ASS_QOSD).</td></tr>
</table></td></tr>
<tr><td>3</td><td>Critical Warning Control: This field indicates critical warning settings of the Virtual Device. If a bit is cleared to '0', then that critical warning event is disabled. If a bit is set to '1', then the event is enabled.
<table>
<tr><th>Bits</th><th>Definition</th></tr>
<tr><td>7:3</td><td>Reserved</td></tr>
<tr><td>2</td><td>Capacity Reduction Event</td></tr>
<tr><td>1</td><td>Out of pSLC Super Block Event</td></tr>
<tr><td>0</td><td>Out of general Super Block Event</td></tr>
</table></td></tr>
<tr><td>7:4</td><td>Current Erase Serial Number (CESN): Indicates the current number of erase operation. This field is cleared to zero when creating the Virtual Device, and is incremented and assigned to the Super Block when it is allocated.</td></tr>
<tr><td>8</td><td>Maximum P/E Count (MAX_PE): Indicates the maximum P/E count index[1].</td></tr>
</table>

©2023 Software-Enabled Flash Project. All Rights Reserved

9	**Average P/E Count (AVE_PE):** Indicates the average P/E count index[1].
11:10	**Number of Dies per Super Block (NSBDIE)**: Indicates the number of dies to compose a Super Block.
12	**Super Block Bit Width (SBBW):** Indicates the number of bits of the Super Block ID field in the flash address. This is a 0's based value.
13	**ADU Offset Bit Width (AOBW):** Indicates the number of bits of the ADU Offset field in the flash address. This is a 0's based value.
15:14	**Recommended Patrol Cycle (RPC):** Indicates recommended patrol cycle in minutes. Each closed Super Block should be checked within the cycle. This number decreases as wear increases.
19:16	**Maximum Number of Open Super Blocks (MXOSB):** Indicates the maximum number of simultaneously open Super Blocks within this Virtual Device. A value of 0 indicates the number is not reported here and reported in a SEF Command Set Identify Controller data.
23:20	**Default Read FIFO ID (DRFID):** Indicates default Read FIFO ID.
27:24	**Number of Read FIFOs (NRF):** Indicates the number of Read FIFOs attached to this Virtual Device.
31:28	**Number of pSLC Blocks (NPBLK):** Indicates the number of regular blocks per die transformed to pSLC. This is not a 0's based value. FFFFFFFFh means all Super Blocks are configured in pSLC mode.
39:32	**Virtual Device Total Capacity (VDTCAP):** Indicates the total number of usable Planes included in this Virtual Device. This number decreases as wear occurs. (SBSTA.TCAP + PSBSTA.TCAP)
47:40	**Virtual Device Guaranteed Capacity (VDGCAP):** Indicates the total GCAP. (SBSTA.GCAP + PSBSTA.GCAP)
55:48	**Virtual Device In-use Capacity (VDUCAP):** Indicates the total number of Planes actually assigned (SBSTA.UCAP + PSBSTA.UCAP)
63:56	Reserved
127:64	**Super Block Statistics (SBSTA):** Statistics of general Super Blocks. Refer to Table 6.41.
191:128	**pSLC Super Block Statistics (PSBSTA):** Statistics of pSLC Super Blocks. Refer to Table 6.41.
193:192	**Number of QoS Domains assigned (ASS_QOSD):** Indicates the number of QoS Domains assigned.
195:194	**Number of Dies (NDIE):** Indicates the number of dies the Virtual Device contains. This is a 0's based value.
197:196	**1st Die: ID** The first die identifier. The numbers are listed in descending order.

©2023 Software-Enabled Flash Project. All Rights Reserved

...	
2n+195:2n+194	**n-th Die ID:** The last, n-th (=NDIE), die identifier.

NOTES:

1. P/E count index indicates the number assuming 200 is the maximum limit of P/E counts. If the index exceeds 255, 255 is reported.

Table 6.41: Virtual Device Information - Super Block Statistics Structure

Byte	Description
7:0	**Total Capacity (TCAP):** Indicates the total number of Planes of the type of the Super Block. This number decreases as wear goes.
15:8	**Total Guaranteed Capacity (TGCAP):** Indicates the total guaranteed number of Planes of the type of Super Blocks including any required rounding/overhead.
23:16	**In-use Capacity (UCAP):** Indicates the total number of Planes actually assigned to the type of the Super Blocks
27:24	**Number of Free Super Blocks (NFRSB):** Indicates the number of Super Blocks in Free state.
31:28	**Number of Closed Super Blocks (NCLSB):** Indicates the number of Super Blocks in Closed states.
35:32	**Number of Open Super Blocks by Erase (NEOSB):** Indicates the number of Open Super Blocks by Erase.
39:36	**Number of Open Super Blocks for NLW (NWOSB):** Indicates the number of Open Super Blocks for Nameless Write with Placement ID indication.
63:40	Reserved

6.4.3 Super Block List (Log Identifier D2h)

This log page returns a list of Super Blocks assigned to the QoS Domain specified in CDW1.NSID in the specified order in LODR field. CDW12.LPOL and CDW13.LPOU are ignored.

Table 6.42: Get Log Page - Command Dword 10 for Super Block List

Bits	Description
31:15	See Figure 197 of NVMe 2.0
14:12	**List Order (LODR):** Specifies the order the list is sorted.

©2023 Software-Enabled Flash Project. All Rights Reserved

	Value	**Description**
	0h	Wear Leveling: Indicates the order of closed Super Blocks to free for Wear Leveling.
	1h	Most Urgent to Refresh: Indicates the order of closed Super Blocks for garbage collection due to poor data retention quality.
	2h	Most Urgent to Check: Indicates the order of closed Super Blocks for inspection to assess the data retention quality.
	3h	Open and Closed: Indicates a list of open and closed Super Blocks assigned to the QoS Domain. Each state is shown in the SBS field.
	4h	Closed: Indicates a list of closed Super Blocks assigned to the QoS Domain in ascending Super Block ID order.
	5h-7h	Reserved
11:00	See Figure 197 of NVMe 2.0	

Table 6.43: Get Log Page - Super Block List Data

Qword	Description
0	**1st Super Block Descriptor**
1	**2nd Super Block Descriptor**
—	
n-1	**n-th Super Block Descriptor**

Table 6.44: Super Block List Data - Super Block Descriptor

Bytes	Description
3:0	**Super Block ID (SBID):** Indicates Super Block ID
4	**P/E Count Index (PECI):** Indicates the P/E count index assuming 200 is the maximum P/E counts. When the index exceeds 255, 255 is returned.
5	**Data Integrity Index (DII):** Indicates how reliable the Super Block is, when CDW10.LODR=0~2, 4. **Super Block State (SBS)**: Indicates the state of the Super Block when CDW10.LODR=3. Refer to DII and SBS of Super Block Information Data for the values.
7:6	Reserved

When SBID=FFFFFFFFh, it indicates the end of the list.

©2023 Software-Enabled Flash Project. All Rights Reserved

6.4.4 Super Block Information (Log Identifier D3h)

This log page returns information of the Super Block specified with CDW12.LPOL in the Virtual Device specified with the Endurance Group Identifier in CDW11. Thus, the log page offset is always treated as '0'. If the Endurance Group Identifier is FFFFh, it returns the information of a Super Block which raises an asynchronous state change event. If no Super Block raises a state change, the controller aborts the command.

