

PC Storage Performance in the Real World

Table of Contents

| | |
|---|----|
| A Quick Primer on I/O Characteristics | 3 |
| Research Methodology | 4 |
| General Workloads | 6 |
| Gaming Workloads | 9 |
| Content Creation Workloads | 12 |
| Trends in Real-World PC Storage Performance | 15 |
| Evaluating Common Benchmarks | 17 |
| Conclusions & Next Steps | 18 |

Introduction

For PC users trying to determine which solid-state drive (SSD) is the best choice for their specific needs, the job is not always a straightforward one. Price, of course, is a primary concern. Depending on intended usage, other drive characteristics, such as endurance or security features, might take priority.

And then there is the question of performance.

Perhaps more accurately, there are several such questions. How much is enough? How is it measured? Does every activity a user might do – we’ll use the term “workload” in this paper to refer to each individual one – demand the same things of the storage device, or are there many kinds of performance to consider?

The answers, in short: It depends; very carefully; and yes, there is more to the story than the single speed figure that manufacturers often put on the retail box.




Figure 1. The performance number on an SSD retail box usually represents maximum bandwidth, which is rarely seen in real-world usage.

That number is both valid and useful for telling a prospective user something about the drive's performance, but as it turns out, it has relatively little to do with how snappy the system will "feel" in day-to-day performance. We'll get into the reasons for that later.

An enterprising shopper might dig a little deeper by locating the SSD's product specification, or spec sheet (it's common practice for vendors to publish these for each drive). There, they would typically find as many as four performance numbers listed.

Now we're getting somewhere. Having reviewed the spec sheet, we find the so-called "four corners" figures, and consequently understand that "performance" is more than one thing – a drive can read and write data at different speeds, and there's a distinction between sequential and random activity (a function of something called transfer size; more on this in the section "A Quick Primer on I/O Characteristics").

These distinctions are useful for getting a more accurate picture of SSD performance. But based on the Solidigm research outlined in this paper, they're still not enough.

As we will discuss, most real-world workloads comprise a mix of read and write activity, often concurrently. Sometimes the storage device is not the bottleneck at all, meaning user experience could be constrained by another system resource, such as the CPU or GPU. And critically, there is the issue of queue depth (QD) to consider – that is, how many I/O requests are pending at the SSD controller at any given time.

We'll spend a lot of time on queue depth, because it is a critical driver of SSD performance. The higher the QD, the more operations an SSD can run in parallel, resulting in higher throughput – that's why the retail box and spec sheet numbers that SSD vendors use are almost always measured at very high QD (32 is common). But it turns out the vast majority of real-world usage occurs at a QD of just 1 or 2, according to our findings. Consequently, theoretical max performance is not as important as low-QD performance in most cases.

We will also briefly discuss popular SSD benchmarking tools and how well their results align to what we know about real-world performance.

- Capacities: 512 GB, 1 TB, 2 TB
- PCIe 3.0x4 NVMe
- NAND: 144-layer Intel® 3D NAND Technology
- Form Factor: 80mm (single-sided) 2280-S3-M
- Thickness: S3 - up to 2.38 mm
- Weight: <10 grams
- Bandwidth Performance^{1,2}
 - Sequential Read: up to 3,500 MB/s
 - Sequential Write: up to 2,700 MB/s
- Read and Write IOPS^{1,2}
 - Random 4KB Reads: up to 310,000 IOPS
 - Random 4KB Reads: up to 340,000 IOPS
- AES 256-bit Encryption
- Pyrite 2.0

Figure 2. Excerpted from the Solidigm 670p product specification.

A Quick Primer on I/O Characteristics

There is no shortage of ways you could examine the characteristics of a given workload. What’s important for the conversation of SSD performance is identifying the few critical metrics that have the greatest influence on user experience.

For example: Using our internal tool to “trace” workloads as they run and record the underlying system activity (see more in the “Research Methodology” section), we observed what happened when we launched Adobe Acrobat and loaded a PDF file. This common activity resulted in a trace output file with literally thousands of individual data points, despite a workload duration of only about 9 seconds.

Not all of those have a major effect on user experience – but among them are certainly the key drivers. Let’s look at each of these three in depth.

- **Read/write mix** is probably the most intuitive for most users. Depending on what you are doing on your PC, it’s making requests to either put data on the drive or retrieve data that’s already there for use. You can likely guess which kinds of activity would cause one or the other: Launching applications and loading files are read-heavy activities, while saving data results predominantly in writes. Something like a file conversion from one format to another would require both – reads from the source file, and writes to the output file.
- **Sequential-random mix** is another critical characteristic. This refers to how data is written

to or read from the drive: Sequential transfers take place over a longer series of consecutive locations on the drive (called logical block addresses, or LBAs), while random operations tend to be composed of smaller reads or writes to non-consecutive locations. The transfer size is usually on the order of several kilobytes. (Astute readers may notice that some SSD vendors list transfer size with their performance specs – for example, “4KB Random Read” vs. “128KB Sequential Read”). An activity that consists of reading or writing many small files will tend to be more random in nature, while workloads with larger individual files will be more sequential.

- **Queue depth** is less intuitive, but no less important to user experience. Any workload, even a small one, will generate a huge number of I/O operations (in the case of our 9-second Acrobat launch and load example, more than 6,000 of them). The SSD will process those as quickly as possible, but sometimes a queue, or line, forms of waiting requests. Thanks to clever architecture, modern SSDs can service multiple requests in the queue at a time (a concept called “parallelism”). Consequently, performance measured in terms of data throughput will be highest at high QD, when multiple requests are processing at once. But as we will learn later, many typical PC workloads just don’t form long queues because the SSD is responding as quickly as the requests come in.

| Workload Characteristics: Important Drivers of User Experience | | |
|---|--|--|
| Read/Write Mix | How consecutive are the data “chunks” being read or written? | Queue Depth |
| How much data is being written to the SSD during the workload vs. being read from it? | How big are the individual “chunks” of data being read or written? | How many I/O requests is the SSD processing at once? |

These three I/O characteristics form the basis of our analysis in this paper. It’s important to keep in mind that a fourth vector unrelated to storage performance can have a huge bearing on user experience: If the bottleneck is actually occurring elsewhere in the system (for example, the CPU is under full load or the local memory has been exhausted), then better storage performance will not result in better user experience.

