

# Intel® Optane™ Memory

---

## ***Performance Evaluation Guide***

June 2020



## Ordering Information

Contact your local Intel sales representative for ordering information.

## Revision History

Revision Number	Description	Revision Date
001	<ul style="list-style-type: none"><li>Initial release</li></ul>	February 2017
002	<ul style="list-style-type: none"><li>Added support for Intel® Optane™ Memory M10 and H10 devices</li></ul>	April 2019
003	<ul style="list-style-type: none"><li>Non content-related format changes</li></ul>	April 2019
004	<ul style="list-style-type: none"><li>Prepared for public release</li></ul>	May 2019
005	<ul style="list-style-type: none"><li>Restructured information</li><li>Add configuration and descriptions of test scenarios</li></ul>	June 2020

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No product or component can be absolutely secure. Check with your system manufacturer or retailer or learn more at [intel.com/optane](http://intel.com/optane).

Intel does not control or audit the design or implementation of third-party benchmark data or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmark data are reported and confirm whether the referenced benchmark data are accurate and reflect performance of systems available for purchase.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

For copies of this document, documents that are referenced within, or other Intel literature please contact your Intel representative.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.



# Contents

---

- 1 About This Guide ..... 5
- 2 Overview ..... 6
  - 2.1 Performance Testing ..... 6
  - 2.2 Scenarios..... 6
  - 2.3 Workloads ..... 6
  - 2.4 Intel® Optane™ Memory Overview..... 7
    - 2.4.1 Intel® Optane™ Memory and Intel® Optane™ Memory M10..... 7
    - 2.4.2 Intel® Optane™ Memory H10 with Solid State Storage ..... 8
    - 2.4.3 Intel® Rapid Storage Technology Driver Operation..... 8
- 3 Benchmarks..... 9
  - 3.1 Application-based Benchmarks ..... 9
  - 3.2 Trace-based Benchmarks ..... 9
  - 3.3 Synthetic Benchmarks ..... 10
  - 3.4 Benchmark Considerations ..... 10
    - 3.4.1 Application and Trace-based Benchmark Considerations ..... 10
    - 3.4.2 Synthetic Benchmark Considerations ..... 11
- 4 Setup ..... 12
  - 4.1 Quick Start..... 12
  - 4.2 Setup Context..... 12
  - 4.3 All Other Things Equal..... 13
  - 4.4 System Setup ..... 13
    - 4.4.1 Processor ..... 13
    - 4.4.2 Memory..... 13
    - 4.4.3 Storage..... 13
    - 4.4.4 Drivers & Storage Management..... 14
    - 4.4.5 BIOS..... 15
    - 4.4.6 Operating System..... 16
    - 4.4.7 Power ..... 17
  - 4.5 Procedures ..... 18
    - 4.5.1 Enabling and Disabling Intel Optane Memory ..... 18
    - 4.5.2 Cache Flush Procedures..... 18
  - 4.6 Pretest Run Setup ..... 19
    - 4.6.1 Achieving Steady State ..... 19
    - 4.6.2 Getting a Good Test Run ..... 20
- 5 Analysis ..... 21
  - 5.1 Median Recommendation ..... 21
  - 5.2 Relative Standard Error ..... 21
  - 5.3 Methodology ..... 21
  - 5.4 Performance Delta..... 22



- 6 Appendix..... 23
  - 6.1 Benchmark Recommendations ..... 23
    - 6.1.1 SYSmark 2018 ..... 23
    - 6.1.2 PCMark ..... 23
    - 6.1.3 IOMeter ..... 24
    - 6.1.4 CrystalDiskMark ..... 25
    - 6.1.5 Other Device Benchmarks..... 25
  - 6.2 Example Test Scenarios ..... 25
  - 6.3 Metrics ..... 26
    - 6.3.1 Providers ..... 27
    - 6.3.2 Key Terms and Concepts for Storage Metrics ..... 27
    - 6.3.3 Recorders and Analyzers ..... 28



## 1 About This Guide

---

If you are evaluating platform or individual component performance in real-world or synthetic scenarios, then this guide presents considerations and recommendations when Intel® Optane™ memory is an ingredient of testing.

This guide also provides an overview of different workloads and benchmarks relevant to Intel® Optane™ memory.

The target audience includes publications, OEMs, technical analysts, academia and any who plan to test or evaluate Intel® Optane™ memory performance.

This guide is divided into the following sections:

- **Overview** provides an overview of Intel® Optane™ memory and establishes a foundation for testing.
- **Benchmarks** explores software designed to mimic a workload on a component or system and provide an indicator of performance.
- **Setup** This section includes a testing quick start plus examines the system and component configuration settings used to measure performance of a platform or component in real-world or synthetic scenarios.
- **Analysis** explores approaches for analysis and handling variability in result sets.
- **Appendix** presents supplementary material not included in the other sections including discussions on executing a set of tests and covers considerations for specific benchmark tools.

§



## 2 Overview

---

### 2.1 Performance Testing

Intel recommends an approach to performance evaluation should be reproducible via controls and observation to ensure fair analysis. This guide considers general principles of performance testing and applies them to Intel® Optane™ Memory.

Consider the below Scenarios and Workloads sections as one set of perspectives for how to approach performance evaluation. This perspective is the basis of methodology used at Intel.

### 2.2 Scenarios

The **real-world** scenario for evaluating performance is to measure common usage where machines are configured as they are out of the box or as in general practice for most users. Simply put, the test evaluates real world usage as machines are naturally configured. Generally, real-world scenarios are associated with **system** or **platform** testing.

The contrasting scenario is to configure machines in a **synthetic** state for measuring usage under manufactured conditions that are uncommon or artificially controlled. Synthetic here does not necessarily mean "best" for a component or platform. In a synthetic configuration for testing you may be exploring a boundary condition of the platform which may put some components at an advantage and others at a disadvantage. The synthetic scenario goal is to evaluate behavior of the test target for a specific condition or set of conditions that attempt to imitate a potential state. Generally, synthetic scenarios are associated with **device** testing.

Deciding when to apply **real-world** or **synthetic** (device optimized) configuration for testing can also be derived from what is measured under test and the goals of the test scenario. When measuring "wall clock time", and by extension the end-user perspective, this implies a real-world setup and scenario. When taking device metrics, such as latency, response time or throughput this implies a synthetic scenario with a device optimized configuration. In general, an optimized configuration will produce the highest possible numbers for a device-to-device comparison or for when measuring the specifications of the device.

