



PHISON

Gaming SSD

A new class of drive
optimized for modern gaming

By Sebastien Jean

Sr. Director, Systems Architecture
Phison Electronics



Gaming SSD

A new class of drive optimized for modern gaming

By Sebastien Jean


Sr. Director, Systems Architecture
Phison Electronics

With the arrival of the latest game consoles, game developers are changing how they approach data management. Historically, games were optimized for slow HDD storage and very little thought was given to optimizing for SSD. We hit a tipping point with 4K textures, which forced the industry to rethink what storage means to gaming. You may have heard of [DirectStorage](#), but that technology only addresses the file system overhead. Changes are also needed on the SSD in order to fully enable modern gaming. The first PCIe Gen 4 SSD was introduced in 2019 and achieved 5 GB/s. The latest iteration can saturate the host interface at 7.4 GB/s. In this article, we will explain how games are changing to make use of this bandwidth.

What has changed that makes Gen 4 worthwhile?

Is Gen 4 storage just a shameless ploy to bilk more cash out of gamers with bigger spec numbers that have no value? As it turns out, that's not the case. The entire gaming ecosystem has been realigning, with a seismic shift in how game developers view storage and system resources.

The change started with 4K video, which greatly increased the size of object texture files, and was further accelerated by [ray tracing](#), which enables highly detailed lighting. Suddenly, games were pushing the limits of PCIe Gen 3 x16 bandwidth. Though PCIe can support 64x or even 128x lanes going to a single device, adding more lanes has a direct impact on CPU silicon area and board design. More PCIe lanes means more cost and increases the technical challenges, thus limiting the number of companies that can participate in gaming. Instead of creating an exclusive club, the industry chose to accelerate their plans for Gen 4, which doubles the bandwidth of each Gen 3 lane from approximately 1 GB/s to 2 GB/s.



Optimizing graphics for **human vision**

Cost and power are always at the front of any technology product-planning roadmap. To manage both critical factors, GPU vendors took a deeper look at how human vision works. Everyone knows that objects that are far away tend to be less clear. This means games can use lower-resolution textures for distant objects, leaving more resources available for near objects while also increasing the overall visual realism. However, it's less obvious how peripheral vision works, because everything we look at is always in focus—or is it?

Though our field of vision spans approximately 120 degrees, we can only see sharp detail in the first 5 degrees to 7 degrees on either side of where we focus our gaze. Everything farther out in our peripheral vision is fuzzy, but when we shift our focus, that object is now in detail. It's easy to overlook how quickly our vision can shift focus, but taking this into account allows for another resource optimization.

Unlike movies and TV shows, gamers move their point of view around constantly to focus on the action. To complete different objectives, gamers must focus on key game elements in order to navigate, unlock a puzzle, or blast a target. Taking a prioritized approach to rendering allows GPUs to push boundaries of realism while managing product cost and power.

So, we're done right? All we needed was a faster **Gen 4 PCIe bus!**

Prioritized rendering helps reduce bandwidth requirements. We can further improve the bandwidth by introducing a read-optimized file system like DirectStorage. We can get another boost through data compression, though decompression is very CPU intensive. Moving the decompression to the GPU extends the bandwidth benefit through DRAM and the PCIe link to the GPU. The catch is that games are constantly getting bigger. Putting all that data into CPU or GPU DRAM greatly increases overall system cost. What if those super-fast Gen 4 SSDs were part of the solution?

While 1 GB of memory costs \$4 to \$5 for DRAM, it only costs \$0.12 to \$0.20 for an SSD. Though an SSD is not a direct replacement for DRAM, the latest Gen 3 SSDs provide 3.5 GB/s of bandwidth and Gen 4 SSDs can now reach 7.4 GB/s. Consoles are pushing the industry to view the SSD more like an L4 or L5 cache that can assist the DRAM. This allows game graphics to continue improving while managing the system thermal load and keeping costs reasonable.



Wait... DRAM, TLC, NAND, Flash... Are those things memory?

Whether we are talking about desktop PCs, laptops, tablets, or smart phones, these devices all have three basic aspects in common: CPU, memory, and storage. These devices typically have 4 GB-64 GB of high-speed DRAM memory that is used when programs are running. The DRAM is considered volatile, because it will lose the data content when power is turned off. The devices also have 128 GB-4 TB SSD storage, which is implemented with nonvolatile NAND flash media. Like older magnetic media, NAND retains its content when power is removed. Flash media was originally implemented using NOR logic, but that technology has been dropped in favor of NAND logic-based flash media. NOR flash is still around for specialized applications, which is why we still use the full name: NAND flash.

NAND flash cells can be configured to store one bit ($2^1 = 2$ voltage levels) per cell for SLC, two bits ($2^2 = 4$ voltage levels) per cell for MLC ("M" = multi), three bits ($2^3 = 8$ voltage levels) per cell for TLC, or four bits ($2^4 = 16$ voltage levels) per cell for QLC ("Q" = quad). The write flow is slower for higher bit configurations because the program operation requires a lot more precision.

The read speed of NAND is roughly equivalent for most applications, which means QLC drives are great for some read-intensive applications. Though gameplay is focused on the read flow, QLC is not currently a good match for this workload. The intense IO activity due to 4K Texture Streaming and Game Patching causes a lot of indirect write activity on the drive.

Program operations stress the NAND cells and wear them down a tiny bit on every write. This is similar to the wear experienced by the brakes on your car. After a few years, they need to be replaced. That's why SSD warranties specify the endurance as max TB Written (TBW). The value is far in excess of what a typical client application will use, but a heavy gamer with frequently patched games will be better served with TLC SSD. It also means that gaming drives will require more reserved capacity to meet TBW targets, so user capacities will likely be 5%-10% lower than a drive tuned for client applications.