Research Methodology

To ascertain what’s going on under the hood during real-world PC usage, we first had to determine two things: which workloads to study, and how to collect and analyze the data.

For workload selection, we opted for a balanced approach that considers three important usage “segments:” general productivity, gaming, and content creation.

Our workload sample is by no means exhaustive, and the results should not be generalized to all real-world PC usage. Instead, we sought to study a collection of activities that were more or less representative of typical day-to-day activities that users in each of these three segments would do. We also wanted to ensure a variety of results were likely, from read- to write-heavy I/O patterns and dealing with a balance of small and large files.

For each workload, we defined “start” and “stop” cues for running the trace, and attempted to ensure that no other activities were running on the system at the same time that might confound the results.

The data collection itself was accomplished using built-in Microsoft Windows features. Microsoft makes available a tool in the Windows Assessment and Deployment Kit (Windows ADK) called the Windows Performance Recorder (WPR), which records Event Tracing for Windows (ETW) events in real-time during workload execution.

| Real-World PC Workloads for Analysis | | |
|--|---|--|
| Productivity | Gaming | Content Creation |
| <ul style="list-style-type: none"> ▪ Launch Zoom and join a virtual meeting ▪ Launch Adobe Acrobat and load a PDF ▪ Launch Microsoft Outlook (with an associated PST file) ▪ Export a Microsoft PowerPoint PPTX to MP4 video | <ul style="list-style-type: none"> ▪ Launch the League of Legends client ▪ Load a new level in Elden Ring ▪ Launch Total War: WARHAMMER III and load a battle while streaming and recording with OBS ▪ Play Grand Theft Auto V for 15 minutes | <ul style="list-style-type: none"> ▪ Launch Adobe Photoshop and load a multi-layer PSD ▪ Launch Blender and load a project file ▪ Import 100 video clips to an Adobe Premiere Pro project ▪ Export a Premiere Pro project to AVI video |

As mentioned earlier, the resulting trace output files contain literally thousands of data points. Among them are:

- How active were various system resources, including the SSD, CPU, GPU, main memory, and network
- Total workload size (in both MB and count of I/O operations), plus read/write mix by size and by count
- QD for every single operation
- Transfer size for every single operation
- Average response time by transfer size
- Data transferred (both size and count) segregated by process
- Individual files that were touched, and how much data was read or written per file
- And much, much more!

We then used an internal tool to organize the trace data into an Excel workbook for analysis and visualization.

In practice, this is what it looked like:

1. Prepare the system for tracing by ensuring a quiescent state (i.e., no other activities running in the background)
2. Start the trace
3. Execute the workload (for example, double-click the Adobe Acrobat icon to launch the program, then load a PDF from the “File” menu)
4. Once the final execution step completes, stop the trace
5. Generate an Excel file with results for analysis

A fair criticism of our approach would be that by pre-selecting which workloads to study, we have biased the aggregate results in a way that renders them less useful for making statements about what typical PC usage looks like “in general.” After all, there are probably many users out there who do none of the 12 things we studied on a regular basis. How useful, then, is a really solid understanding of what’s happening in these specific cases, when it might or might not indicate trends in the larger world of PC usage?

A response in two parts. First: Stay tuned! Solidigm intends to continue this research by moving into another phase of “user session” characterization – in other words, if we just take a look at what a representative group of people does all day and measure that, without asking anyone to run specific workloads, what could it tell us? We anticipate publishing the results of this research in the coming months.

And second: While the aggregate results of the research in this paper might not align precisely to PC usage in the most general sense, we do believe what we’re presenting here is correct directionally – and hence useful as a starting point for this discussion. Some of the results are so overwhelming (for example, 92% of all I/O operations across the 12 workloads occurring at a QD of 1 or 2) that it’s worth considering the implications, even before moving to the broader research phase.

And now, on to the workloads!


General Productivity Workloads

Let's start with a review of the four productivity workloads and discuss our findings in each.

The Acrobat "launch-and-load" workload is a good starting point – a common use case with unsurprising results. The total workload size is approximately 300MB, divided into 88% reads and the remainder writes. The activity is predominantly random with a 4KB transfer size.

While the max QD measured was 191, the geometric mean value was 2.1. (This phenomenon will appear in many of the workloads studied – a sky-high peak QD measured for the tiniest fraction of a second, while an overwhelming majority of operations were at just 1 or 2).

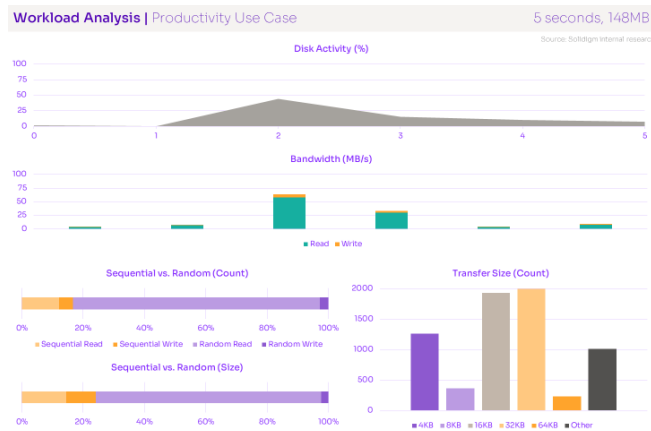
Launch Adobe Acrobat + Load PDF



The second productivity workload tells a similar story. In this one, we launch Microsoft Outlook, which also triggers a load of the associated Outlook Data File (PST) that contains email messages and related items. Again, mostly random read.

Notably, more than 98% of I/O operations during this workload occurred at a QD of 1 or 2.

