Host Memory Buffer
White Paper

December 30, 2022

Version 1.0

Table of Contents

1. Introduction

The design of a modern SSD has become a battle over real estate. Space is at a premium, and mounting more NAND Flash memory is desirable because of the increase in performance, but this is only possible if other components are removed. That said, some manufacturers are happy to remove components if it means cost savings.

One component that is often selected for removal is DRAM. SSDs that are manufactured without DRAM are called, somewhat unsurprisingly, DRAM-less SSDs. However, if DRAM is removed, what about the functionality it provided? How will that be replaced? Engineers developing the NVMe standard came up with a solution. It's known as the Host Memory Buffer, or HMB. However, the HMB was found to have advantages in SSDs that do have DRAM as well.

This white paper will explain the technology and operation of the Host Memory Buffer as it is implemented in NVMe-compliant Apacer SSDs, as well as the advantages it offers.

2. How It Works

2.1 SSDs without HMB

Since the specification that covers HMB was only created in the NVMe standard, many older drives do not support this technology. It will be beneficial to examine how an older SSD works before introducing the concept of HMB and its advantages. The following chart illustrates the communication process between a Host PC and an SSD which does not support HMB.
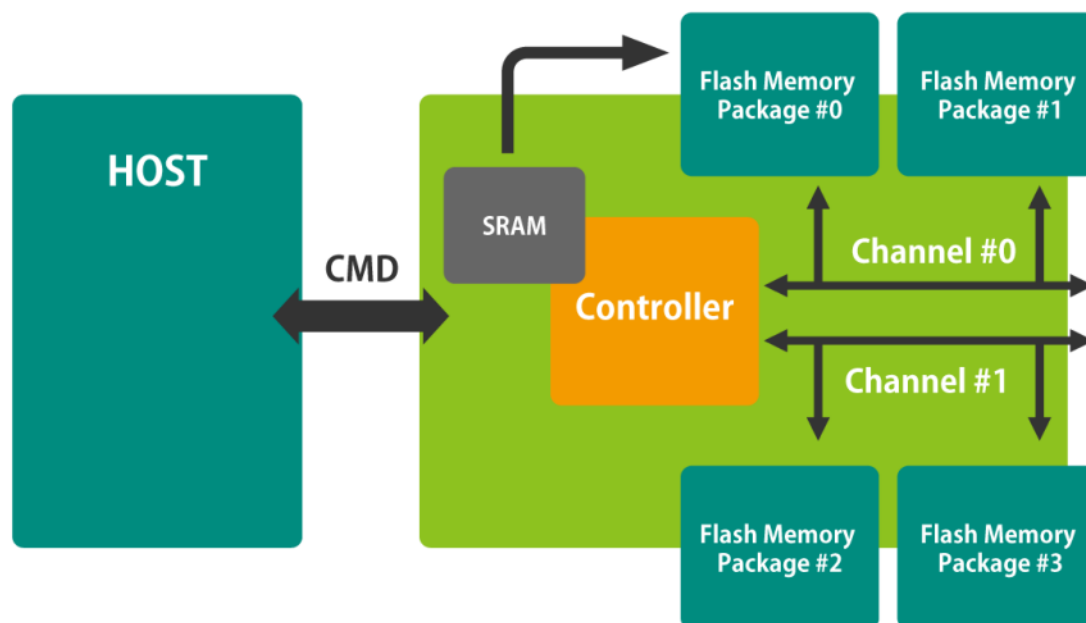
The focal area of this chart is the gray box. The SSD in this example is equipped with SRAM. This small amount of volatile memory performs a valuable function: after the device boots up, the controller copies the mapping table from the Flash Memory Package into the SRAM. Then, as the device is in operation, the mapping table is updated in real time. The mapping table is a small piece of data that tells the SSD where each piece of data is stored in the Flash Memory Package. For example, if the Host asks the controller to retrieve File A, the controller scans the mapping table, and finds that File A is stored in Flash Memory Package #2, so it copies the data from that package and provides it to the host. Later, the Host wishes to save a new file, File B, and sends the file to the controller. The controller scans the mapping table in the SRAM again, and finds there is a free space of suitable size in Flash Memory Package #3, so it saves it there, and it updates the mapping table to reflect this new information.