Table 6.45: Get Log Page - Super Block Information

<table>
<tr><th>Bytes</th><th>Description</th></tr>
<tr><td>1:0</td><td>Virtual Device ID (VDID): Virtual Device ID</td></tr>
<tr><td>2</td><td>State Transition (ST): This field is valid when the Endurance Group Identifier is FFFFh.
<table>
<tr><th>Bits</th><th>Description</th></tr>
<tr><td>7-4</td><td>Reserved</td></tr>
<tr><td>3</td><td>From 'Open by Erase' to 'Closed'</td></tr>
<tr><td>2</td><td>From 'Open for Nameless Write' to 'Closed'</td></tr>
<tr><td>1</td><td>Reserved</td></tr>
<tr><td>0</td><td>Data Integrity becomes 'Unknown'</td></tr>
</table></td></tr>
<tr><td>3</td><td>Super Block State (SBS): Indicates the state of the Super Block. If the Super Block is in Closed state, either of bit2 or 3 is also set to indicate the way of its open.
<table>
<tr><th>Bit</th><th>Description</th></tr>
<tr><td>7</td><td>Offline</td></tr>
<tr><td>6</td><td>pSLC</td></tr>
<tr><td>5-4</td><td>Reserved</td></tr>
<tr><td>3</td><td>Open by Erase</td></tr>
<tr><td>2</td><td>Open for Nameless Write with Placement ID</td></tr>
<tr><td>1</td><td>Closed</td></tr>
<tr><td>0</td><td>Free</td></tr>
</table></td></tr>
<tr><td>7:4</td><td>Super Block ID (SBID): Indicates Super Block ID</td></tr>
<tr><td>9:8</td><td>Placement ID (PID): Indicates Placement ID if bit2 of SBS is set. Otherwise, FFFFh is set.</td></tr>
<tr><td>11:10</td><td>QoS Domain ID (QOSDID): Indicates QoS Domain ID the Super Block belongs to, if either bit of 1 to 3 in SBS is set. Otherwise, it is cleared to '0'.</td></tr>
<tr><td>12</td><td>P/E Count Index (PECI): Indicates the P/E count index assuming 200 is the maximum P/E counts. When the index exceeds 255, 255 is returned.</td></tr>
<tr><td>13</td><td>Data Integrity Index (DII): Indicates how reliable the Super Block in 'Closed' state is.</td></tr>
</table>

©2023 Software-Enabled Flash Project. All Rights Reserved

	Value / **Description** 0 / Unknown 1 / Good 2 / Allowable 3 / Marginal 4-15 / Reserved
15:14	**Time Left (TL):** Indicates minutes left. When DII=0, the host should complete checking (patrolling), when DII=2 or 3, the host should complete refreshing within the limit.
19:16	**Capacity (CAP):** Indicates the capacity in the number of ADUs. It may decrease during Open state. This field is effective when either of bit 1 to 3 in SBS is set.
23:20	**ADU Pointer (ADUPTR):** Indicates the ADU ID that should be written next. The value is the offset following the mode of Perfect, Packed or Fragmented. Is valid only when it is in Open state.
27:24	**Erase Serial Number (ESN):** Indicates the order of erasure. This field is valid for those in Open or Closed state.

Table 6.46: Get Log Page - Super Block Information

Bytes	Description
29:28	**Die Group ID (DGID):** Indicates the identifier of the die group which is a set of dies that the Super Block consists of. The value is between 0 and (VD.NDIE / VD.NSBDIE - 1).
31:30	**Number of Defective Planes (NDPL):** Indicates the number of defective Planes per Super Page. Is valid when its defective management method is Packed or Fragmented. It possibly increases during Open.
127:32	Reserved
255:128	**Defect Bitmap (DBM):** Indicates the bitmap which expresses defective Planes. Is valid when its defective management method is Packed or Fragmented. It possibly increases during Open.

6.4.5 User Address List (Log Identifier D4h)

This log page returns a list of user addresses in the Super Block specified with CDW13.LPOU used by the QoS Domain specified in CDW1.NSID in the order of Flash Address in order to reconstruct LUT after an unsafe shutdown. Only CDW12.LPOL is used for the offset.

The Super block must be in Open or Closed state. When an ADU fails to be read, FFFFFFFF_FFFFFFFFh for User Address is returned.

©2023 Software-Enabled Flash Project. All Rights Reserved

Table 6.47: Get Log Page - User Address Data List

Qwords	Description
0	**1st User Address**
1	**2nd User Address**
—	
m-1	**m-th User Address**

6.4.6 Address Change Order (Log Identifier D5h)

This log page is used to provide Address Change Order information of the ACO ID specified with CDW12.LPOL for the QoS Domain specified with CDW1.NSID. CDW13.LPOU is reserved.

An Address Change Order is generated by the device to notify the host of the Flash address relocation when a program error happens for a Nameless Write command.

After the command completes successfully, the information in the device is deleted, so that the host can no longer retrieve the information. If the number of Dwords is less than six times of CQE DW1.NADU of the Nameless Write command, the command is aborted.

If the ACO ID is invalid, the command is aborted.

Table 6.48: Get Log Page - Address Change Order Log Entry (Log Identifier D5h)

Bytes	Description
07:00	**Starting User Address (SUA):** The starting user address specified with CDW10.SUA in the associated Nameless Write command.
15:08	**Flash Address (FLA):** The flash address specified with CDW14-15 in the associated Nameless Write command.
17:16	**Number of ADUs (NADU):** The number of ADUs specified with CDW12.NADU in the associated Nameless Write command. This is a 0's based value.
19:18	**Number of ACO Entries (NACO):** The number of address change order entries generated. This is a 0's based value.
23:20	**NLW ID (NLWID):** Nameless Write ID from CDW13.NLWID in the Nameless Write command.
27:24	**Number of ADUs Left (NADUL):** The remaining number of ADUs in the last Super Block.
29:28	**Number of Defective Planes (NDP):** The number of defective planes per Super Page in the last Super Block.
31:30	Reserved
55:32	**1st Address Change Order Entry**

©2023 Software-Enabled Flash Project. All Rights Reserved

79:56	**2nd Address Change Order Entry**
—	
24n+31:24n+8	**n-th Address Change Order Entry**

Table 6.49: Address Change Order Entry

Qword	Description
0	**User Address**
1	**Old Flash Address:** This field is set to FFFFFFFF_FFFFFFFFh if the associated Nameless Write command was processed independently.
2	**New Flash Address**

6.4.7 Read FIFO List (Log Identifier D6h)

This log page returns a list of Read FIFOs of the Virtual Device specified with the Endurance Group Identifier (ENDGID) in CDW11 in ascending ID order.

If ENDGID is FFFFh, all Read FIFOs are listed. For the Read FIFO which is not attached to any Virtual Device, its Virtual Device ID field is cleared to '0'. For the Read FIFO which is attached to multiple Virtual Devices, multiple Read FIFO Collections are listed in ascending Virtual Device ID order.

Table 6.50: Get Log Page - Read FIFO List

Entry	Description
0	**1st Read FIFO Collection**
1	**2nd Read FIFO Collection**
—	
m-1	**n-th Read FIFO Collection**

Table 6.51: Get Log Page - Read FIFO Collection

Bytes	Description
3:0	**Read FIFO ID**
5:4	**Virtual Device ID**
7:6	**Weight**

When Read FIFO ID=FFFFFFFFh, it indicates the end of the list.