Launch Microsoft + Load PST



| Read (MB) | Write (MB) |
|-----------|------------|
| 130 | 18 |
| 88% | 12% |

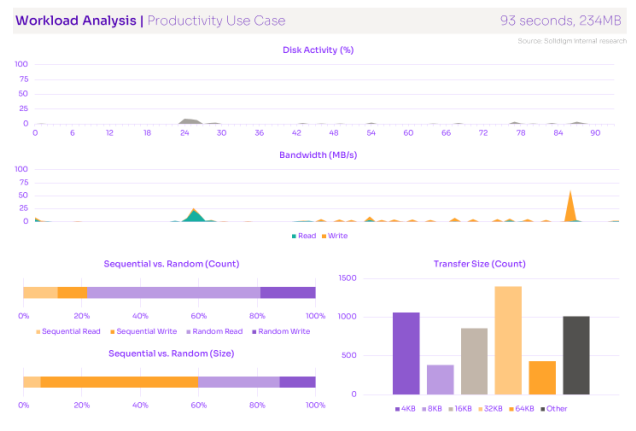
| I/O Attributes | Max | Mean | Notes |
|----------------|--------|-------|-----------|
| Queue Depth | 28 | 1.1 | 98% ≤ QD2 |
| Latency | 53.1ms | 0.2ms | |

| System Resources | Max | Mean |
|-------------------|---------|--------|
| CPU Activity (%) | 16% | 13% |
| Main Memory Usage | 4.2GB | 4.1GB |
| GPU Activity (%) | 3% | 2% |
| GPU Memory Usage | 273MB | 253MB |
| Network Usage | 104KB/s | 22KB/s |

Our first write-heavy workload involves exporting a series of slides from PowerPoint to an MP4 video, a built-in feature of PowerPoint. The process took about 90 seconds.

The bandwidth graph tells a story that is logically consistent with this activity: lots of read activity during the first half, followed by mostly writes as the new video file is generated.

Export Microsoft Powerpoint PPTX to MP4



| Read (MB) | Write (MB) |
|-----------|------------|
| 79 | 155 |
| 34% | 66% |

| I/O Attributes | Max | Mean | Notes |
|----------------|--------|-------|-----------|
| Queue Depth | 20 | 1.1 | 97% ≤ QD2 |
| Latency | 26.0ms | 0.2ms | |

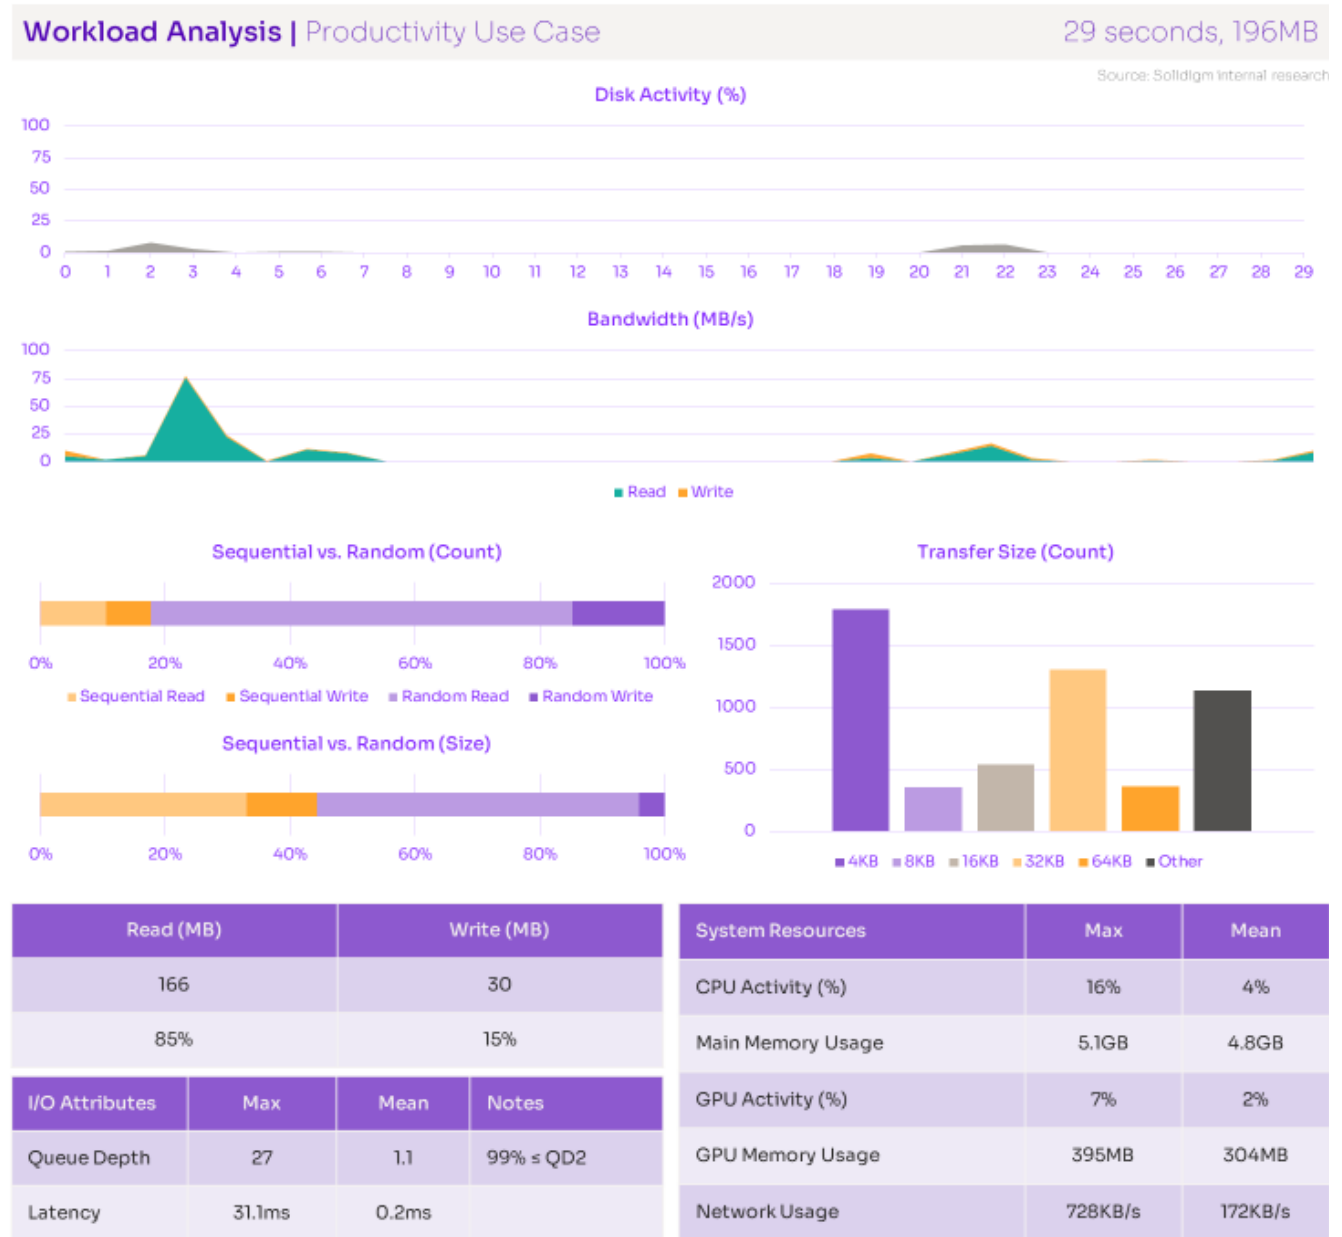
| System Resources | Max | Mean |
|-------------------|---------|--------|
| CPU Activity (%) | 32% | 17% |
| Main Memory Usage | 5.2GB | 4.7GB |
| GPU Activity (%) | 13% | 1% |
| GPU Memory Usage | 602MB | 378MB |
| Network Usage | 567KB/s | 14KB/s |

