

JEDEC STANDARD

Solid-State Drive (SSD) Requirements and Endurance Test Method

JESD218B

(Revision of JESD218A, February 2011)

MARCH 2016

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 2016
3103 North 10th Street
Suite 240 South
Arlington, VA 22201-2107

This document may be downloaded free of charge; however JEDEC retains the copyright on this material. By downloading this file the individual agrees not to charge for or resell the resulting material.

PRICE: Contact JEDEC

Printed in the U.S.A.
All rights reserved

PLEASE!

**DON'T VIOLATE
THE
LAW!**

This document is copyrighted by JEDEC and may not be reproduced without permission.

For information, contact:

**JEDEC Solid State Technology Association
3103 North 10th Street
Suite 240 South
Arlington, VA 22201-2107**

or refer to www.jedec.org under Standards-Documents/Copyright Information.

SOLID STATE DRIVE (SSD) REQUIREMENTS AND ENDURANCE TEST METHOD**Contents**

1	Scope.....	1
2	Reference Documents	1
3	Terms, definitions, abbreviations, and technical background.....	2
3.2	Data Error	2
3.3	Data Retention	2
3.4	Endurance	2
3.5	Endurance failure.....	2
3.6	Endurance Rating (TBW rating).....	2
3.7	Erase block.....	3
3.8	Failure	3
3.9	Functional Failure	3
3.10	Functional Failure Requirement (FFR).....	3
3.11	Gigabyte (GB).....	3
3.12	Host	3
3.13	Host writes	3
3.14	Logical Block Address (LBA)	3
3.15	Non-Volatile Memory (NVM).....	3
3.16	Page	3
3.17	Program/erase cycle (p/e cycle).....	3
3.18	Qualification Family	4
3.19	Retention failure	4
3.20	Solid State Drive (SSD).....	4
3.21	Terabyte (TB)	4
3.22	Uncorrectable Bit Error Rate, or ratio (UBER)	4
3.23	Wear leveling.....	5
3.24	Workload	5
3.25	Write amplification factor (WAF)	5
4	General SSD Ecosystem Factors	6
4.1	Capacity.....	6
4.2	Form Factors	6
5	Application classes and SSD Endurance Rating.....	6
5.1	Purpose of application classes	6
5.2	Definition of Application Classes.....	6
5.3	SSD Endurance Rating	7
5.4	Estimation of the SSD endurance rating.....	7
6	Endurance Test Method	8
6.1	Direct method	8
6.1.1	Sample Size and Acceptance Criteria	9
6.1.2	Categorization of failures.....	11
6.1.3	Endurance Stress Phase.....	11
6.1.4	High Temperature Retention Stress	14
6.1.5	Room temperature retention evaluation.....	15
6.2	Extrapolation method.....	16
6.2.1	Short Stroke Extrapolation Method	16
6.3	Endurance verification given alternative requirements.....	17
7	Apparatus and Precautions.....	17

7.1	Apparatus	17
7.2	Precautions.....	17
8	Summary	18
Annex A (normative): Calculations of Temperature-Accelerated Stress Times.....		19
Annex B (normative): Assessment of Room-Temperature Data Retention.....		21
Annex C (normative): SSD Capacity		23
C.1	SSD Capacity General Overview.....	23
C.2	LBA Counts from Advertised Capacity for Low Capacity Drives	24
C.2.1	LBA Counts from Advertised Capacity for Low Capacity Drives Overview	24
C.2.2	LBA Counts from Advertised Capacity for Low Capacity Drives with logical block size of 512 bytes	24
C.2.3	LBA Counts from Advertised Capacity for Low Capacity Drives with logical block size of 4096 bytes	24
C.3	LBA Counts from Advertised Capacity for High Capacity Drives	24
C.3.1	LBA Counts from Advertised Capacity for High Capacity Drives Overview.....	24
C.3	LBA Counts from Advertised Capacity for High Capacity Drives (cont'd).....	25
C.3.1	LBA Counts from Advertised Capacity for High Capacity Drives Overview (cont'd).....	25
C.3.2	LBA Counts from Advertised Capacity for 5xx High Capacity Drives	25
C.3.3	LBA Counts from Advertised Capacity for 4xxx High Capacity Drives	25
Annex D (informative): Estimated Retention Time for Nonstandard Active Use temperatures		26
Annex E (informative): Estimation of SSD Endurance Rating		27
Annex F (informative) - Differences between JESD218B and JESD218A.....		28

SOLID STATE DRIVE (SSD) REQUIREMENTS AND ENDURANCE TEST METHOD

(From JEDEC Board Ballot JCB-16-09, formulated under the cognizance of the JC-64.8 Subcommittee on Solid State Drives (Items 303.19, 303.20, 303.21, 303.22, 303.23, 303.26, 303.27, 303.28, and 303.32).)

1 Scope

This standard defines JEDEC requirements for solid state drives. For each defined class of solid state drive, the standard defines the conditions of use and the corresponding endurance verification requirements. Although endurance is to be rated based upon the standard conditions of use for the class, the standard also sets out requirements for possible additional use conditions as agreed to between manufacturer and purchaser.

Qualification of a solid state drive involves many factors beyond endurance and retention, so such qualification is beyond the scope of this standard, but this standard is sufficient for the endurance and retention part of a drive qualification. This standard applies to individual products and also to qualification families as defined in this standard.

The scope of this standard includes solid state drives based on solid-state non-volatile memory (NVM). NAND Flash memory is the most common form of memory used in solid state drives at the time of this writing, and this standard emphasizes certain features of NAND. The standard is also intended to apply to other forms of NVM.

2 Reference Documents

The revision of the referenced documents shall be that which is in effect on the date of the qualification plan.

JESD22-A117, *Electrically Erasable Programmable ROM (EEPROM) Program/Erase Endurance and Data Retention Stress Test*

JESD47, *Stress-Test-Driven Qualification of Integrated Circuits*

JEP122, *Failure Mechanisms and Models for Semiconductor Devices*

JESD219, *Solid State Devices (SSD) Endurance Workloads*

SFF-8447, *LBA Count for Disk Drives*

3 Terms, definitions, abbreviations, and technical background

3.1 Cycling pool

Erase blocks used by the SSD during read, write, or erase operations at a specific point in time.

NOTE The SSD may have additional erase blocks besides those in the current cycling pool that may be used as spares or for other purposes. The cycling pool is typically larger than the user-accessible LBA count.

3.2 Data Error

A type of failure in which the drive fails to return correct data to the host.

NOTE One data error occurs if a read of a logical sector causes the drive to return an unrecoverable error message or to return incorrect data. Data errors are counted as such even if they are transient. See 6.1.2 for further discussion of data errors.

3.3 Data Retention

The ability of the SSD to retain data over time. Synonymous with *retention* in this document.

3.4 Endurance

The ability of an SSD to withstand multiple data rewrites.

3.5 Endurance failure

A failure that is caused by endurance stressing.

NOTE 1 Endurance failures may be data error (3.2) types or functional failure (3.9) types.

NOTE 2 In an endurance stress, some failures may occur that are unrelated to endurance. For example, a solder joint could fail. A failure is considered unrelated to endurance if it was not caused by the endurance stress itself (i.e., if it was not caused by the repetitive writing of data to the drive). Such non-endurance failures are not considered as part of endurance verification.

NOTE 3 A number of distinct failure mechanisms are responsible for endurance failures, and in general these are accelerated in different ways by temperature and other adjustable qualification parameters. For example, in floating-gate memories failure may be caused by charge trapping (normally accelerated by lower temperatures) in the charge transfer dielectric or by oxide rupturing (normally accelerated by higher temperatures) in the transfer dielectric or in peripheral dielectrics. For these reasons, endurance may depend on temperature but it is not known *a priori* whether high temperature is worse than low temperature, or vice-versa.

3.6 Endurance Rating (TBW rating)

The number of terabytes that may be written to the SSD such that the SSD meets the requirements defined in 5.2.

NOTE 1 Several factors impact the endurance rating including how optimally the wear leveling has been implemented, write amplification factor and the cycling capability of the NAND components. The relationship between TBW, write amplification factor and the wear-leveling efficiency are highly dependent on the workload applied for the characterization of endurance.

NOTE 2 The TBW rating may also be expressed in an alternative form called drive-writes-per-day (DWPD), with the assumed lifetime also stated. The value of DWPD is $TBW/(C*Y*365)$, where TBW is the endurance rating in terabytes written, C is the capacity in terabytes, and Y is the lifetime in years that is stated along with the DWPD rating.

3.7 Erase block

The smallest addressable unit for erase operations, typically consisting of multiple pages.

3.8 Failure

The noncompliance of an SSD to the electrical or physical requirements specified for the device.

NOTE Failures may be permanent or transient. For the purpose of this standard, a permanent failure is an SSD that fails sometime during a reliability stress and continues to fail at the final test at the end of that same stress. A transient failure is an SSD that fails during a reliability stress but passes in the final test at the end of that stress.

3.9 Functional Failure

A failed drive that fails to function properly in a way that is more severe than having a data error.

NOTE See 6.1.2 for discussion of functional failures.

3.10 Functional Failure Requirement (FFR)

The allowed cumulative functional failures over the TBW rating. See 5.2.