©2023 Software-Enabled Flash Project. All Rights Reserved

6.5 Identify command

The SEF Command Set supports following CNS values.

Table 6.52: Identify - CNS Values

CNS Value	O/M	Definition	NSID	CNTID	CSI	Reference Section
Active Namespace Management						
00h	M	Identify Namespace data structure.	Y	N	N	NVM1.0
01h	M	Identify Controller data structure.	N	N	N	5.5.1
02h	M	Active Namespace ID list	Y	N	N	NVMe 2.0
03h	M	Namespace Identification Descriptor list.	Y	N	N	NVMe 2.0
05h	M	Identify I/O Command Set specific Namespace data structure.	Y	N	Y	5.5.2
06h	M	Identify I/O Command Set Specific Controller data structure.	N	N	Y	5.5.3
07h	M	Active Namespace ID list associated with the specified I/O Command Set.	Y	N	Y	NVMe 2.0

©2023 Software-Enabled Flash Project. All Rights Reserved

08h	M	I/O Command Set Independent Identify Namespace data structure	Y	N	N	NVMe 2.0
Controller and Namespace Management						
10h	M	Allocated Namespace ID list.	Y	N	N	NVMe 2.0
11h	M	Identify Namespace data structure for the specified allocated NSID	Y	N	N	NVMe 2.0
12h	M	Controller List of controllers attached to the specified NSID.	Y	Y	N	NVMe 2.0
13h	M	Controller List of controllers that exists in the NVM subsystem.	N	Y	N	NVMe 2.0
19h	M	Endurance Group List	N	N	N	NVMe 2.0
1Ah	M	I/O Command Set specific Allocated Namespace ID list	Y	N	Y	NVMe 2.0
1Bh	M	I/O Command Set specific Identify Namespace data structure	Y	N	Y	NVMe 2.0
1Ch	M	I/O Command Set data structure	N	Y	N	NVMe 2.0

©2023 Software-Enabled Flash Project. All Rights Reserved

Command Set Identifier for SEF Command Set is tentatively 30h.

Table 6.53: Command Set Identifiers

Command Set Identifier Value	Description
00h	NVM Command Set
01h	Key Value Command Set
02h	Zoned Namespace Command Set
30h	SEF Command Set

6.5.1 Identify Controller data structure (CNS 01h)

Below only shows portions different from NVMe 2.0 and related to SEF.

Table 6.54: Identify - Identify Controller Data Structure, I/O Command Set Independent

Bytes	Description
99:96	**Controller Attributes (CTRATT):** This field indicates attributes of the controller.

©2023 Software-Enabled Flash Project. All Rights Reserved

Bits	Description
31:15	Same as NVMe 2.0
14	**Delete NVM Set:** If set to '1', then the controller supports the Delete NVM Set operation (refer to section 8.3.3 in NVMe 2.0). If cleared to '0', then the controller does not support the Delete NVM Set operation. SEF reports '0'.
13	**Delete Endurance Group:** If set to '1', then the controller supports the Delete Endurance Group operation (refer to section 8.3.3 in NVMe 2.0). If cleared to '0', then the controller does not support the Delete Endurance Group operation. SEF reports '0' in production, may be '1' in pre-production.
12	**Variable Capacity Management:** If set to '1', then the controller supports Variable Capacity Management (refer to section 8.3.3 in NVMe 2.0). If cleared to '0', then the controller does not support Variable Capacity Management. SEF reports '0'.
11	**Fixed Capacity Management:** If set to '1', then the controller supports Fixed Capacity Management (refer to section 8.3.2 in NVMe 2.0). If cleared to '0', then the controller does not support Fixed Capacity Management. SEF reports '1'.
10	**Multi-Domain Subsystem (MDS):** If set to '1', then the NVM subsystem supports multiple domains (refer to section 3.2.4 in NVMe 2.0). If cleared to '0', then the NVM subsystem does not support the reporting of multiple domains and the NVM subsystem consists of a single domain. SEF reports '0'.
9:0	Same as NVMe 2.0

6.5.2 Identify SEF Command Set specific Namespace data structure (CNS 05h CSI 30h)

A structure based on the NVM Command Set Identify Namespace data (CNS 00h) with additional fields is returned. The table below shows portions different from NVM 2.0.

Table 6.55: Identify SEF Command Set specific Identify Namespace Data Structure

Bytes	O/M[1]	Description
07:00	M	**Namespace Size (NSZE):** Indicates the maximum number of ADUs The value is (SBSTA.QUOTA + PSBSTA.QUOTA) in ADUs
15:08	M	**Namespace Capacity (NCAP):** Indicates the number of guaranteed ADUs The value is (SBSTA.GCAP + PSBSTA.GCAP) in ADUs

©2023 Software-Enabled Flash Project. All Rights Reserved

<table>
<tr><td>23:16</td><td>M</td><td>Namespace Utilization (NUSE): Indicates the number of ADUs the QoS Domain is currently using. The value is (SBSTA.UCAP + PSBSTA.UCAP) in ADUs</td></tr>
<tr><td>103:102</td><td>M</td><td>Endurance Group Identifier (ENDGID): This field indicates the Endurance Group (Virtual Device) with which this namespace is associated.</td></tr>
<tr><td>119:104</td><td>M</td><td>Namespace Globally Unique Identifier (NGUID)</td></tr>
<tr><td colspan="3">SEF specific fields</td></tr>
<tr><td>771:768</td><td>M</td><td>QoS Domain Settings:
<table>
<tr><th>Bits</th><th>Description</th></tr>
<tr><td>31:04</td><td>Reserved</td></tr>
<tr><td>03:02</td><td>Defect Management Type:
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>0</td><td>Perfect</td></tr>
<tr><td>1</td><td>Packed</td></tr>
<tr><td>2</td><td>Fragmented</td></tr>
<tr><td>3</td><td>Reserved</td></tr>
</table></td></tr>
<tr><td>01:00</td><td>Read Deadline:
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>0</td><td>Fastest</td></tr>
<tr><td>1</td><td>Typical</td></tr>
<tr><td>2</td><td>Long</td></tr>
<tr><td>3</td><td>Heroic</td></tr>
</table></td></tr>
</table></td></tr>
<tr><td>773:772</td><td>M</td><td>Number of Placement IDs (NPLID): Indicates the number of Placement IDs. This is a 0's based value.</td></tr>
<tr><td>775:774</td><td>M</td><td>Maximum Number of Open Super Blocks (MAXOSB): Indicates the number of Super Blocks this NS is allowed to open simultaneously.</td></tr>
<tr><td>777:776</td><td>M</td><td>Write Weight (WWT)</td></tr>
<tr><td>779:778</td><td>M</td><td>Erase Weight (EWT)</td></tr>
<tr><td>783:780</td><td>M</td><td>Default Read FIFO ID (DRFID): Indicates the default Read FIFO ID used when CDW13.RFID of a Physical Read command is invalid.</td></tr>
<tr><td>831:784</td><td></td><td>Reserved</td></tr>
<tr><td>895:832</td><td>M</td><td>Super Block Statistics (SBSTA): Statistics of general Super Blocks. Refer to Table 6.56</td></tr>
<tr><td>959:896</td><td>O[2]</td><td>pSLC Super Block Statistics (PSBSTA): Statistics of pSLC Super Blocks. Refer to Table 6.56</td></tr>
<tr><td>967:960</td><td>M</td><td>Root Pointer-0</td></tr>
<tr><td>—</td><td></td><td>---</td></tr>
</table>