The final productivity workload represents a use case that is increasingly common since COVID-19 came on the scene: calling into a virtual meeting.

would be the network, which will determine user experience on a call like this more than any other resource.

This activity involves some light I/O traffic, with more than 99% at low QD – but the likeliest bottleneck here

Launch Zoom + join virtual meeting



Gaming Workloads

Gamers are no strangers to waiting for things to load. That’s what makes a thorough analysis of this segment so interesting – there is lots of room for improvement.

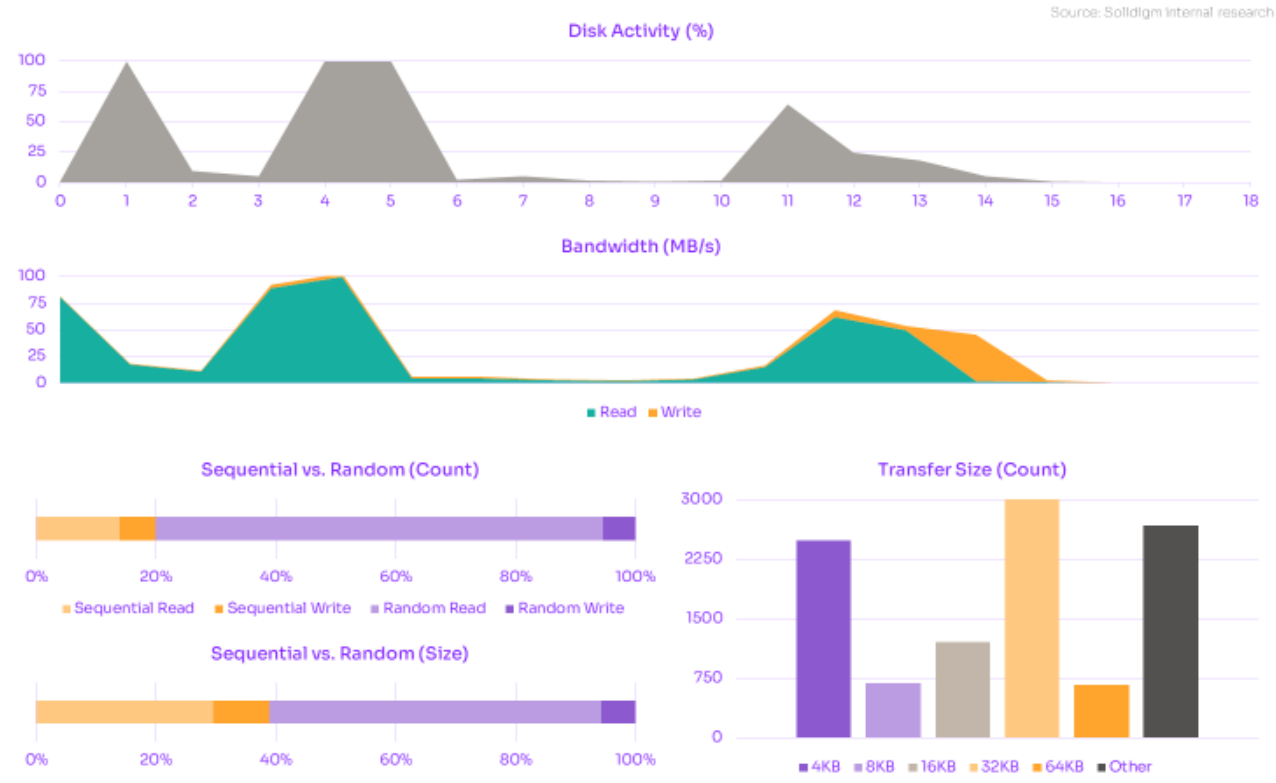
Things like launching games and loading save states and new levels have caused extended wait times for as long as PC gaming has been around. But what’s going on under the hood at times like these? Let’s find out.

The first gaming workload we looked at was relatively light – we simply loaded the League of Legends game client from the desktop. This is just the first step toward gameplay; it only takes the player as far as the menu, so we wanted to examine it in isolation. The entire workload took 18 seconds and involved more than 500MB of data transfer, 85% of which were reads and 80% at low QD.

