

An Innodisk White Paper
March 2015

Garbage Collection & TRIM for SSDs

Data Cleansing through Garbage
Collection, and Increasing Drive Life
and Write Speed by Using TRIM

Revision History

Date	Version	Information
2012.07.09	1.0	First Release
2015.3.5	1.1	New logo instead and products renew

Introduction

This white paper presents a workaround for dealing with a solid-state drive's (SSD) inherent drawback in the way that it reads, writes and deletes data. Garbage collection is the technology used to maintain data consistency and perform routine data cleansing on SSDs. To further increase an SSD's speed and lifespan, the TRIM command enables the SSD to handle garbage collection overhead, significantly reducing write operations and increasing the speed and lifespan of the drive.



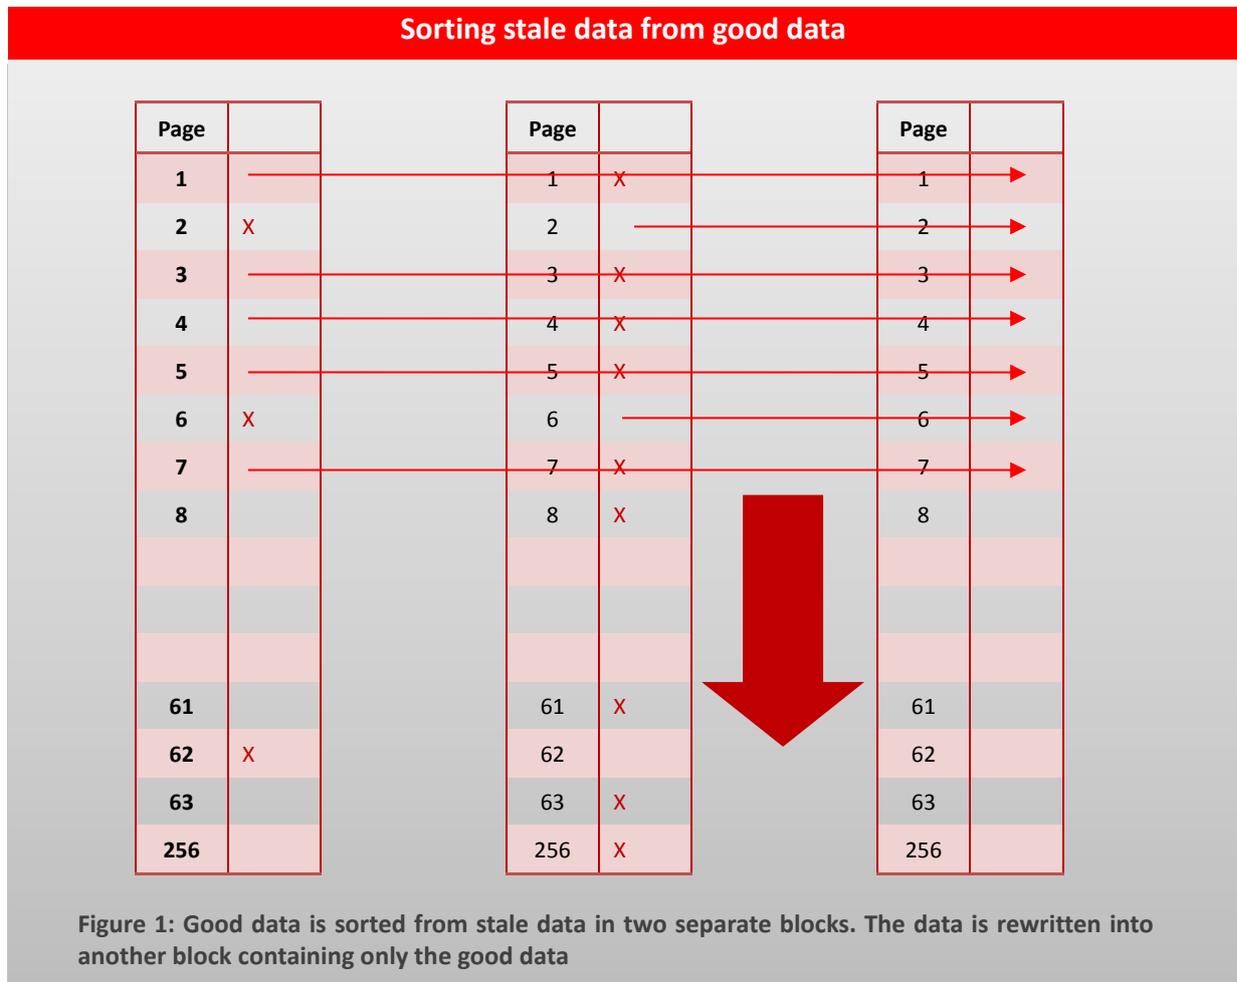
Non-volatile Flash memory is increasingly used in harsh environments like factories where severe vibration can impact operability. SSDs are more robust than traditional HDDs, offering better energy efficiency, longer lifespans, faster access times, and better read performance. The problem, however, is that Flash memory is a write-once and bulk-erase form of data storage, and requires a garbage collection mechanism that sorts good data into available blocks. This ensures efficient disk usage but can significantly reduce performance if not managed correctly. Innodisk employs a linear-access optimized garbage collection along with the TRIM command to guarantee optimal read/write operations and data integrity.

How Innodisk manages garbage collection with TRIM

As stated above, data is written to the Flash memory in units called pages. Pages make up blocks and data can only be erased in entire block-units. Garbage collection involves removing stale data from pages and rewriting the data into newly erased blocks. See Figure 1. When the free block is empty or the system is idle, the SSD will perform garbage collection on the block.