### 2.3 Workloads

Workloads refer to a set of activities & any sequencing. Workloads for performance evaluation can be characterized in a variety of ways:

- **Target:** component (device) or platform (system)
- **Use Case:** categories of usage from the end-user perspective
  - o *Productivity:* Usage of common office or productivity applications such as the Microsoft Office suite of applications, Web Browsers and "light" audio/photo editing. This category of applications generally "produces" something.
  - o *Content Creation:* Professional software and usage by an enthusiast or professional such as Autodesk Computer Aided Design, 3D/Audio/Video editing & rendering among other "large" workloads where something is created.
  - o *Gaming:* Consumer gaming relies on multiple components and processing for a responsive experience. AAA games typically push the boundaries of audio/visual/interactive/processing capabilities of a system.
- **Tasking:** single or multi-tasking; also thought of in terms of foreground and background tasks.



## Intel® Optane™ Memory


### 2.4 Intel® Optane™ Memory Overview

Intel® Optane™ memory is a system acceleration solution. This solution uses Intel® Optane™ memory, based on Intel® Optane™ memory media, along with the Intel® RST driver. This revolutionary memory media is conceptually located between the processor and slower storage devices (HDDs). We can store commonly used data and programs closer to the processor so the system can access information more quickly, thereby delivering improved overall system responsiveness.


Intel® Optane™ memory requires specific hardware and software configurations.

Visit [www.intel.com/OptaneMemory](http://www.intel.com/OptaneMemory) for configuration requirements.

#### INTEL® OPTANE™ MEMORY M10



#### INTEL® OPTANE™ MEMORY H10 WITH SOLID STATE STORAGE



**NOTE:** Celeron™ and Pentium™ processors also supported

#### 2.4.1 Intel® Optane™ Memory and Intel® Optane™ Memory M10

Both the first- and second-generation products are a single M.2 device that provides a PCIe Gen 3x2 solution containing the Intel® Optane™ memory media. Both products are also designed to combine the performance of Intel® Optane™ memory with SATA based storage devices and some NAND-based storage devices such as the Intel SSD 660p/665p. Once enabled, the Intel® Optane™ memory accelerated volume appears as a single device to the end user. The second-generation Intel® Optane™ memory product adds support for low-power modes for mobile platforms.

See the link below for more information on this product.

<https://ark.intel.com/content/www/us/en/ark/products/series/99743/intel-optane-memory-series.html>



## 2.4.2 Intel® Optane™ Memory H10 with Solid State Storage

Intel® Optane™ Memory H10 with Solid State Storage is a single M.2 device that provides a dual ASIC solution. One ASIC is a PCIe Gen 3x2 solution containing the Intel® Optane™ memory media. The other ASIC will also be a PCIe Gen 3x2 solution and will manage the Intel® QLC 3D NAND media.

Prior to enabling Intel® Optane™ memory, this single M.2 device will enumerate as two distinct drives in the operating system. Once enabled, Intel® Optane™ memory driver will present the Intel® Optane™ memory accelerated volume as a single device to the end user.

For more information on this product see:

<https://www.intel.com/content/www/us/en/products/docs/memory-storage/optane-memory/optane-memory-h10.html>

## 2.4.3 Intel® Rapid Storage Technology Driver Operation

The Intel® RST driver combines the Intel® Optane™ memory device with a high capacity storage device to create a single virtual drive. This virtual drive will have regions that vary from very high performance and low latency to lower performance with lower cost per GB of capacity. The Intel® RST driver seamlessly manages this virtual drive, ensuring that frequently used data immediately resides on the fastest storage device, and over time, seamlessly migrates data between the faster and slower storage device.

On systems with traditional storage (e.g. HDD, NAND-based SSD), boot performance, application start, data access and virtual memory paging performance are all limited by the storage performance. This negative effect is even greater when there are other I/O intensive background activities running simultaneously, like system software updates or virus scans.

Intel® Optane™ memory resolves these issues by caching boot data, executables, frequently accessed user data and the system page file. While the user sees the full capacity of the traditional storage, the storage sub system delivers key data with the responsiveness of an Intel® Optane™ SSD.

§





## 3 Benchmarks

---

Benchmarks are designed to mimic a workload on a component or system and provide an indicator of performance. This section provides an overview of three basic types of benchmarks that can be used to measure system storage performance:

- Application-based Benchmarks
- Trace-based Benchmarks
- Synthetic Benchmarks

For specific benchmark recommendations and information please see the [Benchmark Recommendations](#) subsection in the appendix.

### 3.1 Application-based Benchmarks

Application-based benchmarks emulate end-user usage using scripted execution of real-world programs on a system. Application-based benchmarks measure the load and execution time of these applications and present the results as a score. Often the scores of applications that are common to a usage are combined into a subsystem score. These subsystem scores are reported along with an overall performance score that is the combination of the subsystem scores.

[SYSmark 2018](#) is an example of an application-based benchmark. It has 4 subsystem scores and 1 overall score. Another example of application-based benchmarking is the realistic usage guides (RUGs) developed by Intel.

Application-based benchmarks are helpful in determining the User Experience (UX) for a given system. In particular, the scores from an application-based workload are more likely to reflect the real world UX than the results from a purely synthetic workload.

One disadvantage of application-based benchmarks when trying to determine storage device speed is that they focus on CPU, memory, and graphics performance and may not properly weigh storage sub-system speed in the results. They also may not consider end-user perceivable delays such as application loads. Another disadvantage is that the scripted nature of an application-based benchmark is fixed, and the application workload may not be representative of an end-user usage model, especially when multiple iterations are repeatedly in sequence. Lastly, due to practical considerations, rather than constrain the total allowable runtime and total size of the benchmark (e.g., download based distribution, total amount of disk space required to run), many application-based benchmarks have a short-run duration and limited storage device usage. Therefore, these benchmarks will not be representative of storage device usage over time as a practical basis.

### 3.2 Trace-based Benchmarks

Trace-based benchmarks use traces, or recordings of disk I/O operations executed during a certain period of real use or script-based use. The trace is then used to 'playback' the system I/O sequence on the drive to be tested. Trace-based benchmark results vary in format.

[PCMark Vantage HDD](#) test is an example of a trace-based benchmark. Trace-based benchmarking has many of the advantages of application-based benchmarking, if the trace is collected from real-use or a realistic script-based activity. It has the further advantage of highlighting disk I/O behavior while avoiding the bottlenecks caused by the CPU, graphics, and memory subsystems.

One disadvantage of trace-based benchmarking is that the recorded trace may not reflect the true long-term usage of the storage device over weeks or months of time.