Launch League of Legends game client

Workload Analysis | Gaming Use Case 18 seconds, 542MB

Source: Solidigm internal research



| Read (MB) | | Write (MB) | |
|-----------|--|------------|--|
| 460 | | 82 | |
| 85% | | 15% | |

| I/O Attributes | Max | Mean | Notes |
|----------------|--------|-------|-----------|
| Queue Depth | 244 | 2.1 | 80% ≤ QD2 |
| Latency | 58.8ms | 0.9ms | |

| System Resources | Max | Mean |
|-------------------|---------|---------|
| CPU Activity (%) | 27% | 13% |
| Main Memory Usage | 5.6GB | 5.1GB |
| GPU Activity (%) | 10% | 2% |
| GPU Memory Usage | 367MB | 298MB |
| Network Usage | 1.6MB/s | 184KB/s |

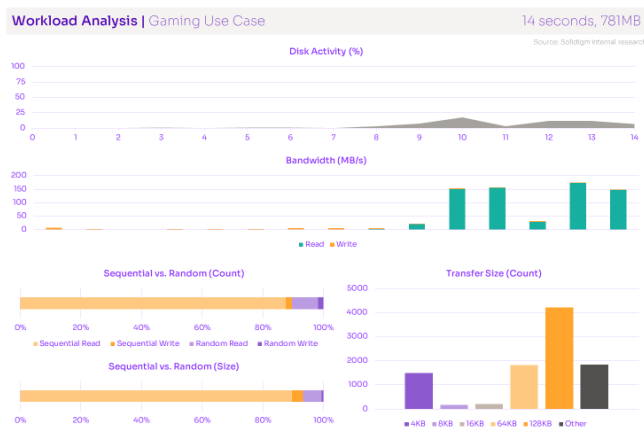
In contrast to the productivity workloads we reviewed earlier, the read activity here is overwhelmingly sequential in nature, with a most-common transfer size of 128KB – this makes sense, given the huge assets (models, textures, etc.) that a modern AAA game employs to create an immersive environment.

Here we also start to see a workload demanding more of the other system resources, particularly main memory and GPU.

This workload involves launching a game and loading a new battle while simultaneously streaming and recording activity using Open Broadcaster Software (OBS). As is typical with AAA gaming workloads, the majority of reads were sequential.

This is also an example of a workload where one of the other system resources, the GPU, reached max utilization – indicating a bottleneck unrelated to storage performance.

Load Elden Ring new level




| Read (MB) | Write (MB) |
|-----------|------------|
| 747 | 34 |
| 96% | 4% |

| I/O Attributes | Max | Mean | Notes |
|----------------|--------|-------|-----------|
| Queue Depth | 4 | 1.2 | 98% ≤ QD2 |
| Latency | 16.2ms | 0.2ms | |

| System Resources | Max | Mean |
|-------------------|--------|--------|
| CPU Activity (%) | 32% | 15% |
| Main Memory Usage | 15.8GB | 14.7GB |
| GPU Activity (%) | 62% | 30% |
| GPU Memory Usage | 3.2GB | 2.3GB |
| Network Usage | 8KB/s | 2KB/s |

Launch Total War: Warhammer III



| Read (MB) | Write (MB) |
|-----------|------------|
| 4,730 | 234 |
| 95% | 5% |

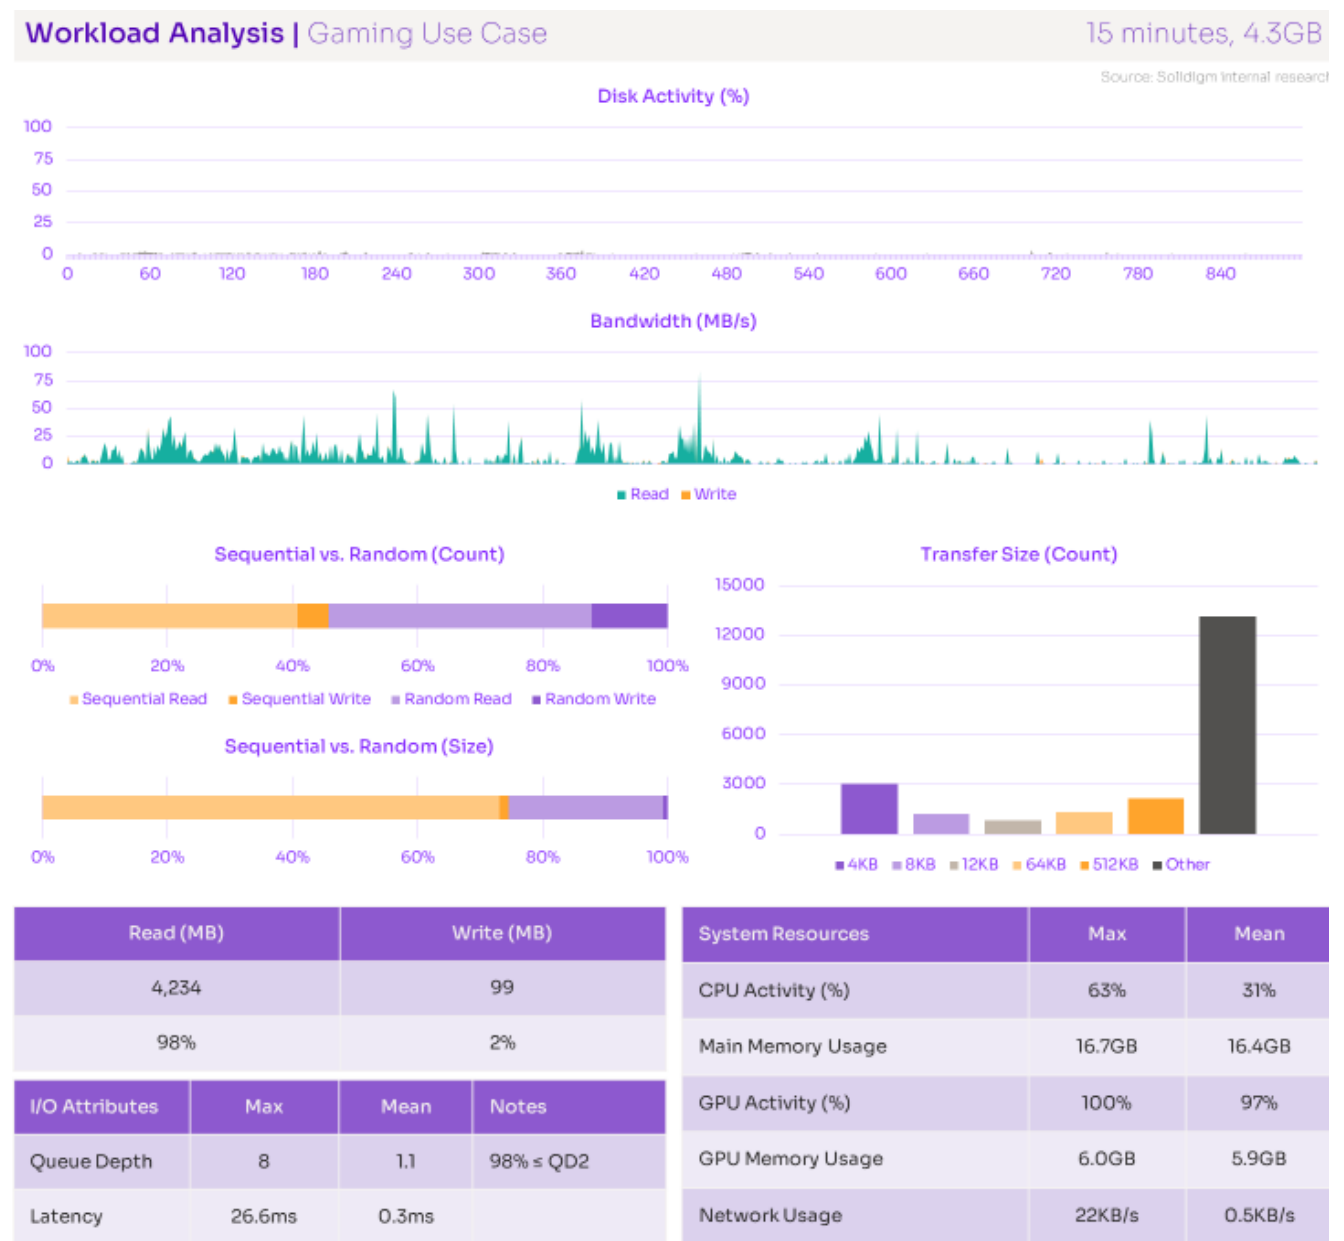
| I/O Attributes | Max | Mean | Notes |
|----------------|--------|-------|-----------|
| Queue Depth | 264 | 1.1 | 97% ≤ QD2 |
| Latency | 98.4ms | 0.9ms | |