Can I use any **Gen 4 SSD** for gaming?

Developing a new AAA game title costs millions of dollars. The storage architecture is decided early in the game's development and becomes a fundamental part of the game. Game developers are not likely to redesign the game for a different platform. In most cases, they would simply reduce the graphic quality which is not what PC gamers want. Consoles provide a much more sophisticated environment, which means PC gaming must also evolve.

Client SSD

Today client SSDs are optimized for a workload that hasn't changed much in 30 years. Though there are variations, and some users deviate from this flow dramatically, most client workloads have short periods of activity with over 80% idle time. The outstanding command queue depth rarely goes above four, and most IOs are either 128K large sequential or 4K random IO.

Due to its burst-y nature, the client SSD benefits from an SLC write cache (SLC = fast write, great endurance), which hides the latency of TLC or QLC writes. Given that the drive is largely idle, there is always time to flush the cache when the user is not interacting with the drive. Most of the SSD maintenance and error recovery flows are designed to wait until idle time, because for a client SSD, that is always the right decision.

Gaming SSD

The Gaming SSD is designed for a very different workload. The duty cycle goes from 80% for a client SSD idle to 100% active on a gaming SSD for eight to twelve hours of continuous gaming. The workloads also change dramatically and now falls into four major use cases: Game Install, Texture Streaming, Game Patch, and Bulk Load.

— Game Install

Game Install is a nice orderly write flow that is generally sequential in nature. Due to the size of the typical game install (60-250 GB), it does not benefit from an SLC write cache. The data is immediately pushed out to TLC/QLC before the entire game is installed. Gaming SSD will reduce drive wear by enabling command hints that allow the SSD to automatically bypass the SLC write cache.

— Texture Streaming

During gameplay, textures are constantly loaded, then dropped and the drive has very little idle time. The pending command queue depth increases to 64-128 on average, but it can jump up to 1000. Games will issue speculative reads and make much greater use of the abort command. The dominant command size is now 64K, which aligns with the typical game asset size even when the data is compressed. Given the new rapid load and drop flow, the same file may be read thousands of times per minute. The typical access window moves from a few GB of disk address space up to several dozen GB. This new access pattern introduces additional stress on the NAND, so the SSD must perform more data refresh operations, which increases the indirect write count on the drive. More internal writes means higher TBW endurance requirements, though these can be offset with internal over provisioning. The lack of idle time means all this extra work has to be interleaved with host IO without creating significant spikes in command completion latency.

— Game Patch

Game patches only update the pieces of each file that have changed, so it turns the nice orderly game install into a giant pile of spaghetti on the NAND. Though a given file may be logically sequential, the write order on the media is now random, and read commands that used to only require one NAND read may now take two or three operations. To address this inefficiency, the SSD introduces a new defrag command that borrows from the semantics of the existing trim command. Instead of indicating the data is not needed, as is the case for trim, the new defrag command tells the SSD to ensure the data on the NAND can be read efficiently.

— Bulk Load

This can be thought of as the classic game or level load sequence. In this flow, the game is loading code and textures to enable gameplay. The IO is primarily focused on large 128K sequential read. Games will now take full advantage of a deeper queue to make best use of the faster SSD bandwidth and ensure the load time is minimized.



What about latency?

Game developers also need tight latency control. Having an occasional read command take a few milliseconds to complete doesn't really impact the client workloads, but it will cause visible texture glitches during gameplay. The Gaming SSD will have to pull in some quality metrics from the enterprise world: Quality of Service (QoS) and IO Consistency.

— Quality of Service (QoS)

This represents the long tail latency and is designed to measure the one-in-a-millionth worst latency, or 99.999. A million may sound like a rare event, but even with 64K commands, a modern Gen 4 SSD that can perform 1.2M 4K IO can still service 115K 64K IO commands before hitting bus saturation. An excessively slow QoS at 99.999 would be visible every 10 seconds of gameplay.

— IO Consistency

This metric defines an upper bound to the QoS measurement. Drives will be required to achieve at least 90% IO consistency. This means that the average latency divided by the 99.999 QoS latency must achieve a ratio of 0.9 or better. An SSD delivering 115K IO per second on 64K commands μ sec time budget, on average, for each command. This means the worst QoS latency cannot exceed 9.5 μ sec to achieve an IO Consistency of 90%. Without an IO Consistency target, some drives can push out the max latency to several hundred millisecond while still maintaining a good average latency. This kind of long tail latency is not acceptable for the new 4K texture Streaming workflow.



What to look for in a gaming SSD?

Initial game titles are expected to target 2400 MB/s, though this will likely increase as developers learn how to work with these new drives.

Look for the **Guaranteed Gaming Bandwidth (GGB)** metric and compare it to game requirements. The workload definition will be established within the gaming industry and published so that everyone is designing to the same target. Some drives will exceed the targets, but gamers can be confident that their gaming experience will not be compromised if their drive meets the targets identified by the game.

Summary

Gaming workloads are very aggressive for the SSD and require a lot of additional internal operations to keep the data error-free. If we compare workloads to cars, the standard client workload is like a station wagon, and the gaming workload is like a high-performance sports car. They can both hit 100 mph, but the gaming-tuned SSD will support a much more aggressive driving style. As we move into late 2021, be sure to keep an eye out for Gaming Class SSD. You'll also want to match the Guaranteed Gaming Bandwidth (GGB) with the game requirements.

Gaming SSD: Class 5000

PCIe Gen-4 x4
Seq: 7.4 GB/s
Random 1.2M IOPS
GGB: 5000 MB/s
IO Consistency: 90%

Gaming SSD: Class 2400

PCIe Gen-3 x4
Seq: 3.5 GB/s
Random 840K IOPS
GGB: 2400 MB/s
IO Consistency: 90%