## 3.3 Synthetic Benchmarks

Synthetic benchmarks measure raw drive input/output (I/O) transfer rates. These benchmarks typically use well-defined, synthetic workloads and target only specific components such as Solid-State Drives (SSDs), Hard Disk Drives (HDDs) or other networking devices. These benchmarks format results as raw megabytes (MBs) in Input / Output Operations per Second (IOPS).

Storage subsystem synthetic benchmarks focus on drive performance without considering bottlenecks from other subsystems such as CPU, memory, or graphics. This makes these benchmarks useful for measuring drive performance for changing parameters such as transfer sizes. However, because these benchmarks exercise components and systems in ways that do not reflect system usage models, the results may not reflect real-usage cases. For example, one SSD might have better synthetic benchmark scores for 512 KB random reads than other SSDs, but 512 KB random reads may not be a good indicator of overall system performance because of the rarity of that particular I/O access size in what is important to a given end user experience.

Synthetic benchmarks can be divided into two sub-groups: long and short.

- Long duration synthetic benchmarks measure performance variation over the entire run time. [IOMeter](#) is a variable length benchmark that when properly configured can perform as a long duration synthetic benchmark
- Short synthetic benchmarks are commonly used to measure component performance in the immediate present. These typically have limited configurability to see performance variations over time. [CrystalDiskMark](#) is considered a fixed, short duration synthetic.

## 3.4 Benchmark Considerations

### 3.4.1 Application and Trace-based Benchmark Considerations

To emulate real end user behavior, usage, and temporal state, there are three main challenges with typical application and trace based benchmarks: controlling the repeatability of results, accounting for the effects of “system aging”, and accurately modeling the storage footprint.

The repeatability of results can be challenging simply due to the complexity and uncertainty of the system with background processes and other runtime services. In addition, the starting state of the cache should be in a known state in order to ensure repeatability.

Accounting for the effects of “system aging” can be thought of as three sub-challenges. First, in normal system usage, the unused part of main memory will be preloaded with a complex set of applications, services, OS references, and prefetched data in anticipation of upcoming usage. The state of system memory will directly impact benchmark behavior. Second, multi-iteration benchmarks perform the same sequence of actions in repeated iterations, a behavior not often found in end users. The result of such benchmarks is an unrealistic situation where a modest amount of main memory can contain most of the necessary data and accessing storage isn’t required. Finally, accurate steady state performance measurements of a storage subsystem that implements any type of caching requires emulation of a warm-up sequence of storage I/O activity.

The storage footprint of users' systems will vary and can be quite large. Emulating a storage footprint in a workload or benchmark can be challenging because of the sheer size of the files required to achieve it.



### **3.4.2 Synthetic Benchmark Considerations**

In addition to the considerations outlined in the section above, benchmarking caching software with synthetic I/O workload generators like IOMeter may lead to non-intuitive results. This is because what data gets inserted and remains in the cache is subject to the cache policy contained within the Intel® RST driver. The cache policy is adaptive in its behavior. When the cache is first initialized and enabled the policy is to aggressively insert IO until it is roughly full. After this point, the policy adapts to a longer term "steady state" algorithm optimized for the long-term usage of the system. Once in this state, sequential data streams may be sent directly to the HDD as well.

The cache policy also differentiates between high-value data and low-value data like that which may be caused by a virus scanner or other similar "one-touch" data which typically does not benefit from caching, and therefore is not inserted into the cache.

Due to all these factors and the intentional design of the cache algorithm to optimize for actual in-system typical client workloads, running synthetic workloads on caching storage subsystems can lead to results that aren't intuitive or appear to offer no performance benefit.

§



## 4 Setup

---

This section examines the system and component configuration settings used to measure performance of a platform or component in real-world or synthetic scenarios.

### 4.1 Quick Start

This section recommends a set of steps to quickly get started with testing and evaluating Intel® Optane™ memory.

#### Plan

- Consider the audience, context and desired learnings for testing to plan your test.
- Identify whether your test target is the platform or a component in a real-world or synthetic scenario.

#### Setup

- In most cases you will either (1) run a benchmark or (2) run a workload noting the start and stop times.
- Prepare your system and storage component as dictated by your test goals, referring to the entries in the [System Setup](#) section.
- See the [Benchmark Recommendations](#) section in the appendix for a list of benchmarks relevant to storage recommended by Intel. Each benchmark noted in the appendix entry includes a description and tips for running.
- Before executing the test be sure to review the [Pretest Run Setup](#) section for considerations and recommendations.

#### Execute

- Run your test(s) and capture the results.

#### Analyze

- Look at the [Analysis](#) section of this guide for Intel's recommendations on interpreting and analyzing the results.

### 4.2 Setup Context

When evaluating the performance of **platforms with Intel® Optane™ memory**, Intel recommends configurations that reflect the way machines are configured when used in the field. For the tables below these are the settings in the "Real World Value" column.

To ensure consistent and repeatable measurements, a stable and deterministic environment is desired. Each configuration entry listed in the tables and their associated settings have been evaluated for their impact on performance results and variability. If disabling a technology will reduce run to run variance without materially impacting performance results it will be disabled. If disabling a feature will create a meaningfully higher or lower performance score it will be left untouched from "normal" OEM default settings and Intel recommends additional runs to enable a statistical analysis to filter out the variability. Some settings are disabled because they have a substantial impact on performance but infrequently execute, an example being "Automatic Windows Updates". Disabling these types of features avoids having to debug unexpected results.

When evaluating the performance of Intel® Optane™ memory devices with some workloads, for example low queue depth synthetic IO, it is critical to ensure that the platform features do not mask the performance of the device.

For the tables below the settings in the "Synthetic Value" column create a test environment that maximizes the performance of the storage device. C-States, for example, are critical to the ability of modern mobile platforms to save power and extend battery life. But because having C-States enabled may have a secondary effect of causing the CPU to be less responsive to IO completion, they are disabled.



## 4.3 All Other Things Equal

Make sure your tests are setup the same or as intended to ensure results do not vary unexpectedly or draw conclusions from hidden causes.

**Ceteris paribus** is a Latin phrase that translates to "other things equal". Assume that unless otherwise mentioned "all other things being the same" is implied for test setup. This is especially important when comparing different test configurations. Care should be taken to identify independent (changed or controlled) and dependent (tested and measured) variables. Recording these variables is essential to ensuring we have a reproducible test scenario that can be objectively verified.

## 4.4 System Setup

The below setup items are recommended to prevent run-to-run and system-to-system variability that impact the ability to reproduce the test and results. If one of the below items is a controlled variable or the target of a test, then modify as desired.