| System Resources | Max | Mean |
|-------------------|---------|---------|
| CPU Activity (%) | 92% | 18% |
| Main Memory Usage | 15.3GB | 9.7GB |
| GPU Activity (%) | 100% | 24% |
| GPU Memory Usage | 5.9GB | 2.3GB |
| Network Usage | 601KB/s | 291KB/s |

The final gaming “workload” is a slightly different look at this segment – instead of defining a specific activity, we asked a volunteer to simply play a game for 15 minutes to examine how data might be loaded in as one moves the character around in an open-world environment.

The result was no great spikes in disk activity, but a steady stream of smaller reads that added up to more than 4GB of data transferred over the time window. Those reads were a mix of sequential and random.

Play Grand Theft Auto V for 15 minutes



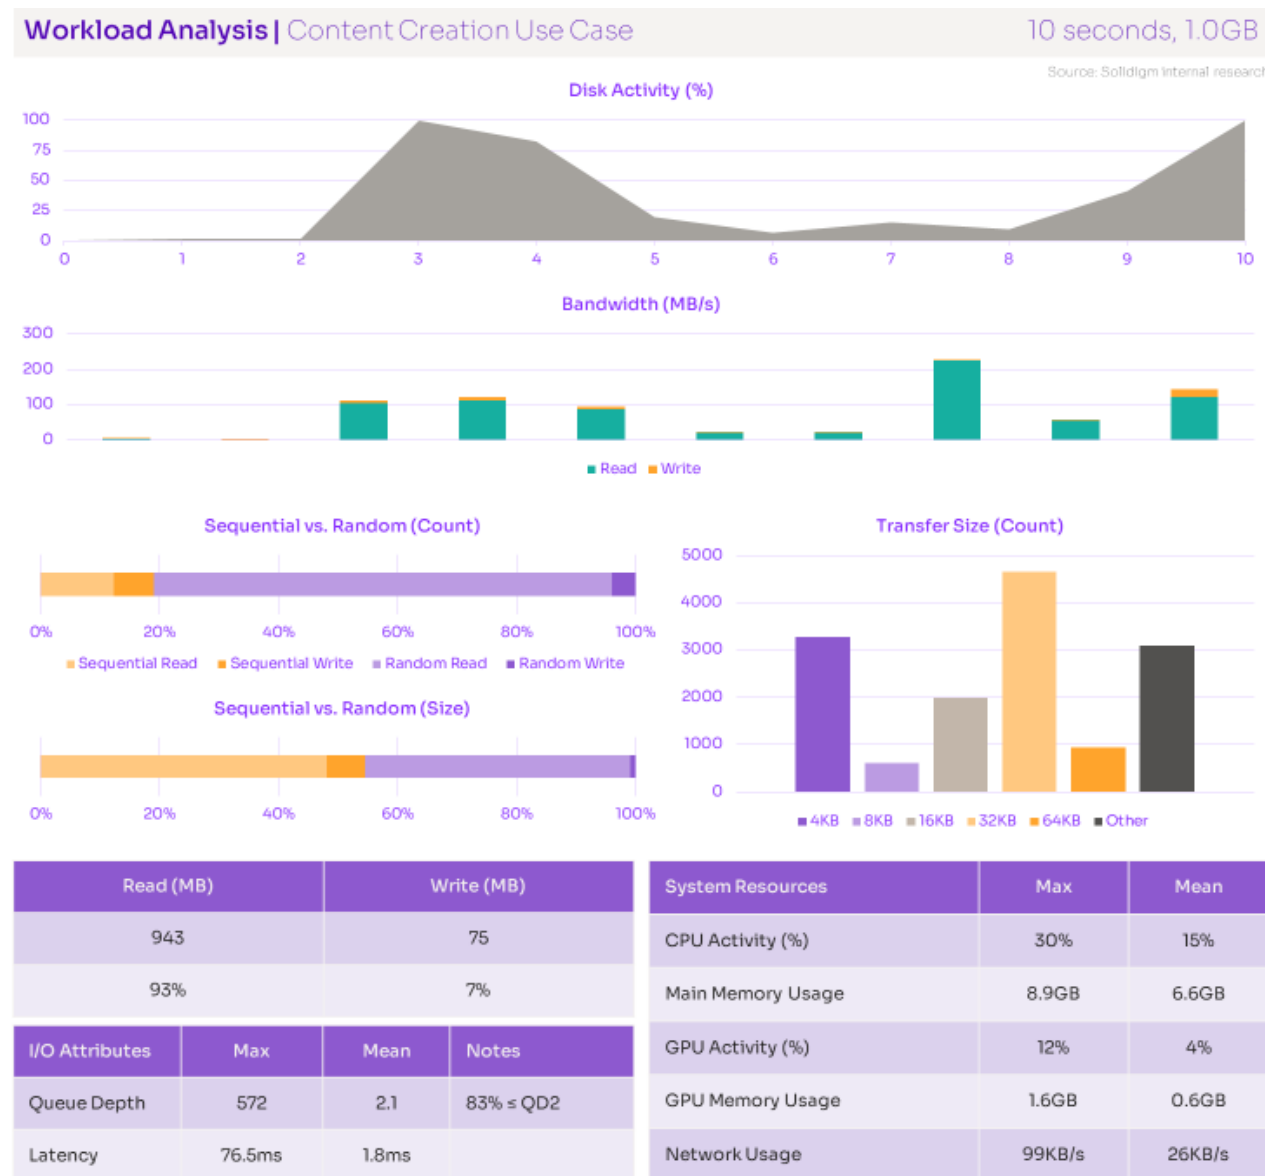
Content Creation Workloads

Finally, let's turn to the content creation segment – a varied category we use to refer to workloads that involve the production of media (video, photo, audio, and so on). Like gaming, this segment can involve some large data transfers and long wait times, depending on the project.

In the first content creation workload, we examine a very common activity: launching Adobe Photoshop and loading a multi-layer PSD (project file).

This 10-second activity involved about 1GB of data transfer; most of the read operations were random, but the smaller proportion of sequential reads nonetheless accounted for about half of the data read from the disk due to the larger size of those operations (mostly 128KB each).

Launch Adobe Photoshop - load PSD




Blender is a free, open-source 3D creation software popular with creators across various industries.

Mean QD creeps all the way up to 2.6 here, but 77 percent of operations remain at QD 1 or 2.

In this particular case, 4KB reads are the predominant I/O type – though depending on the project file used, others might expect to see a greater proportion of sequential transfers.