The process of garbage collection involves reading and rewriting data to the Flash memory. This means that a new write from the host will first require a read of the whole block, a

write of the parts of the block which still include valid data, and then a write of the new data. This can significantly reduce the performance of the system. This is where the TRIM command comes in.



TRIM enables the SSD controller to skip invalid data instead of moving it. Naturally, this frees up a significant amount of resources and extends the lifespan of an SSD by reducing erase and write cycles on the SSD. TRIM tells the controller not to waste resources performing garbage collection on data in its logical block addresses (LBA) that has been designated invalid. See Figure 2.

Garbage Collection with the TRIM command

	1. User writes four new files	2. User deletes file "C" and OS sends TRIM	3. User writes new file "E"
OS Logical View	<div style="display: flex; justify-content: space-around; font-size: small;"> File A File B</div> <div style="display: flex; justify-content: space-around; font-size: small;"> File C File D</div> <div style="display: flex; justify-content: space-around; font-size: small;"> Free</div>	<div style="display: flex; justify-content: space-around; font-size: small;"> File A File B</div> <div style="display: flex; justify-content: space-around; font-size: small;"> File D</div> <div style="display: flex; justify-content: space-around; font-size: small;"> Free Free</div>	<div style="display: flex; justify-content: space-around; font-size: small;"> File A File B</div> <div style="display: flex; justify-content: space-around; font-size: small;"> File D File E</div> <div style="display: flex; justify-content: space-around; font-size: small;"> Free</div>
SSD Logical View (LBAs)	<div style="display: flex; justify-content: space-around; font-size: x-small;"> A1 A2 A3 B1</div> <div style="display: flex; justify-content: space-around; font-size: x-small;"> B2 B3 B4 B5</div> <div style="display: flex; justify-content: space-around; font-size: x-small;"> B6 C1 C2 D1</div>	<div style="display: flex; justify-content: space-around; font-size: x-small;"> A1 A2 A3 B1</div> <div style="display: flex; justify-content: space-around; font-size: x-small;"> B2 B3 B4 B5</div> <div style="display: flex; justify-content: space-around; font-size: x-small;"> B6 D1</div>	<div style="display: flex; justify-content: space-around; font-size: x-small;"> A1 A2 A3 B1</div> <div style="display: flex; justify-content: space-around; font-size: x-small;"> B2 B3 B4 B5</div> <div style="display: flex; justify-content: space-around; font-size: x-small;"> B6 E1 E2 D1</div>
SSD Physical View	<div style="display: flex; justify-content: space-around; font-size: x-small;"> A1 A2 A3 B1</div> <div style="display: flex; justify-content: space-around; font-size: x-small;"> B2 B3 B4 B5</div> <div style="display: flex; justify-content: space-around; font-size: x-small;"> B6 C1 C2 D1</div> <div style="background-color: black; width: 100%; height: 10px; margin-top: 5px;">Over Provisioning</div>	<div style="display: flex; justify-content: space-around; font-size: x-small;"> A1 A2 A3 B1</div> <div style="display: flex; justify-content: space-around; font-size: x-small;"> B2 B3 B4 B5</div> <div style="display: flex; justify-content: space-around; font-size: x-small;"> B6 GC GC D1</div>	<div style="display: flex; justify-content: space-around; font-size: x-small;"> A1 A2 A3 B1</div> <div style="display: flex; justify-content: space-around; font-size: x-small;"> B2 B3 B4 B5</div> <div style="display: flex; justify-content: space-around; font-size: x-small;"> B6 GC GC D1</div> <div style="display: flex; justify-content: space-around; font-size: x-small;"> E1 E2 </div>
	SSD writes new data; only SSD knows about OP	TRIM from OS tells SSD to ignore the data in the location previously holding file "C" during GC	OS writes new file to old location; SSD writes file E to another free area

Figure 2: Garbage collection with the TRIM command

Source: thesdreview.com

In the figure above, when file C in column 2 is erased, the old file C space is immediately marked as invalid. During garbage collection this free file C spaced is used to maximize system performance. The controller does not need to move the invalid data to a free block.

Conclusion

While SSD storage solutions offer substantial value added benefits over traditional HDD data storage, data integrity remains a constant challenge for engineers and industry experts. Innodisk’s handling of garbage collection along with the TRIM command improves write performance on SSDs. Garbage collection with TRIM eliminates the need for whole block data erasing prior to every write operation, and helps prevent performance degradation and increase SSD life spans.

Innodisk SSD series that incorporate Garbage Collection with TRIM

All Innodisk SSDs incorporate Garbage Collection. The following series– 3ME3,3MG-P, 3MG2-P, 3MR-P, 1MR-P, 1SR-P, 3SE-P and 3SR-P series – incorporate Garbage Collection with TRIM.

About us

Innodisk is a worldwide leading provider of data storage and memory module solutions for industrial and mission-critical applications. Leveraging in-house engineering and R&D expertise with a keen insight on industry trends, Innodisk's solid-state drive (SSD) technologies provide enhanced, vertically-integrated data storage solutions. Our advanced Flash-based data storage and DRAM memory solutions meet stringent aerospace and defense application requirements, and are also widely used in industrial applications and embedded systems. Innodisk offers customized solutions, from unique form factors to special firmware designs, and our support team of hardware, software and firmware engineers is always ready to tailor the right solution to each customer's needs. Innodisk continually strives to innovate and provide system integrators and end customers with the best service in the industry.

For more information about Innodisk's product line, technologies and applications, please visit www.Innodisk.com