### 4.4.1 Processor

Processors should be consistent with the use case of the test however different chips will have different platform level performance characteristics. Relevant specifications are base clock, turbo clock, number of cores, hyperthreading, and package power. For example, the Intel® Core i3-9350K has a base clock of 4.0 GHz, where the Core i7-9700K has a base clock of 3.6 GHz. The i3 is a lower bin chip but could potentially exhibit better storage performance on some workloads due to the faster base clock.

### 4.4.2 Memory

System memory should be consistent with the use case of the test, and consistent across multiple system configurations unless system memory is an independent variable and performance effects are being measured (e.g. Intel® Optane™ Memory H10 + 8GB RAM vs TLC NAND + 16GB RAM). In most cases, system memory should be representative of the market segment of the product: 8GB or 16GB for mobile, 16GB for desktop, with DDR4 speeds matching memory specs of the processor being tested.

### 4.4.3 Storage

For installation and troubleshooting of Intel® Optane™ memory or Intel® Optane™ Memory with Solid State Storage consult the [Intel® Optane™ Memory M and H Series Installation Guide](#).

Intel recommends prefilling and conditioning a brand-new or out-of-the-box drive. Drives that are mostly empty and those that have never had at least a majority of the media written are not ideal for testing in that they can both perform far beyond spec and produce much more variable results than a drive that has had the media sufficiently "activated" and filled with use. Perform two full writes to the drive before testing and we recommend formatting the drive then using [IOMeter](#) to create a test file with an infinite span, repeated twice.

For all drives Intel recommends a cache flush before a test run and that the drive is filled to 50% capacity.



#### 4.4.3.1 Prefill

The process of pre-filling the drive will saturate the SLC cache and, unless accounted for, will adversely affect drive performance. The SLC cache must be flushed prior to testing.

Drives must be pre-filled to an acceptable level to standardize drive performance. Unless purposefully deviating from this standard, this should be a 50% capacity pre-fill. To pre-fill use one of the below methods.

1. The simplest acceptable method is to write multiple copies of a large file to the disk. Create multiple copies of a file until the drive's capacity is 50% filled.
2. Create an IOMeter test file that spans half the drive.

#### 4.4.3.2 Cache Flush

The goal in flushing SLC is to ensure cache consistency across multiple test sets.

Intel H10, Intel SSD 760p and 660p/665p Series with NAND media have an SLC cache in addition to the TLC or QLC media. Many competitors' drives also implement an SLC cache.

For the Intel® Optane™ Memory H10 with Solid State Storage, Intel SSD 660p/665P the SLC cache should be flushed for the authenticity of test conditions when balancing consistency and reproducibility of results against the complexity of setup of in a practical time. Intel recommends a flush of SLC as a part of initial test preparation after pre-filling a drive and between test sets during device-specific testing (high performance power plan). This is done to put the drive into the state that is common during normal use, to offset the following:

- Rapid filling of the drive under test
- High duty cycle of benchmarking
- The need to start testing promptly

See the [Cache Flush Procedures](#) for how to flush the SLC cache.

An SLC Cache Flush is recommended for **synthetic** scenarios, such as evaluating Intel Optane memory devices for max I/O performance.

### 4.4.4 Drivers & Storage Management

Storage drivers are of critical importance to Intel® Optane™ memory testing.

- For systems with Intel® Optane™ memory, the storage controller drivers should be managed by the latest public release of Intel® Rapid Storage Technology (Intel® RST). This driver will allow for the creation of an Intel® Optane™ RAID volume and will manage the Intel® Optane™ caching and acceleration.
- Systems without Intel® Optane™ memory should use the Windows Inbox (default) driver.
- Certain test configurations may call for the use of a different driver, but only if the driver is the independent variable in the test (e.g. effect of using Intel RST with Intel SSD 760p while in DC mode vs. Inbox driver with the 760p while in DC mode).

There are currently two applications you can use to manage your Intel® Optane™ memory device that include drivers. The [Intel® Optane™ Memory and Storage Management](#) application can be downloaded from the Microsoft Store.



## Intel® Optane™ Memory

The Intel® RST User Interface Driver is currently available but scheduled to phase out in 2020. Go to the [Intel Download Center](#) and select the 'SetupRST.exe' [download link](#) from the available choices. Both choices include the Intel® RST driver stack which includes mandatory platform and OS drivers.

Intel recommends using the Intel® Optane™ memory and Storage Management application to enable Intel® Optane™ memory to accelerate your system.

Configuration Item	Value
Intel® Optane™ Memory and Storage Management > Intel® Optane™ Memory > Intel® Optane™ Memory Status	Enabled
Intel® Rapid Storage Technology > Intel® Optane™ Memory > Intel® Optane™ Memory	Enabled

On major OEM production systems, all drivers should be readily available via Windows Update. Ensure all updates have been run before disabling the update service for testing. On OEM pre-production systems, software image files will be provided that include appropriate drivers. Either construct a testing image from the provided system image or import the drivers from the image to a fresh install of Windows.

### 4.4.5 BIOS

**Note:** Occasionally an OEM will not expose some BIOS settings.

While specific BIOS options may vary between OEM and BIOS revision, the SATA controller mode should be configured for Intel® RST Premium with Intel® Optane™ System Acceleration. See OEM BIOS documentation for support on setting details. All other settings should *generally* be left to default with any deviations documented along with an explanation of impact.

In the following table a Real World and a Synthetic value are given for each. As you may recall, a *synthetic* value may be employed for specific test scenarios while *real-world* indicates how the system should be normally configured as a default for practical usage.

Configuration Item	Real World	Synthetic
Hyper Threading	Enabled	
EIST (Enhanced Intel Speed Step Technology)	Enabled	Disabled
Intel Turbo Mode	Enabled	
PCIe ASPM (Active State Power Management)	Enabled	Disabled
C-States	Enabled	Disabled
P-States	Enabled	Disabled



### 4.4.6 Operating System

While system and device tests should normally be performed across a standardized OS installation, in some test scenarios it may be more appropriate to evaluate OEM systems in their 'out of the box' configuration, to retain any OEM-specific tuning, etc. When testing multiple systems, ensure the installed KB's (or Microsoft Windows patches) are as consistent as possible. Below is a table of relevant configuration settings for Windows 10.

In the following table a Real World and a Synthetic value are given for each. As you may recall, a *synthetic* value may be employed for specific test scenarios while *real-world* indicates how the system should be normally configured as a default for practical usage.