This activity involved importing many video files from an external drive into an Adobe Premiere Pro project – something a creator might do when creating a family slideshow or editing together a feature film. The bandwidth chart tells the story of many files being read into the program during the first half of the workload, and lots of concurrent read-write activity on the back half.

Launch Blender + load project file

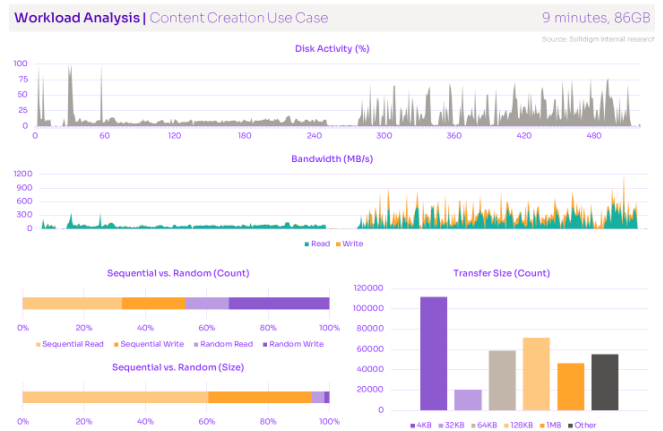


| Read (MB) | Write (MB) |
|-----------|------------|
| 159 | 14 |
| 92% | 8% |

| I/O Attributes | Max | Mean | Notes |
|----------------|--------|-------|-----------|
| Queue Depth | 389 | 2.6 | 77% ≤ QD2 |
| Latency | 20.7ms | 1.5ms | |

| System Resources | Max | Mean |
|-------------------|-------|-------|
| CPU Activity (%) | 14% | 7% |
| Main Memory Usage | 5.3GB | 4.9GB |
| GPU Activity (%) | 4% | 1% |
| GPU Memory Usage | 509MB | 381MB |
| Network Usage | 0KB/s | 0KB/s |

Import 100 video clips to Adobe Premiere Pro



| Read (MB) | Write (MB) |
|-----------|------------|
| 55,405 | 30,235 |
| 65% | 35% |


| I/O Attributes | Max | Mean | Notes |
|----------------|--------|-------|-----------|
| Queue Depth | 1,793 | 1.2 | 97% ≤ QD2 |
| Latency | 54.5ms | 0.9ms | |

| System Resources | Max | Mean |
|-------------------|---------|--------|
| CPU Activity (%) | 33% | 12% |
| Main Memory Usage | 27.0GB | 21.1GB |
| GPU Activity (%) | 62% | 2% |
| GPU Memory Usage | 1.9GB | 1.5GB |
| Network Usage | 131KB/s | 4KB/s |

The final workload in this paper is the largest, by far, in terms of both duration and size. Exporting a large Premiere Pro project file to AVI video took more than two hours and involved moving nearly a half-terabyte of data.

The lion's share of disk activity occurred in the final 20% of the time window, as indicated on the charts. While the process was write heavy (to the tune of 57% of operations), there was a significant proportion of reads involved in the export as well.