3.11 Gigabyte (GB)

Approximately equal to 10^9 bytes when used in reference to SSD capacity. See Annex C (normative): SSD Capacity.

3.12 Host

The computer system, test system, or other device, which writes data to and reads data from the SSD.

3.13 Host writes

Data transmitted through the primary SSD interface to be written to the SSD.

3.14 Logical Block Address (LBA)

The logical address used to reference a data sector (block) in the drive.

NOTE 1 LBA is synonymous with the data sector itself.

NOTE 2 Block in the drive is a logical construct separate from that of an erase block in the NVM.

3.15 Non-Volatile Memory (NVM)

A memory which retains data after the power is turned off.

NOTE 1 Non-volatile memories considered in this standard are capable of being electrically rewritten.

3.16 Page

A sub-unit of an erase block consisting of a number of bytes which can be read from and written to in single operations, through the loading or unloading of a page buffer and the issuance of a program or read command.

3.17 Program/erase cycle (p/e cycle)

The writing of data to one or more pages in an erase block and the erasure of that block, in either order.

3.18 Qualification Family

A group of SSD products which differ only in storage capacity and in minor design details that are directly related to the capacity differences.

NOTE To be considered part of the same qualification family, SSD products must use the same nonvolatile memory products, or different nonvolatile memory products that are themselves part of the same component qualification family (defined in JESD47). The SSD products must also use the same controller and the same firmware, except to the extent that the firmware requires different settings to support the different capacities of the drives. The SSD products must also have the same ratio of TBW specification to capacity; for example, a 100 GB drive with a 100 TBW specification could be in the same family as a 50 GB drive with a 50 TBW specification. Due to the complexity of SSD designs, it is beyond the scope of this specification to completely define what constitutes a qualification family. The burden of proof falls upon the SSD manufacturer who chooses to use the qualification family concept.

3.19 Retention failure

A data error occurring when the SSD is read after an extended period of time following the previous write.

NOTE A number of distinct failure mechanisms are responsible for retention failures, and in general these are accelerated in different ways by temperature and other adjustable qualification parameters. For example, in floating-gate memories, failure may occur due to defects that allow charge to leak through the transfer dielectric or by the detrapping of charge in the transfer dielectric; the former can be weakly accelerated or even decelerated by high temperature, and the latter can be highly temperature accelerated (see JEPI22). For these reasons, retention may depend on temperature but it is not known *a priori* whether high temperature is worse than low temperature, or vice-versa.

3.20 Solid State Drive (SSD)

A solid state drive (SSD) is a non-volatile storage device. A controller is included in the device with one or more solid state memory components. The device should use traditional hard disk drive (HDD) interfaces (protocol and physical) and form factors.

3.21 Terabyte (TB)

For the purpose of this standard, a terabyte is equal to $1 \cdot 10^{12}$ bytes.

3.22 Uncorrectable Bit Error Rate, or ratio (UBER)

A metric for the rate of occurrence of data errors, equal to the number of data errors per bits read. Mathematically,

$$UBER = \frac{\text{number of data errors}}{\text{number of bits read}} \quad (1)$$

NOTE Although the UBER concept is in widespread use in the industry, there is considerable variation in interpretation. In this standard, the UBER values for SSDs are to be *lifetime* values for the *entire population*. The numerator is the total count of data errors detected over the full TBW rating for the population of SSDs, or the sample of SSDs in the endurance verification. A sector containing corrupted data is to be counted as one data error, even if it is read multiple times and each time fails to return correct data. The denominator is the number of bits written at the TBW rating limit, which aligns to the usual definition of errors per bit read when the read:write ratio is unity. See 6.1.1 for a further discussion of UBER calculation.

3.23 Wear leveling

Methods employed by the drive to spread the p/e cycles across the NVM physical locations even when the workload may be unevenly distributed across the logical drive capacity.

3.24 Workload

The detailed sequence of host writes and reads (including data content and timing) applied to the drive during endurance testing. See JESD219 for specification of workload.

3.25 Write amplification factor (WAF)

The data written to the NVM divided by data written by the host to the SSD.

NOTE 1 For the purpose of calculating WAF, data written by the host is considered to be in multiples of drive capacity. For example, if 150GB of data is written to an SSD with capacity of 100GB, the data written by the host is considered to be 1.5. Also, data written to the NVM is considered to be the average number of p/e cycles experienced by NVM blocks in use in the SSD. For example, if the average number of p/e cycles is 3, then the data written to the NVM is considered to be 3. In this example, the WAF would be 2.

NOTE 2 Write amplification factor will depend on the workload and may vary over the lifetime of the device.

4 General SSD Ecosystem Factors

4.1 Capacity

An SSD may replace or co-exist with hard disk drives (HDD's) in an application. See Annex C (normative): SSD Capacity.

4.2 Form Factors

SSD form factors include, but are not limited to, traditional HDD form factors. Many HDD form factors were developed in the SFF Committee. The SFF Committee web site, www.sffcommittee.org, is an excellent resource for finding mechanical specifications of drive form factors and the associated connector locations. Form factor publications that have become JEDEC, EIA, or other recognized standard organization specifications reference the standard number on the first page of the SFF publication.

5 Application classes and SSD Endurance Rating

5.1 Purpose of application classes

Since the endurance of an SSD is dependent up the workload applied to it and the conditions (temperature, duty cycle, etc.) in which that workload is applied, it is necessary to define standard application classes under which endurance for a particular type of device is to be rated and verified. This allows devices within a particular class to be compared to one another as regards the standard endurance rating.

These classes are not all-inclusive and it is understood that variations such as the operating system and application architecture make a significant impact to the workload of the SSD. These classes provide a means to provide parameters for standardized endurance ratings so that the end user may use the endurance rating as a factor in determining if an SSD is suitable for a particular application.

This standard defines two application classes: Client and Enterprise. The quantitative requirements for the application classes are defined in 5.2.

5.2 Definition of Application Classes

This standard is based on a use scenario in which the SSDs are actively used for some period of time, during which the SSDs are written to their endurance ratings (Active Use), followed by a power-down time period in which data must be retained (Retention Use). Table 1 lists the requirements by SSD class.

Table 1 — SSD Classes and Requirements

Application Class	Workload (see JESD219)	Active Use (power on)	Retention Use (power off)	Functional Failure Requirement (FFR)	UBER Requirement
Client	Client	40 °C 8 hrs/day	30 °C 1 year	≤3%	≤10 ⁻¹⁵
Enterprise	Enterprise	55 °C 24hrs/day	40 °C 3 months	≤3%	≤10 ⁻¹⁶

5.2 Definition of Application Classes (cont'd)

The endurance rating shall be determined for the workload and use temperatures required for the particular class of SSD. Since workloads are expected to change as applications evolve, workloads are described in a separate JEDEC standard (JESD219).

The Active-Use and Retention-Use columns specify the temperatures and times for the two time periods of SSD use defined by the standard. For Client, the Retention Use temperature (30 °C) is also the temperature specified for the 16 hrs/day in which the SSD is off (i.e., the client machine is powered down). Temperatures in Table 1 and elsewhere in this standard are SSD case temperatures, as reported by the drive if the drive has temperature reporting capability. The temperatures are intended to represent the relevant temperatures over the respective time periods, for the purpose of endurance and retention estimation, not the absolute max/min specifications to be found on the SSD datasheet. The Retention Use time is the period over which data must be retained with power off. The FFR and UBER columns specify the criteria for functional failures and UBER that must be met for drives written to their endurance rating limits and then being subjected to the specified retention time.

5.3 SSD Endurance Rating

The SSD manufacturer shall establish an endurance rating for an SSD that represents the maximum number of terabytes that may be written by a host to the SSD, using the workload specified for the application class, such that the following conditions are satisfied:

- 1) The SSD maintains its capacity
- 2) The SSD maintains the required UBER for its application class
- 3) The SSD meets the required functional failure requirement (FFR) for its application class
- 4) The SSD retains data with power off for the required time for its application class

This rating is referred to as TBW. Requirements for UBER, FFR, and retention for each application class are defined in 5.2.

In addition to the standard endurance rating determined for the requirements in 5.2, the manufacturer may determine ratings for other sets of requirements as agreed to with purchasers. The methods described in this standard may be used to determine the endurance and retention verification procedures for those additional sets of requirements. See 6.3.

5.4 Estimation of the SSD endurance rating

An informative example of an SSD endurance rating estimation is provided in Annex E (informative): Estimation of SSD Endurance Rating. The estimation method shown in Annex E (informative): Estimation of SSD Endurance Rating is not suitable to general users of the drive because the method requires information about the drive design and data from the drive itself that in general are available only to the manufacturer. The method in Annex E (informative): Estimation of SSD Endurance Rating is also limited in scope to SSDs which use NAND components qualified according to JESD47. The method in Annex E (informative): Estimation of SSD Endurance Rating is not sufficient for verification that an SSD meets the endurance requirements. Verification of an endurance rating is defined in clause 0.

