JEDEC STANDARD

Zoned Storage for UFS

JESD220-5

NOVEMBER 2023

JEDEC SOLID STATE TECHNOLOGY ASSOCIATION

NOTICE

JEDEC standards and publications contain material that has been prepared, reviewed, and approved through the JEDEC Board of Directors level and subsequently reviewed and approved by the JEDEC legal counsel.

JEDEC standards and publications are designed to serve the public interest through eliminating misunderstandings between manufacturers and purchasers, facilitating interchangeability and improvement of products, and assisting the purchaser in selecting and obtaining with minimum delay the proper product for use by those other than JEDEC members, whether the standard is to be used either domestically or internationally.

JEDEC standards and publications are adopted without regard to whether or not their adoption may involve patents or articles, materials, or processes. By such action JEDEC does not assume any liability to any patent owner, nor does it assume any obligation whatever to parties adopting the JEDEC standards or publications.

The information included in JEDEC standards and publications represents a sound approach to product specification and application, principally from the solid state device manufacturer viewpoint. Within the JEDEC organization there are procedures whereby a JEDEC standard or publication may be further processed and ultimately become an ANSI standard.

No claims to be in conformance with this standard may be made unless all requirements stated in the standard are met.

Inquiries, comments, and suggestions relative to the content of this JEDEC standard or publication should be addressed to JEDEC at the address below, or refer to www.jedec.org under Standards and Documents for alternative contact information.

Published by
©JEDEC Solid State Technology Association 2023
3103 North 10th Street
Suite 240 South
Arlington, VA 22201-2107

PRICE: Contact JEDEC

JEDEC Solid State Technology Association
3103 North 10th Street
Suite 240 South
Arlington, VA 22201-2107
https://www.jedec.org/contact

This page intentionally left blank.

ZONED STORAGE for UFS

CONTENTS

Page

This page intentionally left blank.

ZONED STORAGE for UFS

(From JEDEC Board Ballot JCB-20-35, formulated under the cognizance of the JC-13 Committee on Government Liaison.)

1 Scope

Universal Flash Storage (UFS) is a flash storage specification for digital electronic devices. The purpose of the Zoned Storage for UFS standard is to enable higher bandwidth, lower latency and to reduce write amplification. These objectives are realized as follows:

- Zoned storage enables reducing the size of the L2P table for zoned logical units. If the L2P table does not fit in UFS device SRAM without using the zoned storage interface then the switch to zoned storage may make it possible to eliminate L2P paging and hence increases bandwidth and reduces latency.
- It is recommended that vendors make the zone size close to the size of the erase block size. "Close" means that the zone size is identical to the erase block size, a small integer multiple of the erase block size or that there are two or more zones per erase block. Device-side garbage collection and write amplification due to garbage collection are eliminated if the zone size is equal to or a multiple of the erase block size since garbage collection is moved from the UFS device to the host. If there are two or more zones per erase block then the UFS device still has to perform garbage collection.
- Host software can optimize read performance by allocating a contiguous LBA range per file.

2 Referenced Standards

This standard references several other standard specifications by JEDEC and INCITS T10, including:

- [UFS 3.1] JEDEC, Universal Flash Storage (UFS), version 3.1, January 2020.
- [UFS 4.0] JEDEC, Universal Flash Storage (UFS), version 4.0, June 2022.
- [ZBC-2] INCITS, SCSI Zoned Block Commands (ZBC-2), INCITS 550-2023, 2023.

This standard builds on the UFS and the ZBC-2 standards. Zoned UFS (ZUFS) devices shall follow both standards. This means that the SCSI command and VPD page definitions from the UFS standard are combined with the SCSI command and VPD page definitions from the ZBC-2 standard. Additionally, the RC BASIS field shall be implemented in the READ CAPACITY (16) response since this is required by ZBC-2.

Since the SCSI command and VPD page definitions from the UFS standard are compliant with SPC-6 and SBC-5, the references to SPC-6 and SBC-5 in ZBC-2 do not affect any SCSI command or VPD page defined in the UFS standard.