©2023 Software-Enabled Flash Project. All Rights Reserved

1023:1016	M	**Root Pointer-7**
4095:1024		Reserved

NOTES:
1. O/M definition: O = Optional, M = Mandatory
2. Mandatory for controllers that support pSLC Super Block

Table 6.56: Identify SEF Command Set specific Namespace - Super Block Statistics Data Structure

Bytes	Description
07:00	**QUOTA (QUOTA):** Indicates the maximum number of Planes the QoS Domain may use.
15:08	**Guaranteed Capacity (GCAP):** Indicates the number of Planes the QoS Domain is guaranteed.
23:16	**Utilization (UCAP):** Indicates the number of Planes the QoS Domain is currently using.
27:24	**Number of Super Blocks (NSB):** Indicates the number of Super Blocks the QoS Domain is currently holding.
29:28	**Number of Open Super Blocks for Nameless Write (NWOSB):** Indicates the number of open Super Blocks for Nameless Write.
30	**Number of Open Super Blocks by Erase (NEOSB):** Indicates the number of open Super Blocks by Erase.
63:31	Reserved

6.5.3 Identify SEF Command Set specific Controller data structure (CNS 06h, CSI 30h)

The controller returns the information of the SEF Unit.

Table 6.57: Identify - Identify Controller Data Structure, SEF Command Set

Bytes	Description
Static Information	
03:00	**SEF Capability (SEFCAP):** Indicates capabilities the device supports.

©2023 Software-Enabled Flash Project. All Rights Reserved

Bits	Description
31:15	Reserved
14	If set to '1', then the device supports pSLC Super Block. If cleared to '0', then the device doesn't.
13	If set to '1', then the device supports Read FIFOs shared by multiple Virtual Devices. If cleared to '0', then the device doesn't.
12	If set to '1', then the device supports a Super Block operated with Nameless Write and Nameless Copy commands. If cleared to '0', then the device doesn't.
11	If set to '1', then the device supports Nameless Write and Nameless Copy commands explicitly specifying the ADU offset. If cleared to '0', then the device doesn't.
10	If set to '1' then then a paired Nameless Copy and Get Copy Results commands can be issued through different SQs. If cleared to '0', then they must be issued through the same SQ and back to back. The host may use CDW10.FUSE, while controller ignores it.
9	When bit 4 is cleared to '0', and if this bit is set to '1', then a paired Flash Address Request and Nameless Write commands can be issued through different SQs in any order. If cleared to '0', then they must be issued through the same SQ and back to back. The host may use CDW10.FUSE, while controller ignores it.
8	If set to '1', then the device supports SGL for Nameless Write command. If cleared to '0', then the device doesn't.
7	If set to '1', then the device supports SGL for Physical Read command. If cleared to '0', then the device doesn't.
6	If set to '1', then the device supports List indicated Nameless Copy command. If cleared to '0', then the device doesn't.
5	If set to '1', then the device supports User Address Range in Nameless Copy command. If cleared to '0', then the device doesn't.
4	If set to '1', then the device supports independent Nameless Write command. If cleared to '0', then the device doesn't.
3	If set to '1', then the device supports Perfect and other Packed or Fragmented mixed. If cleared to '0', then the device doesn't.
2	If set to '1', then the device supports Perfect ADU Offset. If cleared to '0', then the device doesn't.
1	If set to '1', then the device supports Packed ADU Offset. If cleared to '0', then the device doesn't.
0	If set to '1', then the device supports Fragmented ADU Offset. If cleared to '0', then the device doesn't.

©2023 Software-Enabled Flash Project. All Rights Reserved

04	**Number of Channels (NCH):**Indicates the number of channels in the device. This is a 0's based value.
05	**Number of Banks (NBNK):** Indicates the number of banks in the device. This is a 0's based value.
07:06	**Number of Blocks per Die (NBLK):** Indicates the number of blocks per die. This is a 0's based value.
09:08	**Number of Pages per Block (NPAG):** Indicates the number of pages per block. This is a 0's based value.
10	**Number of Planes per Page (NPL):** Indicates the number of Planes per page. This is a 0's based value.
11	**Plane Data Size (PLDS):** Indicates the Plane data size. The value is reported in terms of a power of two (2^n). A value smaller than 12 (i.e., 4096 bytes) is not supported.
13:12	**Expiration Open Period (EOP):** Indicates the time limit, in seconds, of a Super Block in the Open state (beginning from its being erased to being completely written). If the Super Block is not completely written and closed before this timeout, the device will transition it to the Closed state. The granularity of about 12 hours is expected.
15:14	**Number of Placement IDs (NPLID):** Indicates the maximum number of Placement IDs per QoS Domain. This is a 0's based value.
19:16	**Maximum Number of Open Super Blocks (MXOSB):** Indicates the maximum number of simultaneously open Super Blocks within the device. This is not a 0's based value. A value of 0 indicates the number is not reported and reported in a Virtual Device Information log page.
23:20	**Number of Read FIFOs (NRF):** Indicates the number of Read FIFOs within the device. This is not a 0's based value. The Read FIFO ID of 1 through NRF are valid.
24	**Number of Pages to Secure Integrity (NPSI):** Indicates the maximum number of additional pages per block required to secure integrity of a page. Full-sequence flash returns 0. This is NOT a 0's based value.
63:25	Reserved
Dynamic Information	
67:64	**Number of Read FIFOs in use (NURF):** Indicates the number of Read FIFOs attached. This is NOT a 0's based value.
69:68	**Number of Virtual Devices (NVD):** Indicates the number of present Virtual Devices. This is NOT a 0's based value.
71:70	**Number of QoS Domains (NQOSD):** Indicates the number of present QoS Domains. This is NOT a 0's based value.
4095:72	Reserved

©2023 Software-Enabled Flash Project. All Rights Reserved

6.6 Namespace Management command

The Namespace Management command is used to manage namespaces (refer to section 8.11 of NVMe 1.0), including create and delete operations. CDW11.CSI must be 30h (SEF Command Set). If other value is specified, the controller aborts the command. If SBSTA.GCAP or PSBSTA.GCAP doesn't satisfy those of other QoS Domains in the Virtual Device, the controller aborts the command. The data structure used for the create operation has the same format as the SEF Command Set Identify Namespace data structure. The fields that host software may specify in the create operation are defined in the following table.