6 Endurance Test Method

This section describes two approaches for endurance verification, verification based on the Direct Method (6.1), and verification based on Extrapolation Methods (0). Both consist of endurance verification followed by retention verification. With the direct method, SSDs are written to their TBW ratings and put through retention testing as well. The direct method shall be followed if the full TBW rating may be reached within a 1000 hour stress cycle. If this is not possible, then the extrapolation approach is acceptable. The manufacturer may use the direct method even if this takes longer than 1000 hours. If an SSD product from a qualification family has been qualified using this standard, then subsequent products need only data from a 1000-hour direct method evaluation, even if this results in those drives not being fully stressed to their TBW rating limits.

6.1 Direct method

Figure 1 shows the flow chart of the Direct Method. In this method, SSDs are stressed to their stated endurance rating (in TBW) using specified reference workloads. The endurance stressing shall be performed at both high and low temperatures. Following this endurance stressing, retention testing shall be performed. Since the Retention Use time requirements (Table 1) are long, extrapolation or acceleration is required to validate that the SSD meets the retention requirement.

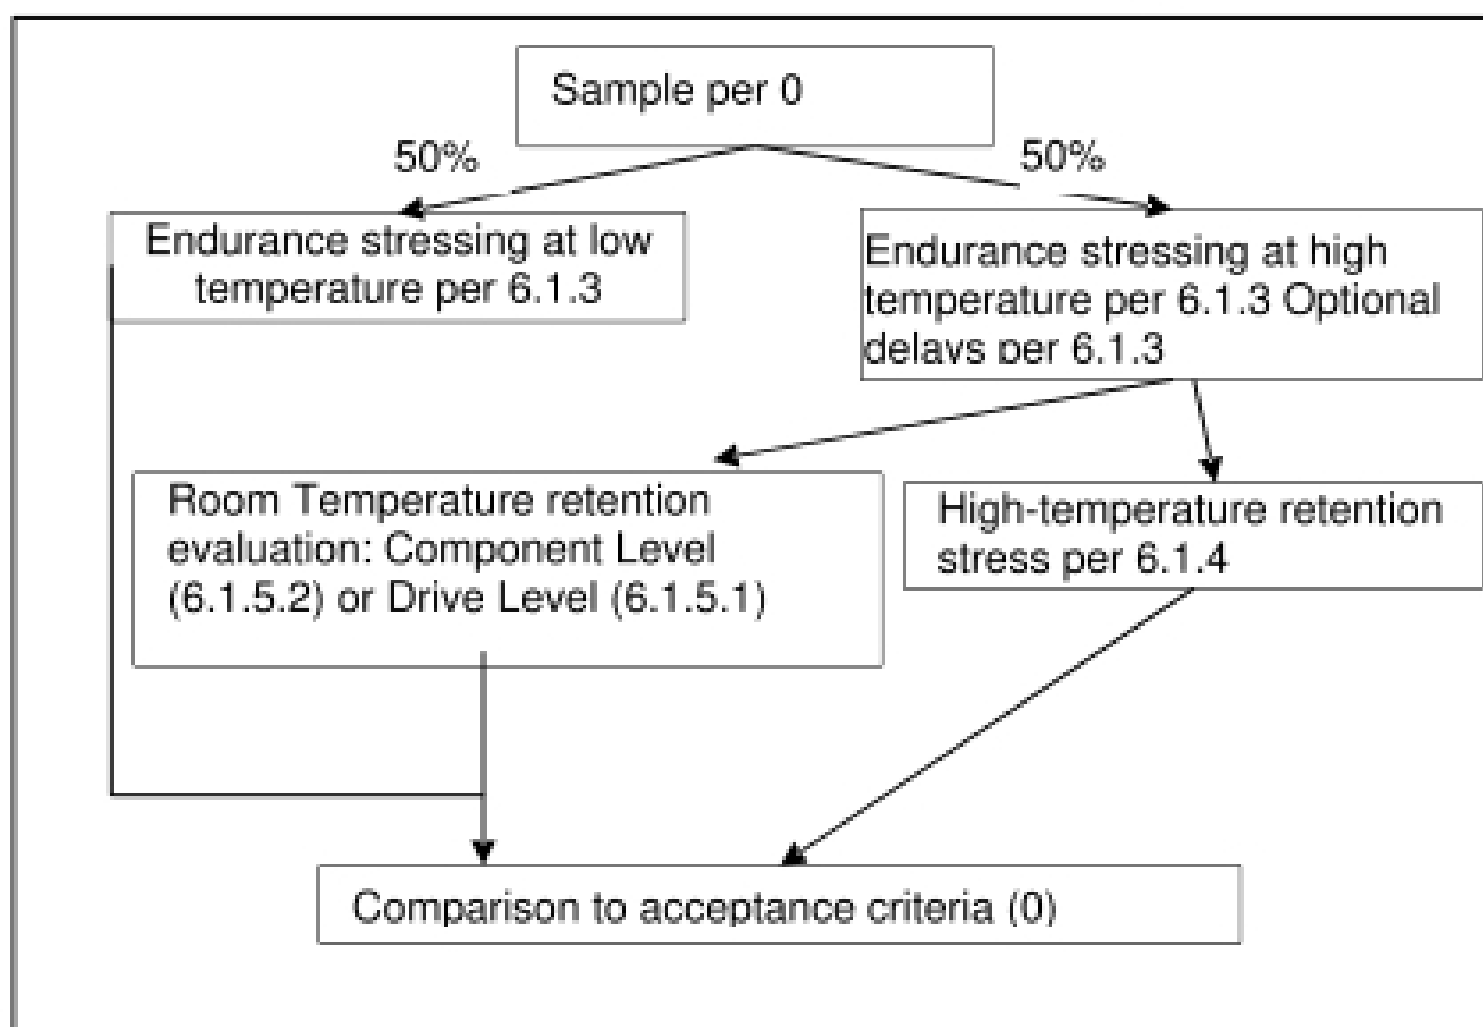


Figure 1 — Flowchart for Direct Method

6.1 Direct method (cont'd)

6.1.1 Sample Size and Acceptance Criteria

For the first product to be qualified in a qualification family, the sample shall consist of SSDs from at least three nonconsecutive production lots and from all fabrication plants responsible for the manufacture of the nonvolatile memory used in the SSD. For subsequent products from a qualification family, a single production lot is sufficient. The number of SSDs in the sample shall be sufficient to establish that both the FFR and UBER requirements are met at 60% confidence.

The sample sizes and acceptance criteria are defined by the following two equations, which mathematically embody the 60% requirement of the previous sentence:

$$UCL(\text{functional_failures}) \leq FFR \cdot SS \tag{2}$$

$$UCL(\text{data_errors}) \leq \min(TBW, TBR) \cdot 8 \cdot 10^{12} \cdot UBER \cdot SS \tag{3}$$

where

functional_failures is the acceptable number of functional failures

data_errors is the acceptable number of data errors

min(x,y) is the minimum of x and y (see also 6.1.3)

FFR and UBER are the drive requirements expressed as fractions

TBW is the endurance rating in terabytes written

TBR is the number of TB read (see 6.1.3)

SS is the sample size in number of drives

UCL(·) is an upper confidence limit function defined in Table 2 (see JESD47 for background)

Table 2 — Values of UCL(x)

AL	n	AL	n	AL	n	AL	n	AL	n
0	0.92	20	21.84	40	42.30	60	62.66	80	82.97
1	2.03	21	22.87	41	43.32	61	63.68	81	83.98
2	3.11	22	23.89	42	44.35	62	64.69	82	84.99
3	4.18	23	24.92	43	45.36	63	65.71	83	86.00
4	5.24	24	25.94	44	46.38	64	66.72	84	87.02
5	6.29	25	26.97	45	47.40	65	67.74	85	88.03
6	7.34	26	28.00	46	48.42	66	68.75	86	89.05
7	8.39	27	29.02	47	49.43	67	69.77	87	90.06
8	9.43	28	30.04	48	50.46	68	70.79	88	91.08
9	10.48	29	31.07	49	51.47	69	71.80	89	92.08
10	11.52	30	32.09	50	52.49	70	72.82	90	93.10
11	12.55	31	33.12	51	53.51	71	73.83	91	94.11
12	13.59	32	34.14	52	54.52	72	74.85	92	95.13
13	14.62	33	35.16	53	55.55	73	75.86	93	96.14
14	15.66	34	36.18	54	56.56	74	76.88	94	97.15
15	16.69	35	37.20	55	57.58	75	77.89	95	98.16
16	17.72	36	38.22	56	58.60	76	78.91	96	99.18
17	18.75	37	39.24	57	59.61	77	79.92	97	100.19
18	19.78	38	40.26	58	60.63	78	80.94	98	101.21
19	20.81	39	41.29	59	61.64	79	81.95	99	102.22

6.1 Direct method (cont'd)

6.1.1 Sample Size and Acceptance Criteria (cont'd)

A common sampling plan is called an accept-on-zero plan because the evaluation will pass if there are zero failures and fail if there are greater than zero failures. For an accept-on-zero plan, the UCL value is 0.92.

NOTE Equations 2 and 3 represent two separate sample-size requirements and two separate acceptance criteria. In a particular situation one of the two equations is the limiting one for sample size purposes. Which equation is the limiting one for the acceptance criteria depends on the number of functional failures and data errors in the verification.