3 Zoned Block Device Model

The data on a block device is organized in logical blocks. A logical block is a contiguous range of bytes that is accessed and referenced as a unit. Each logical block has an integer index called the LBA. The logical blocks of a zoned block device are divided into zones. Each zone contains a contiguous range of logical blocks. Each logical block is included in exactly one zone.

Each zone has the following properties:

- A type. This standard only supports the sequential write required zone type.
- A zone condition: empty, implicit open, closed or full.
- A write pointer.

All zones across all zoned logical units shall have the same length. See also the qZoneSize parameter in the geometry descriptor.

A zoned UFS logical unit reports PERIPHERAL DEVICE TYPE=14h in the INQUIRY response. All zones shall have the sequential write required zone type.

More information about zone types, zone conditions and the write pointer concept is available in ZBC-2. More information about the PERIPHERAL DEVICE TYPE field and the INQUIRY command are available in the UFS standard.

3.1 Open Zones

UFS devices shall support at least 6 zones that are either open or closed. If host software opens a zone if this limit is met then this may result in reduced performance.

3.2 READ Behavior

Zoned UFS devices shall report URSWRZ=1. For more information about URSWRZ, see also ZBC-2.

3.3 Logical Units

In addition to conventional and well known logical units, UFS devices shall support zoned logical units. Zoned logical units are logical units that consist of multiple zones. Each zone is a consecutive range of LBAs and has the sequential write required zone type. All zones shall have the same size across all zoned logical units.

Writing data to all zones shall be supported.

For all logical unit types, the UFS device is responsible for wear leveling, bad block replacement and retrying reads if necessary.

A UFS device is not required to perform garbage collection for zoned logical units except if there are multiple zones per erase block.

3.4 Logical Unit Configuration

The size of a zoned logical unit shall be a multiple of the zone size. A UFS device shall reject attempts to configure a zoned logical unit with a size that is not a multiple of the zone size.

3.5 UFS Application (UAP) Layer – SCSI Commands

Except if indicated otherwise, the behavior of UFS commands defined in the UFS standard is not modified. An exception is the READ CAPACITY command. ZBC-2 defines an RC BASIS field in the READ CAPACITY (16) response that is not present in any version of the UFS specification.

3.6 Inquiry Command

UFS devices shall report one of the following values in the PERIPHERAL DEVICE TYPE field of the INQUIRY response:

- 00h for conventional logical units.
- 1eh for well-known logical units.
- 14h or host-managed zoned block device for zoned logical units.

The values reported in other fields of the INQUIRY response follow the UFS standard.

3.7 Vital Product Data Parameters

A UFS device shall support the following VPD pages for zoned logical units:

- 00h - Supported VPD pages.
- 87h - Mode Page Policy.
- b0h - Block Limits.
- b1h - Block Device Characteristics.
- b2h - Logical Block Provisioning.
- b6h - Zoned Block Device Characteristics for zoned logical units only.

More information about VPD pages 00h, 87h, b0h, b1h and b2h is available in the UFS standard. VPD page b6h is defined in ZBC-2.

3.8 WriteBooster

The WriteBooster functionality shall comply with the UFS standard.

The UFS specification requires that flushing the WriteBooster buffer only happens if the command queue is empty or in the hibernated state. If the WriteBooster buffer becomes full while writing into a zone, a significant amount of data may have to be flushed before writing into the zone can continue. This can cause a latency spike. The initiator is responsible for avoiding these possible latency spikes.

3.9 Device Descriptor

In addition to the bits defined in the UFS standard for dExtendedUFSFeaturesSupport, UFS devices shall report in bit 11 whether or not provisioning type 10h is supported (zoned logical unit).

