Data Compression in the Intel® Solid-State Drive 520 Series

On-board data compression engine helps increase performance and endurance in the Intel® SSD 520 Series.

Utilizing a unique hardware and firmware architecture, the Intel® Solid-State Drive 520 Series (Intel® SSD 520 Series) implements on-board data compression, a feature that helps increase performance and endurance by automatically compressing data sent to the SSD.

What is Data Compression?

Data compression in a storage drive is the process of encoding information in such a way that the resulting data requires less space to be stored on the NAND (the storage component of the Intel SSD 520 Series). The amount of data that can be compressed depends on the type of data.

In the Intel SSD 520 Series, data compression occurs automatically. The SSD uses *lossless data compression* to ensure the original data is preserved and completely reconstructed when accessed.

Figure 1. Data Compression

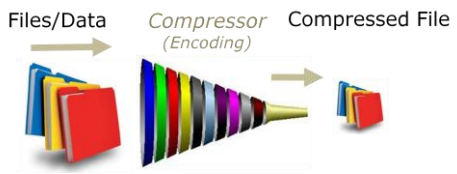

Advantages of Data Compression

Writing to the NAND requires time and is usually the bottleneck of SSD performance. With data compression, less data is written to the NAND, and therefore, effective write speed improves – resulting in better performance.

Additionally, because fewer bits are written to the NAND as compared to an SSD that does not use data compression, the NAND is exercised less. This increases the relative endurance of the SSD, as more lifetime host writes can be performed before the NAND wears out.

What Data is Compressible?

Whether or not a file can be compressed depends on the file type. Various types of files have redundancy in the pattern of data and can be compressed to some degree. Typically, files that can be compressed include system files relating to the operating system, application files and associated user files, and PC utilities and games. Files that typically cannot be compressed include software-encrypted files and image files such as pictures, videos, and DVD movies.

A study of a sample of office users shows a significant amount of data in user drives is compressible. The study shows a typical office user writes Microsoft Outlook* files and documents such as Microsoft Word*, Microsoft PowerPoint*, and Microsoft OneNote* more than 40% of the time as compared to other file types, as shown in Figure 2. Security software (such as virus scan), Internet usage, and instant messaging (IM)-type software such as Microsoft Communicator* comprise approximately another 30% of the activity.

Figure 2. Drive Usage of a Typical Office User[1]

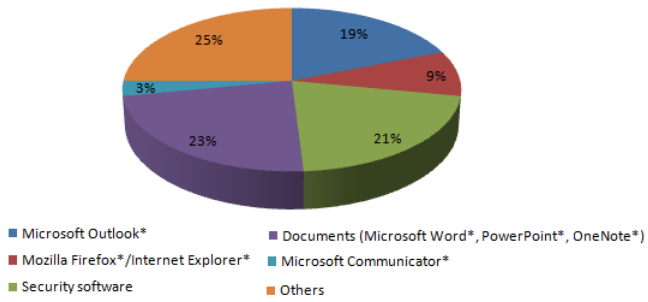

1. Source: Intel labs data.

A closer look at the compression rates of these different file types shows that 75% of the file types observed can be typically compressed 60% or more.

Table 1. Typical Compression Rates[1]

File Type	Typical Compression
Microsoft Outlook*	60%
Microsoft documents	85%
Mozilla Firefox* or Internet Explorer*	80%
Microsoft Communicator*	70%
Security software	95%

1. Source: Intel labs data. Data shown is for illustration purposes only and is not a true representation of the Intel SSD 520 Series.

Improved Drive Writes

Data compression decreases the number of writes to the NAND on an SSD.

For example, Figure 3 shows that with data compression, an installation of Microsoft Windows* 7 and Microsoft Office 2010 results in only about half the number of writes to the NAND as host data was sent.

Figure 3. Compression Benefits with Amount of NAND Data Written[2]

2. Source: Intel labs data.

This decrease in NAND writes directly affects the lifetime drive writes the SSD can support. Additionally, reducing the amount of bits stored on the NAND in general reduces the amount of data clutter on the SSD and improves drive efficiency in maintaining its health through *garbage collection* (the SSD controller process of cleaning up and managing NAND for optimal health).

Improved Performance

With a significant amount of user data being compressible, performance benefits can be substantial. Table 2 shows performance benefits of data compression with the 240 GB Intel SSD 520 Series.

Table 2. Intel SSD 520 Series Performance[3]

Performance Specification	Incompressible Data	Compressible Data
Sequential Write Bandwidth (Mbp/s)	235	520
Sequential Read Bandwidth (Mbp/s)	550	550
Random Write (IOPS)	16,500	60,000
Random Read (IOPS)	46,000	50,000

3. Source: *Intel® Solid-State Drive 520 Series Product Specification*; Random reads based on 4KB Queue Depth 32

Sequential and random write performance is improved in the Intel SSD 520 Series with data compression. In addition, random read performance – where read operations can be bottlenecked due to controller limitations – shows improvement.

Conclusion

Typical user data is compressible, and the Intel SSD 520 Series is the right drive to take advantage of optimizing performance and endurance with data compression.

Solid-State Computing Starts with Intel Inside.® For more information, visit **www.intel.com/go/ssd**

Printed in USA

Please Recycle 326483 -001US