EXAMPLE FFR=3%, UBER= 10^{-16} , TBW=100, all written data read back and verified, accept-on-zero plan.

$$SS \geq 0.92/(0.03)=30.1 \text{ (from equation 2)}$$

$$SS \geq 0.92/(100 \cdot 1 \cdot 8 \cdot 10^{12} \cdot 10^{-16}) = 11.5 \text{ (from equation 3)}$$

Therefore, the required sample size is 31 (the larger of the two results). If the minimum sample size of 31 were chosen, then the verification is passed if there are no functional failures in 31 drives and if the UBER requirement of equation 3 is also met. The latter requirement may be calculated as:

$$UCL(\text{data_errors}) \leq 100 \cdot 1 \cdot 8 \cdot 10^{12} \cdot 10^{-16} \cdot 31 = 2.48$$

From Table 2, up to one data error is acceptable. Therefore, the verification is passed if there are no functional failures and less than or equal to one data error.

NOTE UBER is defined in terms of bits read (see 3.22) but for the purposes of endurance verification equation 3 counts the minimum of bits read and bits written. The rationale is twofold. First, many data errors are transient with respect to rewriting of an SSD, but repeatable with respect to repeated reading. This means that a sector with corrupted data may pass without error if rewritten, however reading non-failing sectors multiple times is unlikely to detect additional errors. This means that if reads are less frequent than writes then many data errors may be missed. All data errors are detected only if all written data is read before those sectors are rewritten. If the TBR is less than the TBW, then the UBER should be increased because of the likelihood that transient data errors went undetected. Using the TBR in place of the TBW accomplishes that goal. Second, this standard is aligned to a reference read:write ratio of unity. If the TBR is equal to the TBW, then the UBER may be considered to be an error rate per bit read or per bit written; both are equivalent. If the TBR in the endurance stress is greater than the TBW then the UBER is based on TBW. Multiple reads of the same written data may be performed for other purposes during endurance verification but may not be used to increase the right-hand-side of equation 3.

6.1 Direct method (cont'd)

6.1.2 Categorization of failures

Failing SSDs are to be divided into three categories: non-endurance failures, endurance functional failures, and endurance data errors. Non-endurance failures are to be excluded from consideration in the endurance verification but should of course be considered if relevant to other parts of drive qualification. The number of functional failures is to be held against the FFR acceptance criterion (equation 2), and the number of data errors is to be held against the UBER acceptance criterion (equation 3).

Failures are to be categorized as non-endurance failures only if compelling evidence exists that they were not caused by the act of writing the drive to its endurance limit, or by the subsequent retention stress. Failures that are not in the circuit path of the written data (for example, failures isolated to power supplies and capacitors) may be considered non-endurance failures. Failures in the circuit path of the written data, particularly the controller and the nonvolatile memory, are more often be considered endurance failures, but there may be exceptions. Failures that are in the circuit path of the written data may be considered non-endurance failures if the cause of the failure is unrelated to the quantity of data written (e.g., if an oxide breakdown event occurred in the controller and this was the result of the time under bias and the temperature rather than the TBW).

The division of endurance failures between functional failures and data errors is clear-cut when the drive has either failed catastrophically or has suffered only a single corrupted sector. For intermediate cases, involving multiple data errors or partial-drive functionality failure, the manufacturer may use discretion when deciding upon the division.

6.1.3 Endurance Stress Phase

The drives are stressed to their full endurance specification (in TBW) using the agreed-upon workload. The stress time is the time required to reach the TBW rating with the stress equipment being used. If performance variation between stress systems or the SSDs themselves cause some SSDs to receive more writes than others in a given stress time, then the average amount of data written per drive shall reach the endurance specification. All drive errors throughout the stress shall be recorded, even if those errors are transient in nature; testing of the drive merely at the end of the stress is not acceptable. Reading back and verifying the data ideally would involve the stress system retrieving the data previously written to the SSD and comparing that data to a separate, master copy stored by the stress system. Such a complete verification of data may be impossible with some stress systems. In such cases, it is sufficient for the purposes of endurance verification for the stress system to rely on the SSDs error-detection capability to verify the data, provided that the SSD is known to contain strong error-detection capability (>90%).

The sample is divided into two groups. One group undergoes endurance stressing at low temperature and the other at high temperature. The low and high temperatures for the endurance stress are defined by SSD class in Table 3 and Table 4. When the temperature varies from one SSD to the other, for example because of different positions in a stress chamber, the average temperature of the group shall be used.

The intent of the low-temperature (25°C) endurance leg is to evaluate endurance near the lower end of use temperature but not so low as to require refrigeration. Therefore, in this case the stress temperature is considered the ambient temperature rather than the SSD case temperature. Best practice is to use sufficient air circulation to achieve as low a temperature as practical for this room-temperature stress.

6.1 Direct method (cont'd)

6.1.3 Endurance Stress Phase (cont'd)

The intent of the high-temperature endurance leg is twofold: 1) to evaluate endurance near the upper end of use temperature; and 2) to accelerate recovery effects by enough to match use conditions. Regarding point one, Table 3 specifies that the temperature used be at least as high as the use temperatures defined in Table 1. Regarding point two, it is known for NAND Flash memory that damage created with each p/e cycle partially recovers or heals during the delays between p/e cycles (see JESD22-A117). Therefore, an endurance stress test that is performed over just a few weeks results in more net damage than would be experienced in normal use over several years. The main effect of this higher net damage is to reduce the data retention capability compared to the capability that would exist in real use. To avoid this problem, Table 3 and Table 4 specify target stress temperatures that have been calculated so that the amount of recovery matches what would occur in the actual use conditions specified in Table 1. The target temperatures vary with the endurance stress time because the effect of temperature is to accelerate the rate of recovery, so a longer stress time does not require as high a temperature as a shorter stress time.

The target time/temperature values in Table 4 are calculated so that the endurance stress time is equivalent to 1.5 years at the Active-Use temperature and hours/day shown in Table 1 assuming an activation energy of 1.1 eV. See Annex A (normative): Calculations of Temperature-Accelerated Stress Times for the detailed calculations. Although in actual use the endurance limit is expected to be reached only after 3-5 years, and matching actual use is the intent of this standard, the calculations were performed for the more severe case in which the full TBW occurs in only 1.5 years. This is a conservative assumption because a shorter stress time allows less recovery between writes and therefore results in higher net damage at the end of the stress. This assumption is made to add margin against possible inaccuracies in the 1.1 eV acceleration model for high temperature data retention.

Ideally, the high-temperature stress leg would be run for times and temperatures that exactly match one of the pairs of values in Table 4, but an exact match may be impossible. For example, the SSDs may reach their endurance limit sooner or later than had been expected, or different models of SSD may share the same stress system and run at different case temperatures even though they share the same ambient temperature. It is acceptable to run an endurance verification test at a different time/temperature than the target values as long as the temperature meets the minimum allowed level specified in Table 3. If the actual temperature or actual time is higher than the target specified in Table 4, it is required to increase the retention bake time or temperature following the method of 6.1.4. If the time or temperature is below the target condition specified in Table 4, the method of 6.1.4 may be used to reduce the time or temperature of the retention bake, at the discretion of the manufacturer.

Table 3 — Endurance and retention times and temperatures by drive class

Class	Active Use Temperature	Retention Use Temperature	Retention Use Time	Endurance Stress Temperature	High Temperature Retention Stress Temperature
Client	40 °C (8 hrs/day)	30 °C	1 year	Low: T≤25 °C ambient High: Min 40 °C Target: Table 4	96hrs / T≥66 °C or 500hrs / T≥52 °C
Enterprise	55 °C (24hrs/day)	40 °C	3 months	Low: T≤25 °C ambient High: Min 60 °C Target Table 4	96 hrs / T≥66 °C or 500hrs / T≥52 °C

6.1 Direct method (cont'd)

6.1.3 Endurance Stress Phase (cont'd)

Table 4 — Target endurance high temperature (T_{tar}) vs. endurance stress times

Actual endurance stress hours	Client	Enterprise
50	82.7	109.9
100	76.0	102.1
150	72.1	97.6
200	69.5	94.6
250	67.4	92.2
300	65.8	90.3
350	64.4	88.7
400	63.2	87.4
450	62.2	86.2
500	61.2	85.1
600	59.7	83.3
700	58.3	81.8
800	57.2	80.4
900	56.2	79.3
1000	55.3	78.3
1100	54.5	77.4
1200	53.8	76.5
1300	53.1	75.8
1400	52.5	75.0
1500	51.9	74.4
1600	51.4	73.8
1700	50.9	73.2
1800	50.4	72.7
1900	50.0	72.2
2000	49.5	71.7
2200	48.8	70.8
2400	48.1	70.0
2600	47.4	69.3
2800	46.8	68.6
3000	46.3	68.0

Intentional delays may be inserted into the endurance stressing as long as this follows the guidelines for such delays in JESD22-A117. In particular, the delays may not be inserted solely *at the end* of the endurance stressing. These delays, combined with the effect of the elevated-temperature endurance stressing, shall stay within the 1.5-year equivalent time discussed in the previous paragraph. If the T_{max} values from Table 4 are used, the full allowance is used up and no additional delays may be inserted. See Annex A (normative): Calculations of Temperature-Accelerated Stress Times

A delay may consist of the drive being powered down, or being powered up but not being written to. If the drive is powered up, the manufacturer shall ensure that the drive does not perform any background maintenance that would not be possible during actual use. The direct method shall be used whenever the active portion (excluding delays) of the endurance method can be completed within 1000 hours. Delays at elevated temperature are not allowed in the low-temperature leg.