Configuration Item	Real World	Synthetic
Settings > Accounts > Sign-in Options > Privacy > "Use my sign-in info to automatically finish setting up my device and reopen my apps after an update or restart"	Disabled	
Settings > Updates & Security > Windows Security > Open Windows Security > Virus & threat protection > Manage Settings > Real-time protection <sup>1</sup>	Off	
Settings > Updates & Security > Windows Security > Open Windows Security > Virus & threat protection > Manage Settings > Cloud delivered protection	Off	
Settings > Updates & Security > Windows Security > Open Windows Security > Virus & threat protection > Manage Settings > Automatic sample submission	Off	
Settings > Updates & Security > Windows Security > Open Windows Security > Virus & threat protection > Manage Settings > Tamper Protection	Off	
Settings > Updates & Security > Windows Security > Firewall & network protection > Domain Network   Private Network   Public Network > Windows Defender Firewall	Off	
Settings > use the search input and type in "Indexing Options"	Remove all from "Included Locations"	
Settings > Personalization > Lock Screen > Screen saver settings	None	
Control Panel > System and Security > System > System Protection > Available Drives > Protection	Off	
Control Panel > System and Security > System > Advanced system settings > Advanced > Performance > Settings > Visual Effects	Let Windows choose what's best for my computer	
Control Panel > System and Security > System > Advanced system settings > Advanced > Performance > Settings > Advanced >	Adjust for best performance of Programs	
Control Panel > System and Security > System > Advanced system settings > Advanced > Performance > Settings > Advanced > Virtual Memory > Change	Automatically manage paging file size for all drives	No paging file
Control Panel > System and Security > System > Advanced system settings > Advanced > Performance > Settings > Data Execution Protection	Turn on DEP for essential Windows programs and services only	
Control Panel > System and Security > Administrative Tools > Defragment and Optimize Drives > Scheduled Optimization	Uncheck "Run on a schedule (recommended)"	





## Intel® Optane™ Memory

Control Panel > System and Security > Administrative Tools > Services > Windows Update	Disabled	
Control Panel > System and Security > Administrative Tools > Task Scheduler > Task Scheduler Library	Disable all tasks	
Registry Editor > HKEY_LOCAL_MACHINE > SYSTEM > CurrentControlSet > Control > Session Manager > Memory Management > PrefetchParameters > EnablePrefetcher	3	0
Anti-Virus	Disabled	

### Note:

1. *Real Time Protection* will reset to “On” after a reboot. Take special care to disable this setting and *Tamper Protection* in Windows Security.

## 4.4.7 Power

For laptops Intel recommends performance testing be done with the system plugged into AC power mode.

In the table below a Real World and a Synthetic value are given for each. As you may recall, a *synthetic* value may be employed for specific test scenarios while *real-world* indicates how the system should be normally configured as a default for practical usage.

Configuration Item	Platform	Real World	Synthetic
Settings > System > Power & Sleep > Screen	Desktop Laptop	Never for both	
Settings > System > Power & Sleep > Sleep	Desktop Laptop	Never for both	
Control Panel > Hardware and Sound > Power Options > Power Plan	Desktop	Balanced	High Performance
Control Panel > Hardware and Sound > Power Options > Power Plan	Laptop	Balanced	
Control Panel > Hardware and Sound > Power Options > Change plan settings > Change advanced power settings > Turn off hard disk after	Desktop	0 minutes (never)	
[INHERIT] > Internet Explorer > JavaScript Timer Frequency	Desktop	Use defaults	Maximum Performance
[INHERIT] > Desktop background settings > Slide show	Desktop	Paused	
[INHERIT] > Wireless Adapter Settings > Power Saving Mode	Desktop	Use defaults	Maximum Performance
[INHERIT] > Sleep > Sleep after	Desktop Laptop	0 minutes (never)	
[INHERIT] > Sleep > Allow wake timers	Desktop Laptop	Use defaults	Enabled
[INHERIT] > USB settings > USB selective suspend setting	Desktop	Use defaults	Disabled
[INHERIT] > Intel Graphics Power Plan	Desktop	Use defaults	Maximum Performance



[INHERIT] > Power buttons and lid > Power button action	Desktop Laptop	Do nothing	
[INHERIT] > Power buttons and lid > Sleep button action	Desktop Laptop	Do nothing	
[INHERIT] > PCI Express > Link Power Management	Desktop	Use defaults	Off
[INHERIT] > Processor power management > Minimum processor state	Desktop	Use defaults	100%
[INHERIT] > Processor power management > System cooling policy	Desktop	Use defaults	Active
[INHERIT] > Processor power management > Maximum processor state	Desktop	Use defaults	100%
[INHERIT] > Display > Turn off display after	Desktop Laptop	0 minutes (never)	
[INHERIT] > Use defaults	Desktop	Use defaults	100%
[INHERIT] > Display > Dimmed display brightness	Desktop	Use defaults	100%
[INHERIT] > Display > Enable adaptive brightness	Desktop	Use defaults	Off

## 4.5 Procedures

### 4.5.1 Enabling and Disabling Intel Optane Memory

Configuration Item	Value
Intel® Optane™ Memory and Storage Management > Intel® Optane™ Memory > Intel® Optane™ Memory Status	Enable   Disable
Intel® Rapid Storage Technology > Intel® Optane™ Memory > Intel® Optane™ Memory	Enable   Disable

### 4.5.2 Cache Flush Procedures

Some SSD manufacturers provide an interface or other means to manage an SSD, allowing for a user to flush the cache with the click of a button. If not, an idle of the system is necessary.

For drives that do not support a manual cache flush operation, including the Intel® SSD 760p or Intel® Optane™ Memory H10, the cache must be flushed by means of a system idle. Use the following experiment to determine an adequate idle time for the flush to occur.

1. Set the system power profile to High Performance.
2. Ensure the SLC cache is fully saturated (via repeated write operations, usually a pre-fill).
3. Allow the system to idle 5 minutes.
4. Copy an 8 GB file to the target drive and observe whether the write happens at SLC speeds or QLC/TLC speeds.
5. If the write speed drops off during the write operation, 5 minutes is an insufficient period to systematically flush the SLC cache. Repeat the experiment with a 10-minute idle. If the write speed continues to drop off, increase the idle time in increments of 5. Once it remains consistent at SLC write speeds, that time interval is treated as an acceptable idle period to systematically flush the SLC cache in High Performance.

**Note:** This conclusion will only hold true while in High Performance mode. Under a Balanced plan, the system may not flush the SLC cache in the same interval.



## 4.6 Pretest Run Setup

The below best practices at Intel are recommended provided they align with your test scenario and goals.

