

# Intel<sup>®</sup> Optane<sup>™</sup> Solid State Drives for Client

---

***Evaluation Guide***

January 2018



## 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>	October 2017
002	<ul style="list-style-type: none"><li>Generalized document from 900P-specific, minor updates, added Appendix D</li></ul>	January 2018

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. Results have been estimated based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

All documented performance test results are obtained in compliance with JESD218 Standards; refer to individual sub-sections within this document for specific methodologies. See [www.jedec.org](http://www.jedec.org) for detailed definitions of JESD218 Standards.

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.

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

Intel, Intel Optane, 3D XPoint, and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

\*Other names and brands may be claimed as the property of others.

Copyright © 2018 Intel Corporation. All rights reserved.



# Contents

	Revision History .....	2
	Terms and Acronyms .....	4
	Related Documentation .....	4
	Related Tools .....	4
1	Overview.....	5
1.1	Intel® Optane™ Technology Overview.....	5
1.2	Evaluation Planning.....	6
2	System Tuning For Optimum Performance.....	7
2.1	Platform Configuration and Settings: .....	7
2.2	Operating System Settings .....	7
2.3	Operating System Preparation Steps .....	9
2.4	Device Preparation Steps .....	9
3	Appendix A – Benchmarking Considerations.....	10
3.1	Application-based Benchmarks .....	10
3.2	Trace-based Benchmarks.....	10
3.3	Synthetic Benchmarks .....	11
4	Appendix B – Why use Intel NVMe Driver for Windows OS? .....	12
5	Appendix C – Operating System Optimization.....	13
5.1	Power Scheme Settings.....	13
5.2	Indexing Service .....	15
5.3	Scheduled Defragmentation .....	16
5.4	System Protection (System Restore).....	18
5.5	Paging File (Swap File) .....	19
5.6	Prefetch and Superfetch .....	21
6	Appendix D – CrystalDiskMark Benchmark Example .....	23

## Tables

Table 1:	Recommended BIOS Settings .....	7
Table 2:	Recommended OS Settings .....	8



## Terms and Acronyms

Term	Definition
AHCI	Advanced Host Controller Interface
ATA	Advanced Technology Attachment
GB	Gigabyte
HDD	Hard Disk Drive
KB	Kilobytes
I/O	Input/Output
IOPS	Input/Output Operations Per Second
MB	Megabytes
NVMe	NVM Express / Non-Volatile Memory Express
PCIe	PCI Express / Peripheral Component Interconnect Express
SATA	Serial Advanced Technology Attachment
SSD	Solid State Drive

## Related Documentation

Title	Download Link
Intel® SSD 750 Series Evaluation Guide (IOMeter* use)	<a href="https://www.intel.com/content/dam/support/us/en/documents/ssdc/hpssd/ssd-750/sb/intel_ssd_750_series_evaluation_guide_332075002us.pdf">https://www.intel.com/content/dam/support/us/en/documents/ssdc/hpssd/ssd-750/sb/intel_ssd_750_series_evaluation_guide_332075002us.pdf</a>
SNIA Solid State Storage (SSS) Performance Test Specification Client	<a href="https://www.snia.org/tech_activities/standards/curr_standards/pts">https://www.snia.org/tech_activities/standards/curr_standards/pts</a>

## Related Tools

Product	Download Link
Intel® NVMe Driver download	<a href="https://downloadcenter.intel.com/download/27231/NVMe-Drivers-for-Intel-SSDs">https://downloadcenter.intel.com/download/27231/NVMe-Drivers-for-Intel-SSDs</a>
Intel® SSD Data Center Tool	<a href="http://www.intel.com/support/ssd">http://www.intel.com/support/ssd</a>
Intel® SSD Toolbox	<a href="http://www.intel.com/go/ssdtoolbox">http://www.intel.com/go/ssdtoolbox</a>
IOMeter*	<a href="http://www.iometer.org/">http://www.iometer.org/</a>



# 1 Overview

---

This guide outlines the best known practices for evaluating Intel® Optane™ Solid State Drives products for client systems in a Microsoft Windows 10\* environment, and considerations for comparing Intel® Optane™ SSDs to NAND-based SSDs. This is not a detailed guide to achieve specific benchmark results with IOMeter\*. For benchmark information, refer to links in the [Related Documentation](#) table.

Intel® Optane™ SSD products for client systems include the Intel® Optane™ SSD 900P Series and other client storage products based on Intel® Optane™ Technology. The Intel® Optane Memory system acceleration solution for client and the Intel® Optane™ SSD DC P4800X Series are not covered in this document.

The primary audience for this document is publications, OEMs, system integrators, technical analysts, and individuals who plan to evaluate an Intel® Optane™ SSD's performance benefits and features.

## 1.1 Intel® Optane™ Technology Overview

Intel® Optane™ Technology features 3D XPoint™ memory media. 3D XPoint™ memory media is a new class of memory technology and the resulting Intel® Optane™ SSD products have characteristics that are fundamentally different from SSD products based on traditional NAND media. Intel® Optane™ technology pairs 3D XPoint™ memory media with Intel Memory or Storage Controller, Intel Interconnect IP (system components architected by Intel), and Intel storage software to deliver products highly optimized for specific markets and workloads demanding large capacity, high endurance and fast storage.