6.1 Direct method (cont'd)

6.1.3 Endurance Stress Phase (cont'd)

In addition to the writing and reading back of data, other stresses and operations may be added to the endurance stressing to suit purposes beyond that of endurance verification, as long as those stresses and operations do not conflict with the limits on delays specified in this section. For example, power interruptions may be added to test the drive for its ability to handle such interruptions.

6.1.4 High Temperature Retention Stress

After the endurance stress, drives are to be powered down and baked at elevated temperature in order to establish the data retention capability. Only the drives stressed at high temperature are to be baked. The drives are to be fully written with data prior to the bake and fully read back afterwards. The number of data errors resulting from the retention stress is to be added to the number resulting from the endurance stress, and the total shall meet the acceptance criterion for UBER (see 3).

Table 3 shows the required conditions for the retention verification. Two equivalent options are given for the bake, both of which are considered equivalent. The qualification retention temperatures and times are chosen to correspond to the required data retention times for the common temperature-accelerated mechanism responsible for data degradation in non-volatile memories, assuming an activation energy of 1.1 eV and the class retention requirements. If a component of the SSD (other than the nonvolatile memory) is not able to tolerate the indicated temperatures, that component may be removed or replaced for the purpose of the retention bake.

As discussed in 6.1.3, the high-temperature endurance stress may be run at a time/temperature pair that differs from the target values in Table 4, and that may call for adjusting the retention bake time or temperature. Whether to make an adjustment in bake time or temperature is up to the discretion of the manufacturer.

If an adjustment is to be made in bake temperature, the bake temperature is to be changed by the difference between the actual endurance temperature and the temperature called out in Table 4, but the temperature is not to be reduced by more than 7 °C. For example, suppose that a client SSD received its endurance stress over 1000 hours at 48 °C, whereas Table 4 specifies a stress temperature of 55 °C for that stress time. The bake condition changes from 96 hours at 66 °C (Table 3) to 96 hours at 59 °C. If the endurance stress temperature had been lower, no further reduction would be allowed, because of the maximum allowed reduction of 7 °C. If the endurance stress temperature is higher than called for in Table 4, the full offset is to be used to increase the bake temperature, not limited to 7 °C.

If the adjustment is to be made in bake time, the time is to be scaled by the following factor:

$$\text{Adjustment Factor} = (\text{Actual stress time}) / (\text{Stress time in Table 4 for that stress temperature}) \quad (4)$$

But not to be less than 0.5

For example, in the previous example, Table 4 specifies a stress time of 2200 hours when the temperature is 49 °C. The adjustment factor would then be 0.45 (1000/2200), rounded up to 0.5 because of the restriction limit of 0.5. From Table 3, the bake time at 66 °C would be 96 hours before the adjustment and 48 hours after adjustment. If the stress temperature were higher than called for in Table 4, increasing the bake time would be mandatory with no upper limit.

The two adjustment methods are considered equivalent and are left to the discretion of the manufacturer. The limits on the reductions in bake time and temperature are imposed to prevent the stress conditions from extending into a range of time and temperature for which the acceleration model for data retention may become inaccurate and optimistic.

6.1.5 Room temperature retention evaluation

Not all mechanisms responsible for data loss are accelerated by temperature, and therefore a second, room-temperature evaluation is required. This requirement holds only for the first product in a qualification family to be qualified; subsequent products are exempt.

Two methods are acceptable for performing this evaluation: drive-level and component-level. Both require mathematical extrapolations, since achieving time acceleration via higher temperature is impossible. The nature of these extrapolations requires access to drive-level or component-level bit-error-rate data that may be available only to the manufacturers.

With this method, one measures the rate of increase of the bit error rate during a qualification retention period and extrapolates that trend to the desired retention time. The details of the method are given in Annex B (normative): Assessment of Room-Temperature Data Retention.

6.1.5.1 Drive-level method for room temperature retention evaluation

As shown in Figure 1, the SSDs are subjected to a room-temperature data retention stress. The step is performed only on the drives undergoing high-temperature endurance stressing. The drives are to be written with data and then set aside at room temperature for a retention period of at least 500 hours (longer times are encouraged for greater accuracy). The drives are to be read back and the bit error rate measured with internal error correction bypassed. The extrapolation is to follow the method of JEP122. There are two variations of this method that are acceptable. The first is to measure the bit error rate of the drives at several time points (for example, 48, 168, and 500 hours) and to extrapolate the trend. The second is to measure the bit error rate only at the end of the retention period, and then to extrapolate the bit error rate by using the time-dependence obtained from characterization of the non-volatile memory components. If the SSD uses multiple types of nonvolatile memory components (i.e., different component product numbers), then the extrapolations shall be done separately for each. If the retention test is done for the full required retention time (such as 3 months for enterprise SSDs) then no extrapolation is necessary; the drives may simply be tested for data errors.

6.1.5.2 Component-level method for room temperature retention evaluation

With this method, the extrapolation method described in 6.1.5.1 (and JEP122) is performed using raw bit error rate data from nonvolatile memory components rather than the actual SSDs. At the completion of the endurance stress, the number of program/erase cycles experienced by the nonvolatile memory components inside the SSDs is to be determined, along with the variation in cycle count from one component or location to another. Retention data for identical types of components are to be obtained after the same number of program/erase cycles and the same amount of variation. The component-level retention data should be obtained from at least three non-consecutive production lots of such components. It is not necessary for the component sample to contain as many components as would be contained in the full sample size of drives calculated using the method of 6.1.1.

6.2 Extrapolation method

An extrapolation method is used if the direct method would require more than 1000 hours of endurance stress. In this version of the specification, only a single extrapolation method is defined. This section describes the acceptable extrapolation method. Note that the extrapolation method may require special access to SSD internal operations, which make such methods possible only for the manufacturer of the SSD. Use of extrapolation methods shall be documented, and agreed upon between manufacturer and purchaser.

6.2.1 Short Stroke Extrapolation Method

The capacity of an SSD is artificially reduced so that some nonvolatile memory components or blocks are not written to, while the remaining ones are cycled faster. Colloquially this method is often referred to as 'short stroke' in similarity to techniques used with hard disk drives to access only a portion of the storage platter. For SSD endurance verification, this method is identical to the direct method (see 6.1) except that the SSDs are stressed until the stressed fraction of the nonvolatile memory components reach the target p/e cycles. Target p/e cycles is defined as the number of p/e cycles expected with the standard workload at the actual TBW rating. For example, suppose that a 100GB SSD has an endurance rating of 100 TBW. Suppose further that it is determined that writing 100GB of data with the standard workload results in two p/e cycles on the nonvolatile memory (i.e., that there is a write amplification factor of two). Finally, suppose that internal control features of the SSD design keep different nonvolatile memory components and blocks within $\pm 10\%$ of the mean number of p/e cycles. In that case, the p/e cycles expected at the end of 100TBW would be expected to be 1000 with a range from 900 to 1100. A short-stroke endurance test that caused this same result (900 to 1100 p/e cycles) would be considered to have brought the SSD to the target p/e cycles.

It is not required that the reduced drive capacity be compensated by an increase in the sample size. Equations (2) and (3) still apply, with the understanding that the TBW term in (3) is to be the terabytes actually written to the reduced-capacity drive, rather than the TBW rating of the drive. Specifically, the requirements become:

$$UCL(\text{functional_failures}) \leq FFR \cdot SS \quad (5)$$

$$UCL(\text{data_errors}) \leq \min(TBW_A, TBR_A) \cdot 8 \cdot 10^{12} \cdot UBER \cdot SS \quad (6)$$

where TBW_A is the actual number of terabytes written to each reduced-capacity drive in the endurance stress, and TBR_A is the actual number of terabytes read.

For this approach to be used, the manufacturer is to ensure that the method of capacity reduction does not significantly distort the normal internal workings of the SSD. For example, the number of spare memory blocks may need to be reduced to ensure that the write amplification factor and the ability of the SSD to tolerate a bad blocks does not change.

Simply reducing the logical span of written data is in general not sufficient, since the SSD controller and firmware will make use of the full nonvolatile memory capacity if not instructed (through modified firmware) to not do so.

6.3 Endurance verification given alternative requirements

6.1 and 0 describe the requirements for verifying that an SSD meets the endurance and retention requirements for the defined JEDEC SSD classes of Table 1. In some cases alternative requirements may be specified by the manufacturer and/or purchaser. In such cases, the equations and tables in this standard may be used to calculate alternative sample sizes and stress conditions. For example, an FFR of 3% could be viewed as representing an annualized failure rate (AFR) of 0.6% averaged over a five-year life, or a mean time between failures (MTBF) of 5 years divided by 0.03, or about 1.5 million hours. Given alternative requirements for failure rate, one could choose an FFR based on such reasoning, and use the sample size and acceptance equations in this document. Another example is that the operating temperature for an SSD may be different than specified for the JEDEC SSD class. In that case, the stress times and temperatures may be adjusted according the Arrhenius acceleration model described in this standard (see Annex A (normative): Calculations of Temperature-Accelerated Stress Times). Alternative requirements and acceptance criteria are up to the manufacturer and purchaser to agree upon.