1. If executing a component specific test, ensure that [Intel® Optane™ Memory](#) has been toggled to flush the Optane cache and that the SLC cache has been either systematically or manually flushed. Trigger volume optimization on systems that contain 32 GB of Intel Optane memory to ensure steady state performance of the File Cache. See the [Achieving Steady State](#) section below.
2. If executing a platform level test, ensure that [Intel® Optane™ Memory](#) has been toggled to flush the Optane cache. Do not flush the SLC cache. Do not manually trigger volume optimization.
3. The system must be rebooted between test runs to ensure test data is completely flushed from system memory. No need to reboot for [SYSMark 2018](#) and [PCMark 10](#) as automated reboots happen as part of the testing suite.
4. The system must be idled after startup to ensure minimal system process interference with test workloads. A 5-minute period is enough for this steady state.

### 4.6.1 Achieving Steady State

The below section is recommended for device or synthetic testing and should not be used during platform or real-world testing.

When evaluating the performance of Intel® Optane™ memory 32 GB or larger SKUs with benchmarks it is recommended to put a portion of the Intel® Optane™ memory driver into a "steady state." Achieving this "steady state" is done by filling the cache with user data to ensure "test" data (from benchmarking tools) is not included in the cache. Failing to fill the cache with user data will skew future results. This is recommended for benchmark tools designed to assess device performance and not platform performance, such as [IOMeter](#) and [CrystalDiskMark](#).

After [Intel® Optane™ memory](#) has been successfully enabled, take the following steps PRIOR to running tests that are designed to evaluate device performance.

Manually trigger execution of the Intel® Optane™ memory – Volume Optimization process:

1. Launch Task Scheduler
2. In Task Scheduler, expand the Task Scheduler Library folder in the left column. Click on Intel. The Intel® Optane™ memory – Volume Optimization task will appear
3. From the Actions column on the right, click Run

Check for completeness after manually triggering the Intel® Optane™ memory – Volume Optimization:

4. Launch Task Manager (Ctrl + Shift + Esc) and click the Performance tab. You will see high activity on Disk 0 (C:)
5. Wait for the activity to revert to idle (0-2%)
6. Verify this is done using the [Intel® Optane™ Memory and Storage Management](#) application. Under the Statistics tab, the Last Optimization time and date should be populated

Manually trigger a second execution of the Intel® Optane™ memory – Volume Optimization process:

7. Follow the previous steps to ensure the activity is done and to verify.
8. Restart the PC.
9. After reboot, confirm the [Intel® Optane™ Memory and Storage Management](#) application reports an unused space of  $\leq 0.1$ GB. Unused Space can be found by launching Intel® Optane™ memory and navigating to the Pinning tab.



## 4.6.2 Getting a Good Test Run

The first run of any test on Intel® Optane™ memory should be considered a "cold" run. Because we are evaluating a caching drive, the first run after flushing the cache will be running solely off the NAND media.

- The [IOMeter](#) configuration detailed in this guide includes a conditioning run. This conditioning run takes the place of a cold run and is automatically discarded. Do not discard the first iteration of IOMeter testing using these configurations.
- Subsequent test runs should be considered "warm" runs. These runs will be representative of the device's performance when using the caching solution. Record the results and proceed to the next run when complete.



## 5 Analysis

### 5.1 Median Recommendation

Intel uses and recommends the use of median values so an actual value from the dataset is reported as a characteristic value. Do not use average as it is a 'calculated' value. A minimum of three 'warm' runs must be collected to calculate a median value.

### 5.2 Relative Standard Error

*"Relative Standard Error: A unit-free measure of the reliability of a statistic, defined as the absolute value of the ratio of the standard error to the sample estimate of the statistic, expressed as a percentage." - J. Black, N. Hashimzade & G. Myles in 'A Dictionary of Economics'*

Intel uses and recommends relative standard error (RSE) to evaluate variability in the dataset. If relative standard error is above 5%, Intel views the variability is unacceptably high and more datapoints must be collected or the test set must be discarded. Depending on the context, RSE below 5% may also be unacceptable—lower is better and more datapoints carry more weight. If a test set takes more than 7 runs to reach acceptable variability, Intel recommends the test set should be discarded.

The RSE of a sample is defined as:

$$\left| \frac{\text{Standard Error (of the sample)}}{\text{Sample Estimate}} \right| = \left| \frac{\text{Standard Deviation (of the sample)}}{\text{Sample Estimate} \cdot \sqrt{\text{Number of Samples}}} \right| = \left| \frac{s}{\bar{x} \cdot \sqrt{N}} \right|$$

In Excel, this can be expressed as a percentage with the following formula:

```
=STDEV.S(range)/SQRT(COUNT(range))/AVERAGE(range)
```

### 5.3 Methodology

When evaluating new workloads, Intel expects to have reached a steady state performance on our caching drives after the first cold run. Additionally, the first three warm runs should have an acceptable relative standard error. Workloads that do not fit this expectation should undergo further analysis. If a suitable explanation is found, the testing window can be expanded to five warm runs so long as the deviation is documented, and all stakeholders are made aware of the workload's characteristics.

When making claims on device or platform performance for storage workloads, Intel evaluates the relative standard error and calculates the median from the most recent three runs. If the relative standard error is unacceptable, a run is added and the sliding window of three runs is shifted. If the RSE of the most recent three runs are not acceptable after at most seven total runs, the test set is discarded.



An example of this sliding window methodology requiring all seven total runs is as follows. The workload is a generic application launch with some background activity, and units are in seconds. The sliding window used to calculate median and RSE is highlighted in blue, and each row represents a successive test run appended to the test set.

Run 1	Run 2	Run 3	Run 4	Run 5	Run 6	Run 7	Median	RSE
29.64	22.67	11.95	--	--	--	--	22.67	24.02%
29.64	22.67	11.95	11.00	--	--	--	11.95	24.61%
29.64	22.67	11.95	11.00	9.12	--	--	11.00	7.78%
29.64	22.67	11.95	11.00	9.12	9.21	--	9.21	6.26%
29.64	22.67	11.95	11.00	9.12	9.21	9.24	9.21	0.39%

### 5.4 Performance Delta

The performance delta between two drives in a test is calculated differently depending on the test. This is because some tests are measured in seconds (lower is better), where others are measured in speed or score (higher is better).

For tests measured in time, relative performance is measured by  $\Delta = \frac{\text{basis}}{\text{comparison}} - 1$  where if we were testing H10 against, for example, a 760p, it would be  $\Delta = \frac{\text{time}_{760p}}{\text{time}_{H10}} - 1$ . This yields a value that should be treated as a percentage improvement in speed (because the values are inverted, it shows an improvement in speed, rather than time).