Table 6.58: Namespace Management - Host Software Specified Fields

Bytes	Description
26	**Formatted LBA Size (FLBAS)**
103:102	**Endurance Group Identifier (ENDGID)**
	SEF Specific Fields
771:768	**QoS Domain Settings**
773:772	**Number of Placement IDs (NPLID)**
775:774	**Maximum Number of Open Super Blocks (MAXOSB):** Specifies the number of Super Blocks this Namespace opens simultaneously. This is not a 0's based value. If a value less than NPLID+2 is specified, then a value of NPLID+2 is adopted. With this value, if the total number exceeds the maximum number within the device or the Virtual Device, the controller aborts the command.
777:776	**Write Weight (WWT)**
779:778	**Erase Weight (EWT)**
783:780	**Default Read FIFO ID (DRFID):** If the specified FIFO is not attached to the Virtual Device, the controller aborts the command. If a value of '0' is specified, the default Read FIFO for the Virtual Device is used.
839:832	**SBSTA.QUOTA:** If SBSTA.QUOTA $<$ SBSTA.GCAP, SBSTA.GCAP is adopted.
847:840	**SBSTA.GCAP**
903:896	**PSBSTA.QUOTA:** If PSBSTA.QUOTA $<$ PSBSTA.GCAP, PSBSTA.GCAP is adopted.
911:904	**PSBSTA.GCAP**

©2023 Software-Enabled Flash Project. All Rights Reserved

7 I/O Commands for the Software Enabled Flash (SEF) Command Set

Table 7.1 shows a list of IO commands the Software-Enabled Flash (SEF) devices support.

Standard commands for NVM Command Set are not supported.

Table 7.1: Opcodes for IO Commands

Opcode	Direction	Command	Remarks
DAh	In	Super Block Management	
D1h	Out	Nameless Write	
D2h	In	Flash Address Request	
D6h	In	Physical Read	
DDh	Out	Nameless Copy	
DEh	In	Get Copy Results	
NOTES: Direction is from the host.			

7.1 Super Block Management command

This command erases (opens), closes, frees, patrols, or flushes a Super Block of the QoS Domain specified in CDW1.NSID. The maximum number of Super Blocks to open at the same time of a QoS Domain is one greater than the number of its Placement IDs. This command uses Command Dword 10, 11 and 15. When CDW11.NUMDU>0, the data of the Super Block Information Log page (Table 6.45) is returned.

Table 7.2: Super Block Management - Command Dword 10

Bits	Description

©2023 Software-Enabled Flash Project. All Rights Reserved

31:0	**Super Block ID (SBID):** Specifies the Super Block to process. When FFFFFFFFh is specified on erasure, the device picks most suitable Super Block.

Table 7.3: Super Block Management - Command Dword 11

Bits	Description
31:20	Reserved
19:16	**Operation (OP):** Specifies an operation to process.

Value	Description
0h, 8h	**Erase:** Erases a Super Block in Free or Close state and makes it Open for Nameless Write and/or Nameless Copy. If CDW10.SBID is FFFFFFFFh, a Super Block in Free state is always selected. If OP is 0h, a general Super Block is allocated. If OP is 8h, a pSLC Super Block is allocated.
1h	**Close:** Closes a Super Block in Open state. Even when this operation is attempted on a Super Block in Closed state, it returns normal end. In case a program error happens for the Super Block for a NLW command with Placement ID indication, its alternative Super Block is also closed. It completes after the target Super Block is closed.
2h	**Free:** Frees a Super Block in Open or Closed state. The command completes after all on-going Read and Nameless Copy commands on this Super Block complete.
3h	**Patrol:** Patrols a Super Block.
4h	**Flush**: Makes NLW/NLC commands which are suspended waiting successive commands return after their data become readable from the media by padding minimum amount of data. Even when this operation is attempted on a Super Block in Closed state, it returns normal end. In case a program error happens for the Super Block for a NLW command with Placement ID indication, its alternative Super Block is also flushed. It completes after the target Super Blocks are all flushed.
5h-7h 9h-Fh	Reserved

15:00	**Number of Dwords (NUMD):** Specifies the number of Dwords to return. This is not a 0's based value. A value of '0' indicates no data returned.

Table 7.4: Super Block Management - Command Dword 15

Bits	Description

©2023 Software-Enabled Flash Project. All Rights Reserved

31:16	**Overriding Weight (OWT):** Specifies Weight value of the operation. When CDW10.OP=0 (Erase), this value is used for the erase operation. When CDW10.OP=1 (Close) or 4(Flush), this value is used for the program operation of padding. In any case, this field is cleared to 0, corresponding weight specified when the QoS Domain was created is used.
15:00	Reserved

7.1.1 Command Completion

Table 7.5: Super Block Management - Completion Queue Entry Dword 0 for other than patrol

Bits	Description
31:00	**Super Block ID (SBID):** Indicates Super Block ID processed.

Table 7.6: Super Block Management - Completion Queue Entry Dword 0 for Patrol

Bits	Description
31:8	Reserved
7:0	**Data Integrity Index:** Indicates how reliable the Super Block is.

Value	Description
0	Unknown
1	Good
2	Allowable
3	Marginal
4-15	Reserved

Table 7.7: Super Block Management - Completion Queue Entry Dword 1

Bits	Description
31	**Caused Closure (CC):** This bit is effective when OP is either 1h (Close), 2h (Free), or 4h (Flush). If CC is set to '1', then the operation caused the Super Block to close and returns after the closure. If CC is cleared to '0', then the operation didn't cause the closure, and returns independently from the closure.
30:00	**Number of ADUs Left (NADUL):** This field indicates the number of ADUs left in the Super Block, and is effective for Erase and Flush.

©2023 Software-Enabled Flash Project. All Rights Reserved

7.2 Physical Read command

This command reads user data of ADUs of the QoS Domain Specified with CDW1.NSID.

Physical Read command uses Metadata Pointer, Data Pointer and Command Dword 10, 11, 12, 14, and 15.

Table 7.8: Physical Read - Metadata Pointer

Bits	Description
63:00	**Metadata Pointer (MPTR):** This field contains the Metadata Pointer, if applicable.

Table 7.9: Physical Read - Data Pointer

Bits	Description
127:00	**Data Pointer (DPTR):** This field specifies the location of a data buffer where data is transferred to.

Table 7.10: Physical Read - Command Dword 10 and Dword 11

Bits	Description
63:00	**Starting User Address (SUA):** Specifies starting User Address to read. If this field is not -1 (all 1), UA part in every ADU is checked with this field. The check can detect certain classes of hardware faults, so it is recommended to specify this field. If SUA is not -1 but the range includes -1, an error returns.

Table 7.11: Physical Read - Command Dword 12

Bits	Description
31	**Limited Retry (LR):** If cleared to '0', the controller follows the setting of the Read Deadline in the QoS Domain Settings in the Identify Namespace. If set to '1', the controller processes in Fastest mode.
30:16	Reserved
15:00	**Number of ADUs (NADU)**: Specifies the number of ADUs to read. This is a 0's based value.

Table 7.12: Physical Read - Command Dword 13

Bits	Description
31:00	**Read FIFO ID (RFID):** Specifies Read FIFO ID to use. If the value is invalid, the default Read FIFO of the QoS Domain is used.

©2023 Software-Enabled Flash Project. All Rights Reserved

Table 7.13: Physical Read - Command Dword 14 and Command Dword 15

Bits	Description
63:48	**Overriding Weight (OWT):** Specifies the Weight to override. When this field is cleared to 0, the weight of the Read FIFO is used.
47:00	**Starting Flash Address (SFLA):** Specifies lower 48bits of the starting Flash Address to read.

7.2.1 Command Completion

Table 7.14: Physical Read - Completion Queue Entry Dword 0

Bits	Description
31:8	Reserved
7:0	**Data Integrity Index:** Indicates how reliable the Super Block is.