7 Apparatus and Precautions

7.1 Apparatus

The apparatus required for this test shall consist of a controlled temperature chamber capable of maintaining the specified temperature conditions to within ± 5 °C. Cables, racks, and/or other mounting and connection means shall be provided within the chamber so that reliable electrical contact is made to the drive connector in the specified circuit configuration. Power supplies and biasing networks shall be capable of maintaining the specified operating conditions throughout the test. Also, the test circuitry should be designed so that the existence of abnormal or failed devices does not alter the specified conditions for other drives on test. Care should be taken to avoid possible damage from transient voltage spikes or other conditions that might result in electrical, thermal or mechanical overstress.

7.2 Precautions

Precautions shall be taken to ensure that no devices can be damaged by thermal runaway and to preclude electrical damage. The test setup should be monitored initially and at the conclusion of a test interval to establish that all devices are being stressed to the specified requirements. The bias voltages and currents on each device shall be noted and corrected prior to further temperature exposure. If a device is not biased properly when checked at the conclusion of a test interval, it shall be determined if the device has changed or if the test circuit has changed so that the validity of the data for qualification may be established.

Because NAND Flash damage partially recovers between cycles, especially at high temperature, delays between the end of the endurance stress and the data write before retention bake should be minimized. That is especially true while the SSDs are in the high-temperature chamber (delays at room temperature after removal from the chamber are not a significant concern). This standard requires the endurance limit to be reached only on average (see 6.1.3), so it is possible to run all SSDs for the same stress time and remove them immediately when the average reaches the endurance limit.

8 Summary

The following details shall be specified in the applicable device specification and/or the manufacturer's internal stress test specification, along with the rationale:

- a) Special mounting, if applicable.
- b) Test condition.
- c) Biasing conditions.
- d) Measurements before, at intermediate test points, (if applicable) and after test.
- e) The workload chosen.
- f) The endurance stress time and temperatures, and temperature-ramp conditions if the ramped-temperature approach is used
- g) Duration and temperature of any delays inserted at intermediate points during endurance stressing
- h) Data retention pattern, duration and temperature.
- i) The details of the calculations of UBER and functional failure requirement, including any extrapolations
- j) If the extrapolation methods of 6.2 are used, the details of the steps taken.

Annex A (normative): Calculations of Temperature-Accelerated Stress Times

The governing equation for the calculations in Table 4 and in the calculation of allowed additional delays is:

$$t_D \cdot A(T_D) + t_S \cdot [FH_S \cdot A(T_{S,H}) + (1 - FH_S) \cdot A(T_{S,L})] \leq t_U [FH_U \cdot A(T_{U,H}) + (1 - FH_U) \cdot A(T_{U,L})]$$

where

t	=	time (in any units as long as all t values are in the same units)
A(T)	=	Arrhenius acceleration term, $\exp(-E_A/KT)$
T	=	Temperature in °K (Celsius temperature plus 273.15)
E _A	=	Activation energy, 1.1 eV
K	=	Boltzmann's constant, $8.6171 \cdot 10^{-5}$ eV/°K
FH	=	Fraction of time spent at high temperature
D	=	Subscript denoting added delays as part of the endurance stress
S	=	Subscript denoting the endurance stress itself
U	=	Subscript denoting the use condition (Table 1)
H	=	Subscript denoting the high temperature of interest
L	=	Subscript denoting the low temperature of interest

This equation expresses the requirement that intentional delays plus the endurance stress time itself add up, after appropriate temperature accelerations are taken into account, to no more than the Active Use time (t_U) which is taken to be 1.5 years as described in 6.1.3. The use of Arrhenius acceleration with activation energy of 1.1 eV for the endurance stress (or NVM p/e cycling) prior to high-temperature retention bake is discussed in JEP122. The subscripts H and L refer to stress or use conditions in which the SSD spends part of its time at a higher temperature and the rest at a lower temperature. When the left side of the equation equals the right side of the equation, the endurance stress performed in endurance testing fairly represents the stress that would occur in the use condition.

For example, consider the client condition with a 1000 hour stress time with no additional delays. The temperature of the endurance stress is to be solved for, such that it fairly represents the client use condition of 8 hours/day at 40 °C and 16 hours/day at 30 °C. As mentioned in 6.1.3, the usage lifetime is assumed to be 1.5 years (rather than a more realistic 3-5 years) to provide some margin for error. The variables in the equation take the following values:

t_D	=	0 (no delays)
t_S	=	1000 hours
t_U	=	1.5 years, or 13149 hours
$T_{S,H}$	=	Endurance stress high temperature, to be solved for
$T_{S,L}$	=	25 °C
$T_{U,H}$	=	40 °C (the active use temperature from Table 1)
$T_{U,L}$	=	30 °C (the power-off temperature from Table 1)
FH_S	=	1
FH_U	=	1/3 (Table 1 shows that the client drive is power up 8 hours/day)

Annex A (normative): Calculations of Temperature-Accelerated Stress Times (cont'd)

$$\begin{aligned}
 1000 \cdot [1.0 \cdot A(T_{S,H}) + 0.0 \cdot A(T_{S,L})] &\leq 13149 \cdot \left[\frac{1}{3} \cdot A(T_{U,H}) + \frac{2}{3} \cdot A(T_{U,L}) \right] \\
 A(T_{S,H}) &\leq \frac{13149}{1000} \cdot \left[\frac{1}{3} \cdot A(T_{U,H}) + \frac{2}{3} \cdot A(T_{U,L}) \right] \\
 &\leq 13.149 \cdot \left[\frac{1}{3} \cdot 1.978 \cdot 10^{-18} + \frac{2}{3} \cdot 5.156 \cdot 10^{-19} \right] \\
 &\leq 1.319 \cdot 10^{-17} \\
 e^{-1.1/KT_{S,H}} &\leq 1.319 \cdot 10^{-17} \\
 T_{S,H} &\leq 338.4^{\circ}\text{K or } 55.3^{\circ}\text{C}
 \end{aligned}$$

Hence the T_{tar} listed in Table 4 for the client flow, for a 1000-hour stress, is 55 °C.

If the choice were made to perform the stress at 50 °C instead of 55 °C, the calculation of the allowed additional delays would be as follows.

$$\begin{aligned}
 \frac{t_D \cdot A(T_D)}{1000} &\leq 1.319 \cdot 10^{-17} - A(T_{S,H}) \\
 &\leq 1.319 \cdot 10^{-17} - 6.9851 \cdot 10^{-18} \\
 t_D \cdot A(T_D) &\leq 6.205 \cdot 10^{-15}
 \end{aligned}$$

If one were interested in inserting 100 hours of additional delays, the temperature of that delays would need to satisfy $A(T_D) \leq 6.205 \cdot 10^{-17}$, from which it follows that $T_D \leq 342^{\circ}\text{K}$ or 69°C . Therefore, the endurance stress would consist of 1000 hours of active endurance stress with the temperature at 50 °C, with an additional 100 hours spend in a non-writing mode of operation at no greater than 67 °C. As discussed in JESD22-A117 and 6.1.3, the 100 hours of delays shall not be inserted after the endurance stress, but rather shall be inserted within the endurance stress.

Annex B (normative): Assessment of Room-Temperature Data Retention

JEP122 documents a reliability model for the retention mechanism referred to as SILC which possesses negligible temperature dependence. This Annex documents an extrapolation method for verifying the retention lifetime for this mechanism without relying on elevated temperatures.

The fraction of bits that are in error, called the raw bit error rate (RBER), depends on the p/e cycle count and the retention time. For $BER \ll 1$,

$$RBER = RBER_0 + B_0 * t^m \quad D1$$

where

- RBER₀ = RBER at the beginning of the retention period, due to effects other than SILC
- B₀ = arbitrary scale factor, dependent upon materials and process details
- t = retention time
- m = retention-time power law coefficient, typically 1 to 2

To verify the useful retention lifetime, the RBER may be measured as a function of time and the parameters RBER₀, B₀, and m fit to equation D1. The resulting fitted equation may then be used to estimate the RBER at the desired retention time in Table 1. The extrapolated RBER shall be below the ECC capability of the SSD controller.

The ECC capability may be calculated from the details of the ECC method used and the allowable UBER from Table 1, using standard methods in the field of ECC. The capability calculated in this way is accurate only if the bit errors are randomly distributed, which is not perfectly the case in real devices. Accordingly, care must be taken to maintain a safety margin between the calculated ECC capability and the RBER. This margin shall be chosen to account for the variation of the BER from device to device and from one location of the device to another.

