

Toshiba SSD User Configurable Over-provisioning

> 1. Introduction

Solid state drive (SSD) technology is arguably the most exciting development to emerge in the storage industry in recent years. With amazingly fast data transfer speeds and lightning quick random access times, SSDs have transformed the worlds of computing and enterprise storage. New applications for SSDs are constantly emerging and adoption continues to grow dramatically.

The continuous evolution of NAND flash, the dominant non-volatile media used in SSDs, has fueled the rise of SSD. Ironically, NAND architecture presents many challenges to overcome when used in a storage device. It is a well-known trend that with each shrink of manufacturing process geometry to increase bit density, certain characteristics of NAND flash media have actually degraded. The keys to overcoming the limitations of NAND flash memory have been advances in controller technology and data management algorithms. These innovations have allowed SSDs to achieve the high performance and high reliability requirements demanded by many OEM applications.

Over-provisioning (OP) an SSD by reserving extra storage space is one such advanced method used to improve SSD performance and endurance. This paper describes OP, and how a user can configure a Toshiba SSD with OP, according to their needs.

> 2. Background

The most challenging aspect of an SSD controller's job is data management on the back end, due to the complexities of NAND flash architecture. The minimum write unit, called a "page", is typically a few kilobytes (KB) in size. The minimum erase unit, called a "block", contains several pages. Before data can be written to a page, the entire block in which the page resides must be erased, including other pages with valid static data. In the case of random write transactions smaller than a page, the controller must copy the valid existing data within a block, as well as the targeted page, and rewrite to another empty block, causing a mismatch between the amount of data writes the host sees, and the actual amount of data writes the NAND sees. This mismatch is called write amplification, which is a major factor that affects the operating life of the SSD.

Also contributing to write amplification, the controller is often working in the background to wear-level, pre-erase blocks, and perform garbage collection. All of these

activities require reserved storage space for the SSD controller to use. During SSD operation, the controller can use any "empty blocks" in the drive for this purpose, but as the drive is filled with user data, fewer empty blocks are available. In this case, the controller must work harder and spend more time actively moving data around, which can reduce performance.

> 3. Description of OP

All SSDs must have a reserved "spare area" for the controller to use for data management purposes. Typically, the minimum reserved spare area in an SSD is the difference between the capacity conventions in storage devices versus memory components. The storage world uses decimal capacities so that one gigabyte (GB) is equal to 10⁹ bytes (1,000,000,000) bytes. On the other hand, the memory world uses binary capacities such that one gigabyte is equal to 2³⁰ (1,073,741,824) bytes (also referred to as a gibibyte, GiB). For example, an SSD listed at a capacity of 128GB will have at least 128,000,000,000 user accessible bytes, although it may have 137,438,953,472 bytes (128GiB) of NAND components on board. This SSD is considered to be a "fully provisioned" or "0% over-provisioned". Therefore, every SSD should have at least 7.37% of reserved spare area in the NAND, and to over-provision is to reserve some additional amount of storage area over this minimum. OP is calculated as follows:

$$OP \% = 100 * (Physical\ capacity - User\ capacity) / User\ capacity$$

where *Physical capacity* = Total amount of NAND flash memory in SSD

User capacity = Amount of NAND flash memory in SSD available to the user

Some SSDs are pre-configured by the manufacturer with some level of OP. Table 1 shows typical SSD capacity configurations and corresponding OP amount.

Table 1. Typical OP Amounts

Advertised Capacity (GB)	Physical Capacity (GB)	User Bytes (GB)	User LBAs	OP %
128	128	128,035,676,160	250,069,680	0%
120	128	120,034,123,776	234,441,648	7%
100	128	100,030,242,816	195,371,568	28%
80	128	80,026,361,856	156,301,488	60%
64	128	64,023,257,088	125,045,424	100%

Typical user LBA amounts based on IDEMA capacity formula from IDEMA Document LBA1-03, LBA Count for Disk Drives Standard.

So why use OP? Why would someone sacrifice valuable user capacity in an SSD? As it turns out, OP can extend drive usable life as well as improve SSD random IOPS performance. Since OP increases the amount of reserved spare blocks available for data management activities, the controller can reduce the amount of background operations and be more efficient when writing and managing data. This improved efficiency lowers the write amplification, which can result in higher endurance as well as increased random IOPS performance. Table 2 shows an example of how the endurance (specified in drive writes per day or DWPD) and random write performance in an SSD can be increased through over-provisioning.

Table 2. OP % vs. DWPD for Toshiba HK3R Enterprise SATA SSD

User Capacity (GB)	OP%	DWPD for 5yrs*	4KB Random Write IOPS**
480 GB (Default)	7%	1	16,000
400 GB	28%	3	34,000
320 GB	60%	4	47,000
256 GB	100%	5	59,000

*DWPD based on 4KiB 100% random write workload, 4KiB aligned, **Based on internal Toshiba testing. DWPD and random write IOPS performance for non-default over-provisioned capacities are for example purposes only, and do not constitute a guarantee or warranty. SSD endurance and performance will vary depending upon environment and workload to the drive.