Value	Description
0	Unknown
1	Good
2	Allowable
3	Marginal
4-15	Reserved

7.3 Flash Address Request command

This command is issued as a pair with a Nameless Write command. The command returns tentative destination Flash Addresses (ARR: Address Record Request) where the data of the Nameless Write command which is issued as a 2nd command in the operation will be written. The destination addresses are returned in the buffer pointed by DPTR. If the data will be written in a single Super Block, one pair is returned, when the data will be written in multiple Super Blocks, two or more pairs are returned. When the second or later Super Block cannot be allocated, or the destination locations span more seven Super Blocks, the complete status of Partial Write is returned with information of the secured region(s) in ARR.

The buffer size is fixed at 128 bytes.

Multiple commands can be issued in advance. When they specify Super Block indication (CDW12.SBI = 1), and the Super Block is closed, they are aborted. This command uses Data Pointer and Command Dword 12-15.

Table 7.15: Flash Address Request - Data Pointer

©2023 Software-Enabled Flash Project. All Rights Reserved

Bits	Description
127:00	**Data Pointer (DPTR):** This field specifies the location of a data buffer where ARR (Address Record Request) is transferred to.

Table 7.16: Flash Address Request - Command Dword 12

Bits	Description
31:26	Reserved
25	**Super Block Indication (SBI):** If cleared to '0', the destination Super Block is implicitly specified with CDW14. PLID. If set to '1', the destination Super Block is explicitly specified with CDW14.FLA.
24:16	Reserved

Table 7.17: Flash Address Request - Command Dword 13

Bits	Description
31:00	**NLWID:** Specifies the ID of the associated Nameless Write command. ARR for the Nameless Write command with the same CDW13.NLWID and CDW12.PLID or the Super Block ID in CDW12.FLA is returned.

Table 7.18: Flash Address Request - Command Dword 14 and 15 in case of CDW12.SBI='0'

Bits	Description
63:48	**Overriding Weight (OWT):** Specifies the Weight to override for Erase operations. When this field is cleared to 0, EWT specified when the QoS Domain was created is used.
47:32	**Placement ID (PLID):** Specifies the Placement ID
31:0	Reserved

Table 7.19: Flash Address Request - Command Dword 14 and 15 in case of CDW12.SBI='1'

Bits	Description
63:48	**Overriding Weight (OWT):** Specifies the Weight to override for Erase operations. When this field is cleared to 0, EWT specified when the QoS Domain was created is used.
47:00	**Flash Address (FLA):** Only the Super Block ID field is valid.

7.3.1 Command Completion

©2023 Software-Enabled Flash Project. All Rights Reserved

Table 7.20: Flash Address Request - Completion Queue Entry Dword 0

Bits	Description
31:00	**Number of FLAs (NFLA):** Indicates the number of Flash Addresses returned. The value is equal to the total number of FLAs in the ARR. This is NOT a 0's based value.

Table 7.21: Flash Address Request - Status Code

SCT	SC	Description
7	F0h	**Partial Write:** Indicates that all CDW12.NADU has not been processed. It implies at least one ADU has been processed.

Table 7.22: Flash Address Request - Address Record Request (ARR) Data

Bytes	Description
00	Number of Super Blocks to be written to
01	Number of bytes of the defective plane bitmap
03:02	Number of Planes per Super Page
07:04	Number of ADUs left in the last Super Block. The value of '0' indicates it will be closed when the associated Nameless Write normally processed.
15:08	Reserved
31:16	1st Super Block information
...	
127:112	7th Super Block information

Table 7.23: Flash Address Request - Super Block information

Bytes	Description
03:00	Number of FLAs to be written
05:04	Number of defective Planes per Super Page. This is a 1's based value. May be increased in the later ARRs/ACOs.
07:06	Reserved
15:08	Starting Flash Address

7.4 Nameless Write command

The Nameless Write command writes data and metadata, if applicable, in the QoS Domain specified in CDW1.NSID. The destination Super Block can be explicitly specified, or implicitly specified with a Placement ID. In case of specifying Super Block, the host needs to allocate a destination Super Block

©2023 Software-Enabled Flash Project. All Rights Reserved

with the Super Block Management command in advance, and the requested number of ADUs is greater than the Super Block, the status of partial write returns. In case of specifying Placement ID, the device automatically allocated Super Blocks as much as possible. The ADU offset in the destination Super Block can be explicitly specified, or let it assigned by the device. When explicitly specified, and a program error happens, the command is aborted in the middle. When the ADU offset is set all '1', the device relocates the data as much as possible and returns with ACO ID. This command can be issued with a Flash Address Request command. If a Flash Address Request command associates, the same NLWID needs to be set in the CDW13. This command returns once it completes processing even if its associated Flash Address Request command is not issued by then.

After the Flash Address Request command completes, the host is allowed to issue Read commands with the new addresses, and the SEF device returns data from the host memory buffer of this command. When all data becomes readable from Flash media, this command completes, returning the number of ADUs and ACO ID. When ACO ID = 0, it indicates all data is successfully written to the Flash Addresses the associated Flash Address Request command returned. Even when Flash Address Request command returned Partial Write, it returns normal as long as the data is successfully written to the Flash Addresses the associated Flash Address Request command returned.

Figure 7.1: Nameless Write - Typical Operation

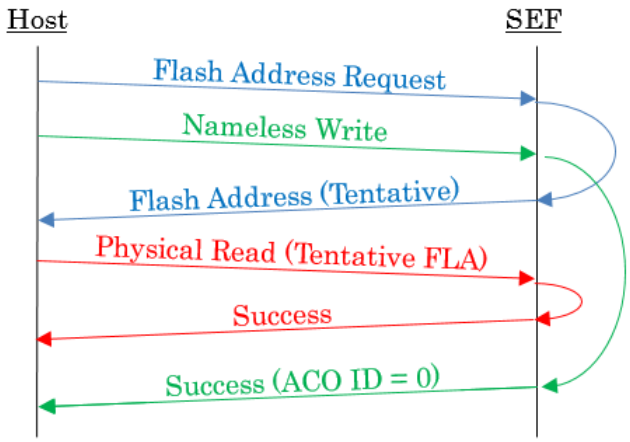

When ACO ID <> 0 with a Flash Address Request command associated, it indicates a program error happens and the number of ADUs relocated is returned. The host needs to get Address Change Order Log page for the relocation. The host should replace with the new FLA in its LUT only when the FLA in the LUT matches the old FLA.

Uncorrectable Read Error could return for a Read command to a Flash address notified on return of the Flash Address Request command in case this command returns at almost the same timing which requests the address to be changed. In this case, the host must retry the Read command with the new Flash address notified on return of this command.