For example, suppose that the following RBER data are obtained:

Retention Time (Hours)	RBER
0	0
48	5.63E-08
168	2.23E-07
500	7.41E-07
1000	1.59E-06

These data points would be extrapolated in time as shown below. In this example, the data fit a pure power-law relationship with zero for the RBER₀ term. The ECC capability calculated from ideal zero-variation assumptions is determined to be 4×10^{-5} , and based on an evaluation of the variation from device to device and location to location the decision is made to guardband the result by a factor of two. The extrapolated RBER is therefore compared against a requirement of 2×10^{-5} .

Annex B (normative): Assessment of Room-Temperature Data Retention (cont'd)

As shown in Figure 2, the RBER reaches this level after 10,000 hours. Since this exceeds the retention times in Table 1, the room-temperature retention requirement is met.

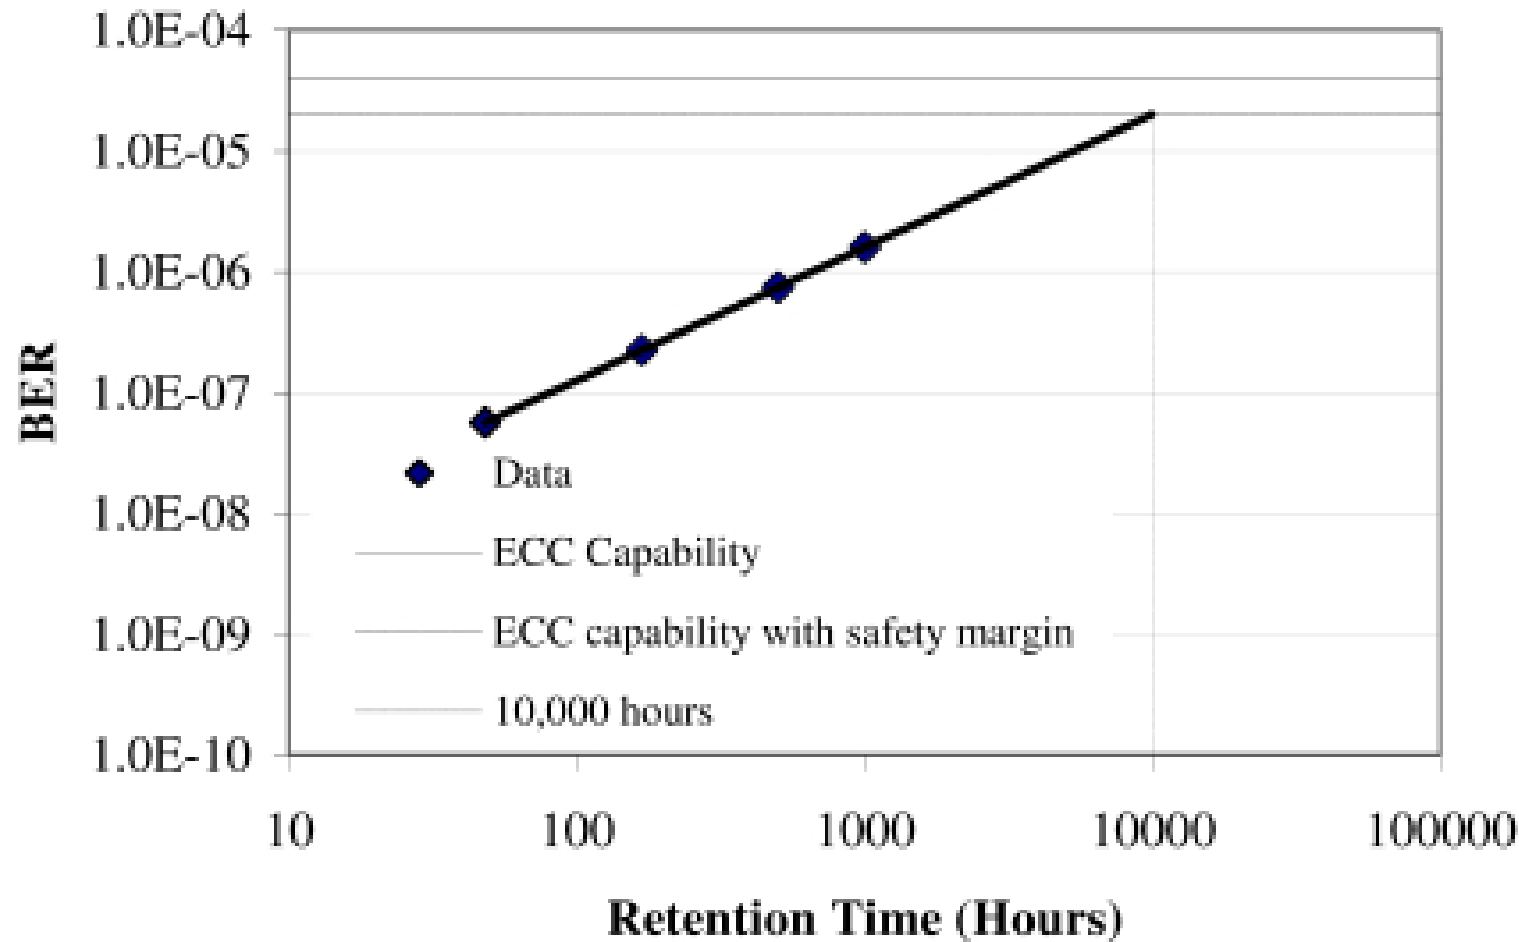


Figure B.1 — Example of BER Extrapolation

Variations of this approach are possible based on the power-law time dependence of Equation D1. For example, an SSD may be able to report the distribution of error counts (e.g., what fraction of sectors have 1 failing bits, 2 failing bits, 3 failing bits, etc.). Based on Equation D1, this distribution is expected to shift towards higher cycles with time, following a power-law relationship which may be extrapolated.

Annex C (normative): SSD Capacity

C.1 SSD Capacity General Overview

An SSD may replace or co-exist with hard disk drives (HDD's) in an application. To be consistent with the established for HDD capacity, SSD capacity shall be defined as user-addressable capacity as calculated using the following formulas (nonvolatile memory areas reserved for device use are not included in this calculation) defined in SFF-8447, where the following definitions apply:

- **^:** the caret character (^) is used as a surrogate symbol for superscript and exponentiation (e.g., b^n for b raised to the n -th power).
- **4xxx:** a reference to logical block sizes of 4096, 4160, 4192, or 4 224 bytes.
- **5xx:** a reference to logical block sizes of 512, 520, 524, or 528 bytes.
- **advertised capacity:** a number of bytes used to describe the capacity of a drive expressed in IEC 60027-2 units (e.g., 250 GB (250,000,000,000 bytes), or 3 TB (3,000,000,000,000 bytes)).
- **ceiling:** a function with two operands. If the first operand is an integer multiple of the second then the function returns the value of the first operand; otherwise, the function returns the first integer multiple of the second operand which is higher than the first operand. That is, $\text{ceiling}(x,y)$ resolves to x rounded up to the nearest multiple of y .
- **floor:** a function with two operands. If the first operand is an integer multiple of the second then the function returns the value of the first operand; otherwise, the function returns the first integer multiple of the second operand which is lower than the first operand. That is, $\text{floor}(x,y)$ resolves to x rounded down to the nearest multiple of y .
- **logical block size:** the number of bytes in a logical block that are user data and is neither SCSI Protection Information nor other information that may not be accessible to the application client.
- **low capacity drive:** a drive with an advertised capacity of 8000 GB or less for a 3.5" form factor (see SFF-8300) device or a 2.5" form factor (see SFF-8200) device.
- **high capacity drive:** a drive with an advertised capacity of greater than 8000 GB for a 3.5" form factor (see SFF-8300) device or a 2.5" form factor (see SFF-8200) device.
- **SCSI Protection Information (PI):** a group of fields at the end of each logical block that contains a logical block guard, an application tag, and a reference tag (see SBC-4).

The size of SCSI Protection Information is defining by the following:

- a) For ATA devices: SCSI Protection Information size = 0 bytes.
- b) For SCSI devices with Protection Information Type 0: SCSI Protection Information size = 0 bytes.
- c) For SCSI devices with SCSI Protection Information Type 1, 2 or 3: SCSI Protection Information size = 8 bytes.

For additional information, including example calculations and calculation tools, see SFF-8447. SFF-8447 is available for download at www.sffcommittee.org. Due to the confusion in the industry and operating systems possibly using a binary calculation to report drive capacity, it is strongly recommended that product specifications clearly define how the SSD capacity is calculated.

C.2 LBA Counts from Advertised Capacity for Low Capacity Drives

C.2.1 LBA Counts from Advertised Capacity for Low Capacity Drives Overview

For the purposes of calculating a low capacity drive's LBA count, the following logical block sizes are encompassed:

- 512 bytes
- 4096 bytes

Drives formatted for SCSI Data Integrity Field (DIF) or SCSI Protection Information (PI) have sector sizes of 520 bytes (512 + 8) or 4104 bytes (4096 + 8). The additional eight bytes are defined as protocol overhead, not additional user capacity. An SSD formatted to support DIF or PI shall have the same logical block count for the reported user-addressable capacity as it does when not formatted to support DIF or PI.

The number of LBAs is an integer multiple of 8 for 512 byte block size low capacity drives to maximize the value of devices with physical sectors that each accommodate 8 logical blocks.