Please contact your local Toshiba Customer Support Engineer for more information about how OP can be used with your specific Toshiba SSD product.

> 4. User Configurable OP

Toshiba services the client and enterprise market with a broad range of SSD products, offered with common OP settings pre-configured, depending upon product family. In addition, all SATA Toshiba SSD products allow the user to utilize standard ATA commands to increase OP from the factory default settings to suit the needs of their application. Current products do not allow decreasing the OP from factory defaults.

Table 3. Toshiba SSD Product Families and Default OP %

SSD Product Series	Product Class	Default OP	DWPD for 5 years
HK3R	Main Stream SATA	7%	1*
HG6	Read Intensive SATA	0%	< 1**

*DWPD based on 4KiB 100% random write workload, 4KiB aligned, **DWPD based on JESD219A Client Workload

> 5. Summary of the steps to over provision a Toshiba SSD

1. Prepare SSD for Secure Erase Operation
2. Perform Secure Erase of SSD
3. Issue command SET MAX ADDRESS (F9h) to set new maximum user area in number of LBAs

In a Windows® environment, users can use the free HDDEraser DOS tool to Secure Erase the drive and the Microsoft® DiskPart tool (included with Windows 7 installations) for partitioning. Various off-the-shelf third party utilities are also available for Secure Erase and Partitioning a drive. For example, Parted Magic and EaseUS® Partition Master are useful and freely downloadable tools².

When using a system with a Linux OS, the following steps are used:

1. Check Security status

```
[root@localhost ~]# hdparm -I /dev/sdb
Security:
    Master password revision code =
    65534
    supported
    not enabled
    not locked
    not frozen
    not expired: security count
    supported: enhanced erase
    2min for SECURITY ERASE UNIT. 2min
for ENHANCED SECURITY ERASE UNIT.
```

2. Enable Security features by setting user password

```
[root@localhost ~]# hdparm --user-master
u --security-set-pass <password> /dev/
sdb
[root@localhost ~]# hdparm -I /dev/sdb
Security:
    Master password revision code =
    65534
    supported
    enabled
    not locked
    not frozen
    not expired: security count
    supported: enhanced erase
    Security level high
    2min for SECURITY ERASE UNIT. 2min
for ENHANCED SECURITY ERASE UNIT.
```

3. Perform Secure Erase

```
[root@localhost ~]# time hdparm --user-
master u --security-erase <password> /
dev/sdb
security_password="<password>"
/dev/sdb:
    Issuing SECURITY_ERASE command,
password="<password>", user=user
real 0m0.550s
user 0m0.000s
sys 0m0.000s
```

4. Check Security status

```
[root@localhost ~]# hdparm -I /dev/sdb
Security:
    Master password revision code =
65534
        supported
    not    enabled
    not    locked
    not    frozen
    not    expired: security count
        supported: enhanced erase
    2min for SECURITY ERASE UNIT. 2min
for ENHANCED SECURITY ERASE UNIT.
```

5. Check Current Drive Capacity

```
root@localhost ~]# hdparm -N /dev/sdb
/dev/sdb:
    max sectors = 234441648/234441648, HPA
is disabled
[root@localhost ~]# hdparm -Np117231408
--yes-i-know-what-i-am-doing /dev/sdb
/dev/sdb:
    setting max visible sectors to 117231408
(permanent)
    max sectors = 117231408/234441648, HPA
is enabled
```

> 6. Conclusion

Over-provisioning is one of the essential tools used by Toshiba to maximize the potential of SSDs. Savvy system designers that understand their system workloads can take advantage of user-configurable OP with Toshiba SSDs to tailor the characteristics of the SSD to the particular needs of the system. Ultimately, utilizing OP can help the designer save cost by simplifying qualifications with multiple configurations of a single Toshiba SSD product.

Please contact your Toshiba technical representative for more information about how you can use OP with a Toshiba SSD.

¹ One Gigabyte (1GB) means $10^9 = 1,000,000,000$ bytes using powers of 10. A computer operating system, however, reports storage capacity using powers of 2 for the definition of $1\text{GB} = 2^{30} = 1,073,741,824$ bytes, and therefore shows less storage capacity. Available storage capacity will also be less if the computer includes one or more pre-installed operating systems, pre-installed software applications, or media content. Actual formatted capacity may vary.

² It is important to note that Toshiba does not recommend or endorse any 3rd party utility, including HDDEraser, DiskPart, Parted Magic, or Easeus Partition Manager. The use of such utilities is entirely at the user's own discretion and risk, and Toshiba is in no way responsible for such use.