©2023 Software-Enabled Flash Project. All Rights Reserved

Figure 7.2: Nameless Write - Race condition in Write and Read Operations

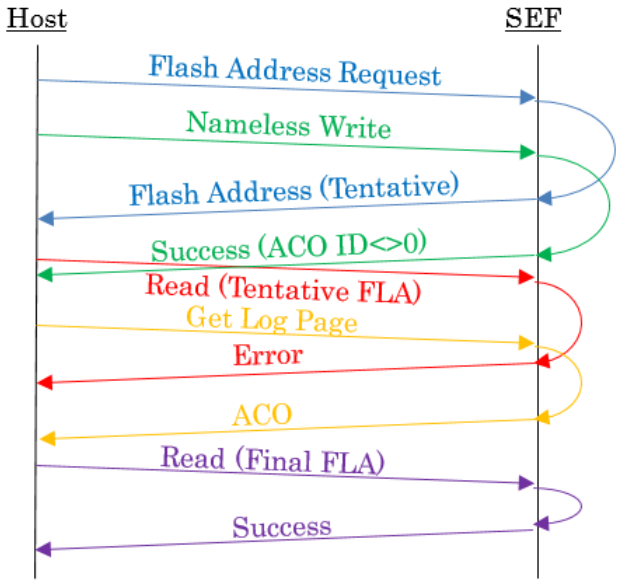

When this command is processed independently, ACO ID is always returned. The host needs to get destination flash addresses by reading Address Change Order Log page. This command uses Metadata and Data pointers and Command Dword 10 through 15.

Table 7.24: Nameless Write - Metadata Pointer

Bits	Description
63:00	**Metadata Pointer (MPTR):** This field contains the Metadata Pointer, if applicable.

Table 7.25: Nameless Write - Data Pointer

Bits	Description
127:00	**Data Pointer (DPTR):** This field specifies the location of a data buffer where data is transferred from.

Table 7.26: Nameless Write - Command Dword 10 and Command Dword 11

Bits	Description
63:00	**Starting User Address (SUA):** This field indicates the address of the first logical block to be written. When SUA is -1 (all 1), it stays at -1 throughout the region. If SUA is not -1 but the range includes -1, an error returns.

©2023 Software-Enabled Flash Project. All Rights Reserved

Table 7.27: Nameless Write - Command Dword 12

Bits	Description
31	**Limited Retry (LR):** If set to '1', the controller doesn't relocate the data when a program error happens, and just returns an error. If cleared to '0', the controller relocate the data when a program error happens, and notify the information through an ACO.
30	**Force Unit Access (FUA):** If set to '1', the controller completes the command without waiting for successive data.
29:26	Reserved
25	**Super Block Indication (SBI):** If cleared to '0', the destination Super Block is implicitly specified with CDW14.PLID. If set to '1', the destination Super Block is explicitly specified with CDW14.FLA.
24:00	Reserved
15:00	**Number of ADUs (NADU):** This field specifies the number of Flash Addresses to be secured. This is a 0's based value.

Table 7.28: Nameless Write - Command Dword 13

Bits	Description
31:00	**NLW ID:** Specifies the ID to associate a Flash Address Request command. If cleared to '0', it indicates the command is issued independently. The Flash Address Request command with the same CDW13.NLWID, and CDW14.PLID or the Super Block ID field in CDW14.FLA returns its ARR.

Table 7.29: Nameless Write - Command Dword 14 and 15 in case of CDW12.SBI='0'

Bits	Description
63:48	**Overriding Weight (OWT):** Specifies the Weight to override for Program operations. When this field is cleared to 0, WWT specified when the QoS Domain was created is used.
47:32	**Placement ID (PLID):** Specifies the Placement ID
31	**pSLC (PSLC):** If cleared to '0', a general Super Block is used. If set to '1', a pSLC Super Block is used. If an open Super Block already exists, and the mode is different, then the command is aborted.
30:0	**ADU Offset (ADUO):** Specifies the starting ADU offset in the destination Super Block. The bit width is specified AOBW in the Virtual Device Information Log page. If this field is set to all '1', the device assigns the location (Nameless). If this field is not set to all '1', the location is specified by the host (Named).

©2023 Software-Enabled Flash Project. All Rights Reserved

Table 7.30: Nameless Write - Command Dword 14 and 15 in case of CDW12.SBI='1'

Bits	Description
63:48	**Overriding Weight (OWT):** Specifies the Weight to override for Program operations. This field only affects Write Scheduler and doesn't affect Die Scheduler. When this field is cleared to 0, WWT specified when the QoS Domain was created is used.
47:00	**Flash Address (FLA):** Specifies the starting destination flash address. The bit widths of fields are specified SBBW and AOBW in the Virtual Device Information Log page. If ADU offset field is set to all '1', the device assigns the location (Nameless). If this field is not set to all '1', the location is specified by the host (Named).

7.4.1 Command Completion

The device returns a Status Code corresponding to whether to associate Flash Address Request and the number of Super Blocks the data span.

Table 7.31: Nameless Write - Status Code

SCT	SC	Description
0	00h	**Successful Completion:** The data has successfully been written with a Flash Address Request command associated. CQE is Type-1.
7	80h	**Successful Completion:** The command has been processed independently and the data has successfully been written in just one Super Block. CQE is Type-2.
7	81h	**Successful Completion:** The command has been processed independently and the data has successfully been written in two or more Super Blocks. CQE is Type-1.
others(TBD)		**Error:** Some kind of error has occurred. CQE is TBD.

Table 7.32: Nameless Write - Completion Queue Entry Dword 0 and 1 (Type-1)

Bits	Description
63:48	Reserved
47:32	**ACO ID (ACOID):** If the value is 0, it expresses that the data has successfully been written to the flash addresses notified with the associated Flash Address Request command. If the value is other than 0, the command has been processed independently and the data has been written to multiple Super Blocks, or the data has failed to be written to the flash addresses notified with the associated Flash Address Request command. The host can retrieve the relocated flash addresses in the ACO Log page with this ID as a key.
31:17	Reserved

©2023 Software-Enabled Flash Project. All Rights Reserved

16:00	**Number of ADUs (NADU):** When ACOID=0, it expresses the number of ADUs written. When ACOID<>0, it expresses the number of ADUs written to the different flash addresses from ones notified with the associated Flash Address Request command. This number is used to determine the size of the buffer for getting the ACO Log page.

Table 7.33: Nameless Write - Completion Queue Entry Dword 0 and 1 (Type-2)

Bits	Description
63:48	Indicates the number of defective planes per Super Page.
47:00	**Starting Flash Address (SFLA):** This field indicates the starting flash address. The host needs to translate to physical flash addresses by reading the Super Block Information log page if defective planes exist.

7.5 Nameless Copy command

This command is issued as a pair with a Get Copy Results command following. This command copies data from locations specified with Source Information to the Super Block specified with CDW14.FLA in the QoS Domain specified with CDW1.NSID. This command completes once the number of ADUs to copy is determined, while the associated Get Copy Results command completes after all destination data becomes readable from the flash media. Then the host needs to issue pairs of the commands until the destination Super Block is filled, or Super Block Management command to flush or close. After a certain period passes without an additional command, the controller completes the pending commands after padding minimum dummy data.

The number of ADUs to copy is determined when any of the following conditions is met.

- The command consumes the whole bitmap or FLA list
- FLA of the bitmap reaches the end of the source Super Block
- The destination Super Block is filled
- Reaches CDW10.NACRE of the associate Get Copy Results command

The destination Super Block needs to be allocated by 'Erase' with the Super Block Management command.

Parameters of the copy are specified in Command Dwords and the buffer pointed by DPTR.

Source ADUs are specified with Valid Bitmap or FLA list, and a User Address Region. CDW10.LIST specifies the indication to use, Bitmap or List. User Address Region is specified with starting User Address and the number of ADUs in the region. When the number of ADUs is '0', the region is disabled

©2023 Software-Enabled Flash Project. All Rights Reserved

and the whole source ADUs are copied. When the region is enabled, only ADUs, UA of which in the ADU in the media is inside or outside the region, are copied. CDW10.UAR specifies inside or outside.