For tests measured in speed/score, relative performance is measured by  $\Delta = \frac{\text{comparison}}{\text{basis}} - 1$  where if we were testing H10 against, for example, a 760p, it would be  $\Delta = \frac{\text{speed/score}_{H10}}{\text{speed/score}_{760p}} - 1$ . This yields a value that should be treated as a percentage improvement in speed/score.

§



## 6 Appendix

This section presents supplementary material not included in the other sections including discussions on executing a set of tests and covers considerations for specific benchmark tools.

### 6.1 Benchmark Recommendations

Intel® Optane™ memory will cache data covering many weeks, even months into the past. All benchmarks are approximations of the user experience.

Intel recommends application-based and trace-based benchmarks to evaluate system performance with Intel® Optane™ memory products.

The following table summarizes how Intel uses each benchmark.

Benchmark Type	Usage
Application-based benchmarks	Used as part of the functional and performance validation flows. Recommended for ranking storage solutions due to its approximation of real-world usages.
Trace-based benchmarks	Used as part of the functional and performance validation flows.
Synthetic benchmarks	Measure raw input/output (I/O) performance.

#### 6.1.1 SYSmark 2018

The feature of SYSmark 2018 which enables a warmup run **must** be used.

[SYSmark 2018](#) is an example of an application-based benchmarking tool and can be used to 'score' a platform on its ability to meet user expectations.

Application-based benchmarks require more than one run to notice substantive gains. A drive's behavior will more closely match its real use behavior when tested with an application-based workload, than it will with a purely synthetic workload. Application-based benchmarks can sometimes produce a single measure for system performance, which can be used as a product ranking index. It is critical to follow the steps below to facilitate repeatability and accuracy.

Ensure the default conditioning run is performed to populate caches with frequently accessed data.

#### 6.1.2 PCMark

##### 6.1.2.1 PCMark 10

[PCMark 10](#) has two storage tests for device-oriented benchmarking which Intel recommends:

- The **Full** System Drive Benchmark uses a wide-ranging set of real-world traces from popular applications and common tasks to fully test the performance of the fastest modern drives. This benchmark includes a comprehensive set of tests executed in a very short period of time spanning >200GB of data. Consider use of this benchmark carefully as it might not show the true day-to-day performance of hybrid SSD solutions.
- The **Quick** System Drive Benchmark is a shorter test with a smaller set of less demanding real-world traces. You can use this benchmark to test smaller system drives that are unable to run the Full System Drive benchmark. The Quick test is not a shorter version of the full test running just a subset of tests and is more representative of typical daily usage that access a relatively shorter span of the SSD. This test shows what would be expected regarding hybrid SSD performance on a day to day basis when compared to the Full test.



### 6.1.2.2 PCMark Vantage HDD

Intel recommends not to run PCMark Vantage HDD benchmark on mostly empty drives. Performance of the benchmark is skewed under this condition as it performs a replay of existing data on the drive to produce load. This has been fixed in PC Mark 10 as the benchmarks first puts data on the device and then starts testing.

The [PCMark Vantage HDD](#) storage test is an older device-oriented benchmark.

PCMark Vantage HDD benchmark is a highly sensitive test which produces large differences in score with only minor storage performance changes. The benchmark predates many recent advances in SSD storage, such as TRIM, and consequently needs procedural modification to most effectively measure modern storage configurations.

The most important of these is to prevent the benchmark from issuing any read operations to drive sectors with no valid data (i.e. sectors that are trimmed).

This is critically important because reading data that was trimmed:

- Is not the intent of the benchmark developers
- Increases dependency on non-storage platform components (e.g. CPU speed and DRAM speed)
- Does not represent a realistic user workload.

To overcome this benchmark behavior Intel recommends that 100GB of data be written to the target drive. Doing so means that most random reads will be issued to initialized data during the benchmark trace replay workload. To accomplish this writing of data Intel uses IOMeter to create the file.

The Intel recommended procedure to execute PCMark Vantage HDD when using Intel® Optane™ Memory is as follows:

1. Secure Erase disk
2. Install operating system and benchmark software
3. Intel recommends filling the drive with 100 GB of data. If using IOMeter IOBW.TST file to do so, set 'Maximum Disk Size' of 209,715,200 sectors when creating the file.
4. 30min Idle Time with PCIe link active
5. Enable Intel® Optane™ Memory
6. Ensure the File Cache is in a steady state as described in previous chapters
7. Run PCMark Vantage HDD
8. Allow for a 10 min Idle Time. For additional iterations go to step #7.

### 6.1.3 IOMeter

[IOMeter](#) is a flexible tool that enables storage device application of user-defined workloads.

IOMeter test sequences should be crafted with the following considerations:

- Test file size – Intel recommends a test file size of 8GB (16777216 sectors). Before testing, additional files may be copied to the device to achieve desired amount of fill.
- Preconditioning for steady state – The common practice of applying a higher QD sequential or random write before running measured tests of the same type, may not attain steady state performance for some hybrid caching SSDs. In addition to the above conditioning pass, in order to emulate the first (potentially not cached) access of stored data, consider performing an unmeasured 'cold run' of each desired workload.
- Workload duration – Heavy workloads applied at a very high duty cycle may push smaller form factor SSDs beyond their expected thermal envelope. This can potentially trigger performance throttled conditions that would not have occurred during normal use. Thermal throttling can be alleviated by:
  - Reducing the duty cycle of the tests to better match real-world usage.
  - Active cooling may be required during testing, if duty cycle cannot be reduced.





## 6.1.4 CrystalDiskMark

[CrystalDiskMark](#) is a simple disk benchmark software for measuring sequential and random performance (Read/Write/Mix) and contains profiles for real world and peak performance.

Intel recommends using the "Real World Performance" and "Real World Performance [+Mix]" profiles when evaluating storage devices. CrystalDiskMark should also be run as administrator (right click – run as admin). CrystalDiskMark 7 introduced a new feature that allows the user to set processor core affinity. Using the ``-ag`` option sets round robin affinity per thread. This option is recommended for higher random read throughput for 4K transfer sizes.

## 6.1.5 Other Device Benchmarks

The following list presents additional examples of device-oriented benchmarks.

- [AS SSD](#)

- Anvil's Storage Utilities

- [ATTO](#)

- [HD Tune](#)

- [TxBENCH](#)

**Note:** The above benchmarks perform minimal if any preconditioning, minimal to no cold runs, may not be suitable for evaluating cached performance of hybrid storage devices (test files are typically deleted / trimmed after test completion), and may be coded in a way that does not scale results proportionally when testing very low latency storage devices.