Export Adobe Premiere Pro PRRROJ to AVI



Trends in Real-World PC Storage Performance

For PC users trying to determine which solid-state drive (SSD) is the best choice for their specific needs, the job is not always a straightforward one. Price, of course, is a primary concern. Depending on intended usage, other drive characteristics, such as endurance or security features, might take priority.

And then there is the question of performance.

Perhaps more accurately, there are several such questions. How much is enough? How is it measured? Does every activity a user might do – we’ll use the term “workload” in this paper to refer to each individual one – demand the same things of the storage device, or are there many kinds of performance to consider?


The answers, in short: It depends; very carefully; and yes, there is more to the story than the single speed figure that manufacturers often put on the retail box.

Beyond that, the other I/O characteristics were highly workload dependent. While the aggregate view of our workloads suggests a higher proportion of reads than writes in general, and more random than sequential activity – and while we would expect that to be true of general PC usage across all users and workloads – any definitive statement to that effect is pending the further stages of this research project mentioned in “Research Methodology.”

In terms of how sequential or random a workload is, it’s clear that the results are highly workload dependent – although these early results do suggest some distinction by usage segment, with gaming relying on more sequential transfers and productivity on predominantly random activity.


Queue Depth

92% of I/O activity is at QD1 or 2; trivial amounts beyond QD4




Read/Write Mix

R/W mix depends on workload; anecdotally, users spend more time waiting on read activities



Sequential/Random Mix

S/R mix is workload dependent, but productivity is more random while gaming is more sequential




Evaluating Common Benchmarks

With these results in hand, we can have a brief discussion about the usefulness of various popular SSD benchmarks in evaluating or predicting actual end-user experience.

- **ATTO Disk Benchmark** presents a results screen that focuses on SSD read and write performance across various transfer sizes (user configurable). Based on our research, the most important ones to real-world usage are 4KB, 16KB, 32KB, and 128KB.
- **Anvil's Storage Utilities** presents a default test that covers 4KB random transfers at various queue depths (including 4 and 16), as well as sequential operations of 32KB, 128KB, and 4MB at indeterminate QD. It also outputs aggregate read and write “scores” as well as an overall score.
- **AS SSD Benchmark** has a clean, simple interface with a limited number of tests by default. The UI does not specify transfer size for the sequential tests or queue depth for either sequential or random tests.

- **CrystalDiskMark** is popular with users and the tech press for its balance of approachability and thoroughness. The current version, 8, presents a default test set that includes sequential transfers of 1MB at QD 1 and 8, and random 4KB transfers at QD 1 and 32. The tool also includes other test “profiles” that permit users to measure a mix of read and write activity, a useful feature.
- **PCMark 10** includes a set of four storage benchmarks: Full System Drive, Quick System Drive, Data Drive, and Drive Performance Consistency. Unlike the other tools listed here, PCM10 tests rely on real-world traces from common applications to calculate a score. While this approach permits less customization up front – users can’t specify a given queue depth or transfer size to focus on – it has the benefit of aligning well by nature with real-world usage. To better understand this approach, we traced the PCM10 Quick System Drive Benchmark as it ran, with the following results.

PCMark 10 Quick System Drive Benchmark



| I/O Attributes | Max | Mean | Notes |
|----------------|-------|-------|-----------|
| Queue Depth | 115 | 1.2 | 96% ≤ QD2 |
| Latency | 357ms | 0.4ms | |

| System Resources | Max | Mean |
|-------------------|---------|-------|
| CPU Activity (%) | 60% | 7% |
| Main Memory Usage | 7.8GB | 7.3GB |
| GPU Activity (%) | 35% | 0% |
| GPU Memory Usage | 447MB | 407MB |
| Network Usage | 2.3MB/s | 7KB/s |

The PCMark 10 approach aligns fairly well to the aggregate results of our 12-workload sample. Mean QD is 1.2 (with 96% of operations at QD 1 or 2), and most operations are random reads.

Given this alignment, our conclusion is that the results of this test are useful for gauging an SSD's performance on real-world applications – although, as with any test that outputs a calculated score, users should ensure they know how it's calculated before using it as the sole basis for any comparison.

PCMark 10 storage benchmarks accelerate the traces' I/O patterns, eliminating other factors in the source workload, such as compute time and host idle time. This serves to stress the storage system to a greater degree and provides more sensitivity to improvements in storage performance, which is part of what makes the benchmark useful. This difference shows in the Solidigm trace of PCMark 10, where average CPU time and GPU time are 7% and 0%, respectively.

One notable difference is the measured proportion of 4KB transfer sizes in the PCMark 10 workload as compared to the Solidigm-measured application traces. While the Solidigm measurements often have a majority of 4KB transfers, they fall short of the overwhelming majority in PCMark 10. It's difficult to speculate on this difference without further study; keep an eye out for the next Solidigm white paper on real-world performance

Conclusions & Next Steps

- The way an SSD will perform under real-world circumstances is almost impossible to communicate in one number. Even the “four corners” numbers listed in most product specifications only tell part of the story.
- It is important to know what an SSD is capable of at maximum performance, and it is useful for certain real-world activities. But in the workloads we studied here, it was clear that many other factors can and do influence what a user would perceive as storage performance or a generic sense of system “responsiveness.”
- Users can make an informed decision about which SSD is right for them based on their intended use. Our research suggests a few useful rules of thumb.
- Queue depth of 1 or 2 was overwhelmingly represented in real-world I/O operations; consequently, a drive's low-QD performance is worth considering in addition to the “four corners” performance numbers
- The activities studied here indicated a mix of read and write activity; tests that focus on mixed activity should be used alongside ones that stress pure operations of one type or the other
- While there is some apparent differentiation by segment, both sequential and random operations were present in each of the workloads studied. Consequently, it's useful to understand which transfer sizes are prevalent in real-world usage and evaluate SSDs accordingly
- This publication represents Solidigm's first foray into the study of real-world PC usage, but it certainly won't be the last. In future white papers, expect to hear more from us about the emerging trends in general PC usage, outside of prescribed workloads such as these. We will also continue to study common benchmarking tools to better understand their relative usefulness in predicting the quality of end-user experience on a given SSD.