When the destination Super Block transitions to Closed, Nameless Copy / Get Copy Results commands already processed complete normally, and successive commands not yet processed complete abnormally. To continue copy, the host needs to be allocate a new destination Super Block and reissue the commands. It is recommended for the host to update the affected elements of the LUT all at once when the destination Super Block is filled.

The controller doesn't start operation until the host fetches the associated Get Copy Results command to prevent race conditions.

This command uses DPTR and Command Dword 10, 12-15.

Table 7.34: Nameless Copy - Data Pointer

Bits	Description
127:00	**Data Pointer (DPTR):** This field specifies the location of a data buffer where copy information is transferred.

The buffer pointed by DPTR consists of two areas.

Table 7.35: Nameless Copy - Data Buffer

Bytes	Description
15:00	**User Address Range (UAR):** User Address Range
a+15:16	**Source Information (SINFO):** Bitmap or FLA List.
Note In case of Bitmap indication: a = (CDW10.NELM + 1) x 4 + 8 In case of FLA List indication: a = (CDW10.NELM + 1) x 8	

Table 7.36: Nameless Copy - User Address Range

Bytes	Description
07:00	**Starting User Address (SUA):** Specifies the staring User Address of the region.
15:08	**Length of Range (LR):** Specifies the number of ADUs of the region. '0' indicates the region is disabled.

©2023 Software-Enabled Flash Project. All Rights Reserved

Table 7.37: Nameless Copy - Source Information for Bitmap Indication

Bytes	Description
07:00	**Starting Flash Address of Bitmap (SFLA):** Specifies the starting Flash Address of the Bitmap. The Super Block must be in 'Closed' state. The address masked lower 5 bits expresses the Flash Address of the top of the Bitmap. The lower 5 bits expresses the starting bit position from the top of the Bitmap.
a+7:8	**Bitmap (BM):** Valid Bitmap
Note a = (CDW10.NELM + 1) x 4	

Table 7.38: Nameless Copy - Source Information for List Indication

Bytes	Description
7:0	**Source Flash Address-0 (SFLA0):** Flash Address 0
–	
a+7:a	**Source Flash Address-n (SFLAn):** Flash Address n
Note • The Super Block must be in 'Closed' state. • a = (CDW10.NELM + 1) x 8	

Table 7.39: Nameless Copy - Command Dword 10

Bits	Description
31:18	Reserved
17	**UA Range (UAR):** Specifies the indication of the User Address range. '0' targets inside the range, and 1' targets outside the range.
16	**FLA List or Bitmap (LIST):** Specifies the indication of the source ADUs. '0' specifies Bitmap indication, and '1' specifies the FLA List indication.
15:00	**Number of Elements (NUMEL):** Specifies the number of Dwords of the Bitmap, or the number of ADUs in the list. This is a 0's based value.

Table 7.40: Nameless Copy - Command Dword 12

Bits	Description
31	**Limited Retry (LR):** If set to '1', the controller stops operation when a program error happens. If cleared to '0', the controller keep processing if possible.

©2023 Software-Enabled Flash Project. All Rights Reserved

30:00	Reserved

Table 7.41: Nameless Copy - Command Dword 13

Bits	Description
31:00	**NLC ID:** Specifies the ID to associate a Get Copy Results command. The Get Copy Results command with the same CDW13.NLCID from the same QoS Domain returns the ACR.

Table 7.42: Nameless Copy - Command Dword 14 and 15

Bits	Description
63:48	**Overriding Weight (OWT):** Specifies the Weight to override for Write operations. When this field is cleared to 0, WWT specified when the QoS Domain was created is used.
47:00	**Flash Address (FLA):** Specifies lower 48 bits of the starting destination flash address. The bit widths of fields are specified SBBW and AOBW in the Virtual Device Information Log page. If ADU offset field is set to all '1', the device assigns the location (Nameless). If this field is not set to all '1', the location is specified by the host (Named).

7.5.1 Command Completion

Table 7.43: Nameless Copy - Completion Queue Entry Dword 0

Bits	Description
31:24	**Copy Status (CSTS):** Indicates the interim status of the Copy.
23:00	**Number of ADUs Left (NADUL):** Indicates the estimated number of ADUs left in the destination when the copy completes.

Bit	Description
7	Reserved
6	Detected a non-closed source Super Block in the list
4:3	Reserved
2	Consuming whole ACR entries.
1	Closing the destination Super Block.
0	Consuming whole source Bitmap or FLA List.

©2023 Software-Enabled Flash Project. All Rights Reserved

7.6 Get Copy Results command

This command is issued as a pair with a Nameless Copy command as its first. This command completes after all destination data of the associated Nameless Copy command become readable.

Get Copy Results command uses DPTR, Command Dword 10 and 13. The buffer size must be (CDW10.NACRE + 2) * 24 in bytes.

Table 7.44: Get Copy Results- Data Pointer

Bits	Description
127:00	**Data Pointer (DPTR):** This field specifies the location of a data buffer where ACR (Address Change Request) is transferred to.

Table 7.45: Get Copy Results - Address Change Request Data

Bytes	Description
03:00	**Number of Processed ADUs in Source (NPADU):** This field indicates the actual number of ADUs in source processed, including read errors, excluding out of range.
07:04	**Next Source Pointer (NSPTR)**: For bitmap indication, this field indicates the next ADU offset of the source flash address. For list indication, this field indicates the next entry number in the source flash address list.
11:08	**Number of Read Error ADUs (NRERR):** This field indicates the number of ADUs that couldn't be read due to a read error.
15:12	**Number of ADUs Left (NADUL):** This field indicates the actual number of ADUs left in the destination Super Block. The number can be less than that reported with the Nameless Copy command.
16	**Copy Status (CSTS):** Bit 7: Reserved Bit 6: Detected non closed source Super Block for list indication. Bit 5: The destination Super Block has defective planes. Bit 4: Read error happened on the source. Bit 3: ADUs eliminated with the UA region exist. Bit 2: Consumed whole ACR entries. Bit 1: Closed the destination Super Block. Bit 0: Consumed whole source Bitmap or FLA List.
23:17	Reserved
47:24	**1st Address Change Request Entry**
—	—

©2023 Software-Enabled Flash Project. All Rights Reserved

24*n+23:24*n	**n-th Address Change Request Entry** (n = ACR.NPADU)

Table 7.46: Get Copy Results - Address Change Request Entry

Qword	Description
0	**User Address**
1	**Old Flash Address**
2	**New Flash Address**
Note User Address has a special meaning at the following values. FFFFFFFF_FFFFFFFFh: Indicates its source FLA couldn't be read due to a read error.	

Table 7.47: Get Copy Results - Command Dword 10

Bits	Description
31:00	**Number of Address Change Request Entries (NACRE):** Specifies the number of Address Change Request entries in the ACR. This is a 0's based value.

Table 7.48: Get Copy Results - Command Dword 13

Bits	Description
31:00	**NLC ID:** Specifies the ID to associate a Nameless Copy command. The ACR for the Nameless Copy command with the same CDW13.NLCID is returned.

©2023 Software-Enabled Flash Project. All Rights Reserved