## 6.2 Example Test Scenarios

The below examples are provided to demonstrate best practice for how to evaluate performance. Each example describes how the test should be setup and can be reproduced through conditions, what measures will be taken, a description of the steps or instructions for the experiment and provides a verifiable conjecture.

For each condition below apply "all other things being equal"

### Platform & Real World

- *Conditions:* For each test employ 4/8/16/32 GB of DRAM, no constraints on foreground/background applications.
- *Metrics:* Collect time/duration during test run, from launch to application becomes interactive; Size of free DRAM; # of running processes; # of running foreground applications.
- *Experiment:* Time to launch Application A with different amounts of DRAM with commonly installed capacities of DRAM.
- *Conjecture:* Amount of DRAM in system has a linear relationship with Application A launch time.

### Platform & Real World

- *Conditions:* Disable/Enable uncritical Windows Services, Virus Scanner, Firewall, Search Indexing; No foreground applications; Execute immediately following Windows startup from cold boot.
- *Metrics:* Collect time/duration during test run, from launch to application becomes interactive.
- *Experiment:* Launch application "A" with and without overhead of computing maintenance processes.
- *Conjecture:* Application "A" launches faster unconstrained by computing maintenance processes.



Component & Real World

- *Conditions:* For each test employ a storage device with different media and interface in the following combinations: Hard Disk + SATA, NAND + SATA, NAND + PCIe, Optane + PCIe; Size and number of images files consistent across tests.
- *Metrics:* Disk read bytes/sec, Disk write bytes/sec, average disk reads/sec, average disk writes/sec
- *Experiment:* Resize a directory of 1000 images totaling 4.5GB from current resolution to 1024x768 into a new directory.
- *Conjecture:* Storage devices with PCIe interface have higher peak bandwidth and lower average latency.

Component & Synthetic

- *Conditions:* Generate sequential disk reads and writes to create a sustained average queue length over 30 seconds for queue lengths of 0, 2, 4, 8, 16, 32.
- *Metrics:* Average Queue Length, Disk reads/sec, Disk writes/sec, Disk read bytes/sec, Disk write bytes/sec
- *Experiment:* Take read throughput measurements while target drive operates at various queue lengths.
- *Conjecture:* Peak read throughput is achieved at lower queue lengths.

### 6.3 Metrics

Most commonly one will look at the amount of **time** a workload took to run, or the **score** provided by a benchmark tool.

To perform an in-depth analysis consider capturing the below metrics from the Operating System.

Metrics	Category	Notes
Disk Writes/Sec	Disk	Write IOPs
Disk Reads/Sec	Disk	Read IOPs
Avg Disk Sec/Write	Disk	Write Latency
Avg Disk Sec/Read	Disk	Read Latency
Avg Disk Write Queue Length	Disk	Indicator of Queue Depth
Avg Disk Read Queue Length	Disk	Indicator of Queue Depth
Write Bytes/Sec	Disk	Write Throughput
Read Bytes/Sec	Disk	Read Throughput
Queue Length	CPU	# of threads delayed; waiting to be executed
% DPC Time	CPU	Deferred procedure calls due to interrupts
% Interrupt Time	CPU	Hardware interrupts
% Privileged Time	CPU	Operating System and Hardware drives etc.
% Processor Time	CPU	Application or Operating System processes
% User Time	CPU	Applications
Available (Mega)Bytes	Memory	As percentage of total memory
Cache Bytes	Memory	Portion of cache resident in memory



Committed Bytes	Memory	Committed memory is the physical memory in use for which space has been reserved in the paging file should it need to be written to disk.
Pages/Second	Memory	The number of pages read from the disk or written to the disk to resolve memory references to pages that were not in memory at the time of the reference. The sum of Pages Input/sec and Pages Output/sec.
Page Writes/Sec	Memory	Rate pages are written to disk
Page Reads/Sec	Memory	Rate pages are read from disk to resolve page faults

### 6.3.1 Providers

- [Windows Performance Counters](#): Provide information as to how well the operating system, or an application, service or driver is performing.
- [S.M.A.R.T.](#): Self-Monitoring, Analysis and Reporting Technology is a monitoring system included in storage devices.
- [Windows Management Instrumentation](#): is an implementation of a standard to access management information (computing assets) such as systems, applications, networks, device and other managed components. Although WMI is largely considered a configuration management system built into the Operating System, it also includes some capability to access performance metrics. See [WBEM](#) and [CIM](#) for related information.

### 6.3.2 Key Terms and Concepts for Storage Metrics

- **IOPS** (Throughput): Input/output operations per second; can be characterized as either sequential or random; also, as read or write. Is a measure of throughput as operations.
- **Queue Depth** (Concurrency): represents the number of outstanding Input/output tasks waiting to be serviced by the disk, essentially the queue of disk work.
- **Read or Write Bytes per Second** (Throughput): Rate at which bytes of data are transferred. Is a measure of throughput as cumulative size of payloads.
- **Sequential and Random**: Reads and writes to storage can be characterized as either random or sequential and relates to how the data is stored and addressed as blocks on the device. Simply, random typically refers to the addresses of data read/written as having a large span while sequential means the data read/written is stored in sequential addresses.



### 6.3.3 Recorders and Analyzers

If you are using time as your metric, then a simple stopwatch will service the test. Benchmarks that present a score have within them tools to record, analyze and derive that score. For those performing an in-depth analysis we recommend the tools in the table below.

Name	Usage	Description
<a href="#">Windows Performance Recorder</a>	In-depth tracing	Performance recording tool that is based on <a href="#">Event Tracing for Windows (ETW)</a> . Replaces Microsoft XPERF.
<a href="#">Windows Performance Analyzer</a>	In-depth analysis	Performance Analyzer to interpret Windows Performance Recorder results and provides optimization recommendations
<a href="#">Microsoft PerfView</a>	In-depth tracing and analysis	PerfView is a CPU and memory performance-analysis tool.
<a href="#">Intel VTune Platform Profiler</a>	High level system tracing and analysis	Platform Profiler is a lightweight standalone tool providing rapid insights into overall system configuration, performance and behavior, with specific focus on identifying platform-level memory and storage bottlenecks and imbalances
<a href="#">Intel Storage Performance Snapshot</a>	High level system tracing and analysis	Analyzes storage, CPU, memory and network usage and displays basic performance enhancement opportunities
<a href="#">Windows Performance Monitor</a>	High level system tracing and analysis	Monitors performance on Windows-based Operating Systems