Offset	Size In bytes	Name	MDV	User Conf.	Description
4Fh	4	dExtendedUFSFeaturesSupport	Device specific	No	Extended UFS Features Support. This field indicates which features are supported by the device. This field value will be exactly the same value and same functionality as defined in the bit[0~7] of bUFSFeaturesSupport device descriptor. Since bUFSFeaturesSupport will be obsoleted, it is recommended to refer this descriptor to find out device feature support. A feature is supported if the related bit is set to one. bit[0]: Field Firmware Update (FFU) bit[1]: Production State Awareness (PSA) bit[2]: Device Life Span bit[3]: Refresh Operation bit[4]: TOO_HIGH_TEMPERATURE bit[5]: TOO_LOW_TEMPERATURE bit[6]: Extended Temperature bit[7]: Reserved for Host-aware Performance Booster (HPB) Extension Specification bit[8]: WriteBooster bit[9]: Performance Throttling bit[10]: Advanced RPMB bit[11]: Zoned logical units bit[12~13] : Reserved bit[14] : Barrier bit[15] : Clear Error History functionality bit[16] : EXT_IID bit[17] : Reserved for File Based Optimization (FBO) Extension Specification bit[18~31] : Reserved

3.10 Geometry Descriptor

In addition to the rows defined in the UFS standard, UFS devices shall report the following rows:

Offset	Size In bytes	Name	User Conf.	Description
57h	1	bMaxZonedLUCount	No	Maximum number of zoned logical units.
58h	8	qZoneAllocationUnit	No	The smallest number of bytes that can be erased by an erase operation from a zoned logical unit.
60h	8	qZoneSize	Depends on the device	Size of a zone in bytes. Additionally, the zone size shall either be a divisor or a multiple of qZoneAllocationUnit.

3.11 Unit Descriptor

Offset	Size in bytes	Name	Description
17h	1	bProvisioningType	**Provisioning Type** 00h: Thin Provisioning is disabled (default). 02h: Thin Provisioning is enabled and TPRZ = 0. 03h: Thin Provisioning is enabled and TPRZ = 1. 10h: Thin Provisioning is disabled, TPRZ=0 and zoned logical unit support is enabled. Others: Reserved.

Annex A (informative) Zoned Storage and Log-structured File Systems

A log-structured file system is a file system that writes all modifications to disk sequentially in a log-like structure. New data is appended at the end of the log instead of overwriting data in place. Space occupied by logical blocks that hold data that is no longer current is reclaimed during a process called garbage collection.

Log-structured file systems are well suited for zoned storage. If data is written into multiple files simultaneously, the data from these files will be interleaved in the log. In other words, data from multiple files may end up in a single zone. During the garbage collection process the layout of files is optimized. Files that occupy dis-contiguous LBA ranges are transferred to a contiguous LBA range. It is the responsibility of the file system to perform garbage collection in such a way that it does not cause latency spikes for host software.

See also Rosenblum, Mendel, and John K. Ousterhout. "The design and implementation of a log-structured file system." ACM Transactions on Computer Systems (TOCS) 10, no. 1 (1992): 26-52.

Standard Improvement Form JEDEC JESD220-5

The purpose of this form is to provide the Technical Committees of JEDEC with input from the industry regarding usage of the subject standard. Individuals or companies are invited to submit comments to JEDEC. All comments will be collected and dispersed to the appropriate committee(s).

If you can provide input, please complete this form and return to:

JEDEC
Attn: Publications Department
3103 North 10th Street. Suite 240S
Arlington, VA 22201

Fax: 703.907.7583

1. I recommend changes to the following:

☐ Requirement, clause number _____

☐ Test method number _____ Clause number _____

The referenced clause number has proven to be:

☐ Unclear ☐ Too Rigid ☐ In Error

☐ Other _____

2. Recommendations for correction:

_____

_____

_____

_____

3. Other suggestions for document improvement:

_____

_____

_____

_____

Submitted by

Name: _____ Phone: _____

Company: _____ E-mail: _____

Address: _____

City/State/Zip: _____ Date: _____


JEDEC®