C.2.2 LBA Counts from Advertised Capacity for Low Capacity Drives with logical block size of 512 bytes

For a logical block size of 512 bytes:

- User-addressable logical block count = $21168 + (1953504 \times \text{SSD Capacity in Gbytes})$; or
- SSD Capacity in Gbytes = $(\text{User-addressable logical block count} - 21168) / 1953504$.

C.2.3 LBA Counts from Advertised Capacity for Low Capacity Drives with logical block size of 4096 bytes

For a logical block size of 4096 bytes:

- User-addressable logical block count = $2646 + (244188 \times \text{Capacity in Gbytes})$; or
- Capacity in Gbytes = $(\text{User-addressable logical block count} - 2646) / 244188$

C.3 LBA Counts from Advertised Capacity for High Capacity Drives

C.3.1 LBA Counts from Advertised Capacity for High Capacity Drives Overview

For the purposes of calculating a low capacity drive's LBA count, the following logical block sizes are encompassed:

- 512 bytes
- 520 bytes
- 524 bytes
- 528 bytes
- 4096 bytes
- 4160 bytes
- 4192 bytes
- 4224 bytes

A fit adjustment factor (defined as constant value = 0.995) is used for high capacity drives when the block size is not equal to:

- 512 bytes with zero bytes of SCSI Protection Information; or
- 4096 bytes with zero bytes of SCSI Protection Information.

C.3 LBA Counts from Advertised Capacity for High Capacity Drives (cont'd)

C.3.1 LBA Counts from Advertised Capacity for High Capacity Drives Overview (cont'd)

The high capacity drive number of LBAs is an integer multiple of 2^{21} for 5xx and 2^{18} for 4xxx to maximize the value to common space allocation methods such as power-of-two block groups in ext and other file systems, and for shingled magnetic recording (SMR) with Host Aware and Host Managed methodology. These are two examples where fractional capacity between useful granularity points has no value. This granularity provides support for Zoned-device ATA Commands (ZAC) and SCSI zoned block commands (ZBC). The LBA count granularity is a constant.

Informational: Drive capacity is an integer multiple of 1 GiB for the nominal logical block sizes of 512 and 4096 bytes with SCSI Protection Information size of zero bytes.

C.3.2 LBA Counts from Advertised Capacity for 5xx High Capacity Drives

For high capacity drives with logical block size of 512, 520, 524 and 528 bytes the following formulas shall apply:

For logical block size of 512 bytes with SCSI Protection Information size of zero bytes, the following formula applies:

- $\text{LBA count} = \text{ceiling}(\text{advertised capacity}/512, 2^{21})$

For logical block size of 512 bytes with SCSI Protection Information size of eight bytes, and logical block sizes of 520, 524, and 528 bytes, the following formula shall apply:

- $\text{LBA count} = \text{floor}(\text{ceiling}(\text{advertised capacity}/512, 2^{21}) * (512 / (\text{logical block size} + \text{PI size})) * \text{fit adjustment factor}, 2^{21})$

C.3.3 LBA Counts from Advertised Capacity for 4xxx High Capacity Drives

For high capacity drives with logical block sizes of 4096, 4160, 4192, and 4224 bytes the following formulas shall apply:

For logical block size of 4096 bytes with SCSI Protection Information size of zero bytes, the following formula shall apply:

- $\text{LBA count} = \text{ceiling}(\text{advertised capacity}/4096, 2^{18})$

For logical block size of 4096 bytes with SCSI Protection Information size of eight bytes, and logical block sizes of 4160, 4192, and 4224 bytes, the following formula shall apply:

- $\text{LBA count} = \text{floor}(\text{ceiling}(\text{advertised capacity}/4096, 2^{18}) * (4096 / (\text{logical block size} + \text{PI size})) * \text{fit adjustment factor}, 2^{18})$

Annex D (informative): Estimated Retention Time for Nonstandard Active Use temperatures

An SSD whose data retention is verified to the standard class requirements of Table 1 will have different data retention capability if the actual use temperature differs from that listed in the table. This is in particular true for the temperature-accelerated data retention mechanism focused on this standard, which has an Arrhenius acceleration with activation energy of 1.1 eV as documented in JEP122.

Using the model documented in JEP122, the retention capability expected under other use conditions is summarized in the Table D.1.

Table D.1 — Expected retention (weeks) at different use temperatures

Power Off Temperature	55	[Redacted]						8
	50	[Redacted]					9	15
	45	[Redacted]				10	17	27
	40	[Redacted]			14	20	31	52
	35	[Redacted]		20	26	38	61	101
	30	[Redacted]	32	39	52	76	120	199
	25	58	65	79	105	155	244	404
		25	30	35	40	45	50	55

Client

Power Off Temperature	55	[Redacted]						2
	50	[Redacted]					2	4
	45	[Redacted]				2	4	7
	40	[Redacted]			3	4	7	13
	35	[Redacted]		3	5	8	14	25
	30	[Redacted]	4	6	10	16	28	50
	25	7	9	12	20	33	58	101
		25	30	35	40	45	50	55

Enterprise

For example, the standard client use condition is 40 °C active temperature and 30 °C power-off temperature, at which condition the data retention verified by this standard is the required 52 weeks (one year). On the other hand, if the actual power-off retention temperature were 25 °C, then this standard would ensure that the retention would be at least 105 weeks, or a little over two years.

Annex E (informative): Estimation of SSD Endurance Rating

Using the appropriate workload (see JESD-219), the SSD manufacturer may determine the relationship between host writes and NAND cycles, the latter being the number of program/erase cycles applied to any NAND block, and use this relationship to estimate the SSD endurance rating. If the SSD employs more than one type of NAND component with different cycling capabilities, then a separate relationship should be obtained for each type of NAND. If operating the SSD to the desired TBW is impractical because time required would be excessive, then the relationship between NAND cycles and host writes should be extrapolated. In performing the extrapolation, any nonlinearities in SSD operation, such as those resulting from a reduced cycling pool at extended cycles, should be accounted for.

The estimated endurance rating is the TBW such that:

$$f(\text{TBW}) < \text{NAND cycling capability} \quad (1)$$

where $f(\text{TBW})$ expresses maximum NAND cycles as a function of TBW. The relationship may be different for different types of NAND components used in the SSD.

Consider an SSD containing only one type of NAND and no features of the drive design that would make the WAF change over the lifetime of the drive. Suppose further that the design of the wear leveling method is expected to result in the most heavily-cycled erase block receiving twice the average number of cycles. In that case, WAF would be a constant (for a given workload), and

$$f(\text{TBW}) = (\text{TBW} \times 2 \times \text{WAF}) / C$$

where C is the SSD capacity (see section 4) and the factor of two is the guard band for the wear leveling effects. The SSD endurance rating would then become

$$\text{TBW} < (C \times \text{NAND cycling capability}) / (2 \times \text{WAF})$$

In the more general case, WAF may not be a constant. More extensive characterization would be needed to determine $f(\text{TBW})$ in equation (1) before estimating the endurance rating.

The NAND cycling capability is obtained from component qualification data. The WAF may be obtained from SSD data using the specified workload for endurance testing.

Measurement of WAF requires access to information about NAND program/erase cycles which is generally not available to third parties. Under the assumption in this example where WAF is constant, WAF may be measured after operating the SSD long enough to reach a steady state, without needing to operate the drive to its full endurance rating. The guard band for wear leveling effects (two in this example) may be measured from similar SSD data or estimated from the design of the wear leveling scheme.

Annex F (informative) - Differences between JESD218B and JESD218A

This annex briefly describes most of the changes made to entries that appear in this standard, JESD218B, compared to its predecessor, JESD218A (February 2011). If the change to a concept involves any words added or deleted (excluding deletion of accidentally repeated words), it is included. Some punctuation changes are not included.

- Item 303.19: Change the assumption of active endurance life (the time over which an SSD is exercised to its full endurance rating) from 1.0 year to 1.5 years, to conform with the JEDEC component reliability standard (JESD47).
- Item 303.20: Provide more flexibility in the setting of temperature during an endurance verification test, for example, adding new degrees of freedom to compensate for temperature fluctuations that can arise in high-volume evaluations.
- Item 303.23: Define “drive writes per day” as an alternative unit of measurement for SSD endurance. This term is in common usage in the industry today.
- Item 303.26: Provide better guidelines to restrict the length of any time delay after endurance stress is complete and before a retention test is performed, preventing a test from being run under conditions that would be unintentionally lenient.
- Item 303.21: Remove the “ramped” method of endurance verification.
- Item 303.22: Eliminate all of the “extrapolation method” options except for the reduced capacity (or “short stroke”) method.
- Item 303.27: Clarify the definition of “data error”, which is equivalent to defining the dividing line between an uncorrectable error rate failure and a functional failure.
- Item 303.28: Refine the low-temperature data retention flow to accommodate the realities of the SILC retention mechanism in modern NAND
- Item 303.32: Updated section on how to count advertised LBA capacity of large capacity drives, as adopted in SFF-8447, to ensure consistency between HDDs and SSDs.



Standard Improvement Form

JEDEC _____

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 240 South
Arlington, VA 22201-2107

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: _____