2.2 SSDs with HMB

If an SSD supports HMB, the method of operation is mostly similar. The following chart illustrates the communication process between a Host PC and an SSD which does support HMB.
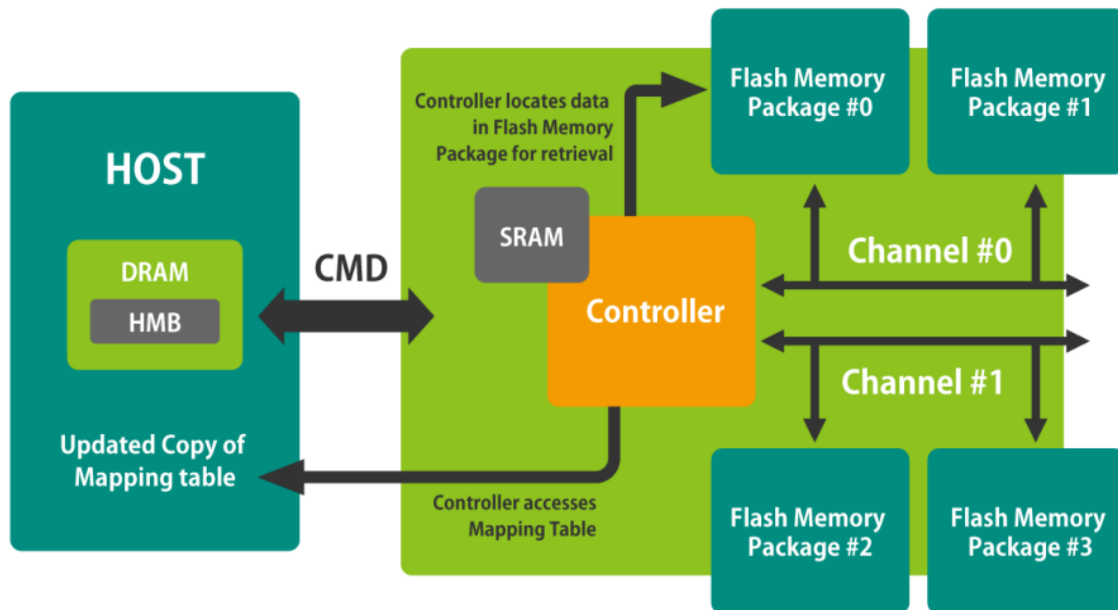
HMB was developed in response to a common trend among modern computers. Even if an SSD had both SRAM and DRAM components, a modern Host computer (especially an industrial computer) would be likely to have much more powerful DRAM modules installed. The change in performance when tasking the SSD's DRAM with maintaining the mapping table was noticeable. The theory was then advanced that using even a small portion of the Host's DRAM for maintaining the mapping table would provide an even greater performance boost. So the specifications were written into the NVMe standard, and SSD manufacturers began to develop HMB-compliant SSDs.

The workflow for Hosts and SSDs when addressing whether or not to use HMB technology is rather simple. Both devices must be compatible with HMB for the technology to work. In some cases, this requires the Host to have a relatively modern operating system installed.

When the Host is powered up and connected to an SSD, it will first examine its own DRAM resources and decide how much size it will allocate to the maintenance of the mapping table. Next, the Host will query the SSD to see if it does or does not support HMB. If it doesn't, the SSD will need to continue to maintain the mapping table as shown in the first chart, using its internal SRAM or DRAM. If it does, the Host will assign a certain range of the available DRAM resources to the SSD for use when running the HMB technology. Then, the SSD will use the assigned range to read and write the mapping table during operation. At certain points, when the SSD is idle, it will still copy the mapping table to a portion of the Flash Memory Package, to guard against the possibility of a sudden unexpected power loss.

2.3 The Advantages of HMB

SSDs which support HMB have certain advantages. The first advantage can be a reduction in price. This is due to the fact that the SSD manufacturer may choose not to install DRAM onto the SSD, creating what is known as a DRAM-less SSD. The SSD will still have SRAM to maintain the mapping table if it is connected to a Host which doesn't support HMB. And if it's connected to a Host which does support HMB, then the performance should be the same or better than if the SSD had internal DRAM anyways.

So the cost of installing a DRAM module on an SSD is saved, but also a small but valuable amount of real estate on the SSD is now available, which can be used for some other application. Since modern SSDs are already very small and very crowded with components, this may be considered as a second advantage of HMB.

The third and final advantage of HMB is in certain Host behavior. Specifically, during random read operations, the load on the Host's DRAM may be higher, but the Host's DRAM may be better equipped to handle the greater load than internal DRAM on the SSD would be. However, the amount in which performance may be increased is also dependent on the SSD's controller. With some controllers, the performance may be markedly improved. With others, it may only be slightly improved or not improved at all.

3. Conclusion

The powerful DRAM modules installed in modern home and industrial PCs have led the SSD industry to respond by adopting HMB technology. Manufacturers of NVMe-compliant SSDs can decide to forgo installing DRAM on their devices to cut costs, optimize real estate on the SSDs, and improve random read performance (as long as controller compatibility is carefully selected). These advantages can then be passed on to the customer, creating a true win-win situation for all.

Revision History

Revision	Description	Date
1.0	Official release	12/30/2022

Global Presence

Taiwan (Headquarters)
Apacer Technology Inc.
1F., No.32, Zhongcheng Rd., Tucheng Dist.,
New Taipei City 236, Taiwan R.O.C.
Tel: 886-2-2267-8000
Fax: 886-2-2267-2261
amtsales@apacer.com

U.S.A.
Apacer Memory America, Inc.
46732 Lakeview Blvd., Fremont, CA 94538
Tel: 1-408-518-8699
Fax: 1-510-249-9551
sa@apacerus.com

Japan
Apacer Technology Corp.
6F, Daiyontamachi Bldg., 2-17-12, Shibaura, Minato-Ku,
Tokyo, 108-0023, Japan
Tel: 81-3-5419-2668
Fax: 81-3-5419-0018
jpservices@apacer.com

Europe
Apacer Technology B.V.
Science Park Eindhoven 5051 5692 EB Son,
The Netherlands
Tel: 31-40-267-0000
Fax: 31-40-290-0686
sales@apacer.nl

China
Apacer Electronic (Shanghai) Co., Ltd
Room D, 22/FL, No.2, Lane 600, JieyunPlaza,
Tianshan RD, Shanghai, 200051, China
Tel: 86-21-6228-9939
Fax: 86-21-6228-9936
sales@apacer.com.cn

India
Apacer Technologies Pvt Ltd,
1874, South End C Cross, 9th Block Jayanagar,
Bangalore-560069, India
Tel: 91-80-4152-9061/62
Fax: 91-80-4170-0215
sales_india@apacer.com