3D XPoint™ memory media features a cross point structure which allows each memory cell to store a single bit of data which can be accessed individually. This is fundamentally different from NAND media, which is read and written in pages, or "large chunks" of data. 3D XPoint™ memory media uses fast switching material, combined with a low-delay metal interconnect design, delivering a significant performance boost, when compared to NAND. Together, these building blocks deliver a combination of performance, high endurance and quality of service, along with low latency. Because of the unique characteristics of SSD products based on Intel® Optane™ technology, some scenarios where they are expected to outperform NAND-based SSDs are:

- High random read and write performance.
- High performance at low queue depths.
- High read performance under write load.
- High performance at small capacity points.
- High performance maintained over time and as the drive fills.
- High endurance.

With their unique high performance characteristics, the performance of Intel® Optane™ SSDs can be impacted by some variables that NAND SSDs are less impacted by. These variables include CPU speed, and relevant hardware and OS settings. This guide introduces those variable that enable users to evaluate optimum performance with Intel® Optane™ SSDs.



## 1.2 Evaluation Planning

How do you define performance? Evaluating the performance of an SSD starts with understanding which storage performance characteristics are important to you. Sequential read performance, random write performance, low latency, simultaneous read/write operations, or some mix? Which factor is most important? It's often difficult to know without extensive analysis. The more you know about your workload(s) and performance bottlenecks, the easier it will be to determine the best SSD for your needs.

Perhaps the first and most important question about your workload is: Where is the performance bottleneck? Is your workload sensitive to storage performance, or is it limited by other factors such as processor frequency, processor core count, or DRAM? One method to explore the performance considerations of your workload is to capture and analyze your workload with the Intel® Storage Performance Snapshot tool:

<https://software.intel.com/sites/products/snapshots/storage-snapshot/>. If your workload is sensitive to storage performance, then take a closer look using benchmark testing and workload analysis.

Benchmarks are designed to mimic a particular type of workload on a component or system and provide an indicator of performance. The complexities of meaningful benchmarking are often masked with “easy” tools and generalities. Easy-to-run benchmark tools can be a convenient way to provide quick guidance on certain performance characteristics of an SSD, but should be thought of as only an initial step in meaningful evaluation. IOMeter is a common tool used for in-depth performance analysis. While it allows you to test a wide variety of test scenarios, it requires more effort to design and execute the tests. Whichever tool you use, remember that benchmarks are intended only as a guide to SSD performance and may not demonstrate the value (or lack of) of certain performance characteristics on your workload.

Queue Depth (QD) is an important factor to understand when planning your benchmark testing. QD is defined as the number of outstanding I/O requests at the storage controller. As a carry-over from test tools and practices designed for hard-disk drives (HDDs), it is common to see benchmark tool configurations and storage device performance specifications listing random performance at a queue depth of 32 (QD32). While your specific workload may vary, Intel-internal and 3<sup>rd</sup>-party analysis of various real-world client workloads shows that for general use in a Windows environment, the workload is heavily random, read/write mixed, at low queue depths. To realistically test drive performance and see how it can really influence the end user, testing should to be done at low queue depths. Intel recommends including testing random read and write performance at queue depths between 1 and 4. If your tool supports the capability, testing with a mixed read/write workload at low QD is the most realistic.

While benchmarks are a good place to start, real-world workload testing and analysis is the recommended approach to understand the true impact of an SSD on your workload with your configuration. Running a real task with your configuration in a real environment is the best way to understand how a particular SSD can impact workload performance. The downside to this approach is that it can be time-consuming, expensive or difficult to set up, execute, and analyze. While this approach may require more effort, because of complex system interactions, it is usually the best way to really understand certain workloads or configurations.



## 2 System Tuning For Optimum Performance

In order to obtain consistent and repeatable data a stable and deterministic environment is necessary. To ensure stability and the deterministic nature of the testing environment, it is important to constrain the variables (platform, device, and OS settings, etc.). The variables are categorized as:

- Platform configuration and settings
- Operating System settings
- Operating System preparation steps
- Device preparation steps

### 2.1 Platform Configuration and Settings:

Configure these settings in the platform BIOS before deploying an OS. The settings usually address, but are not limited to, the platform, CPU and device power states, as well as other CPU features (ie. Hyperthreading).

- CPUs with a frequency of 3.0GHz and compatible with X99, Z170, Z270, X299 or later generation chipset
- Ensure the Intel® Optane™ SSD is connected to the processor's PCIe Gen 3.0 port and not the chipset's PCIe port. Consult your motherboard manual to identify the correct PCIe slot(s).
- Update your motherboard to the latest UEFI BIOS.
- Whenever possible, test the Intel® Optane™ SSD as a secondary drive (install the OS onto a separate drive).
- BIOS setting recommendations will vary based on the purpose of your testing. When testing a real-world workload, refer to the table below. Disable Hyper-threading, EIST, and Intel Turbo Mode when setting up a platform for synthetic benchmark testing in which you desire the most stable, consistent, and repeatable results across platforms.

**Table 1: Recommended BIOS Settings**

BIOS mode	UEFI
Hyper-threading	Enabled*
EIST (Enhanced Intel Speed Step Technology)	Enabled*
Intel Turbo Mode	Enabled*
PCIe ASPM (Active State Power Management)	Disabled
C-States	Disabled
P-States	Disabled

\* Disable this setting for stable, consistent, and repeatable results, especially across platforms.

### 2.2 Operating System Settings

Set the desired OS settings before reading the device and installing the benchmark. These desired settings usually address the disabling of background processes, changing priority for IO, and ensuring that there aren't any unfinished tasks running that could invalidate benchmark results. When setting up the OS for benchmarking, configure these general settings to ensure stable, consistent, and repeatable results.

- If possible, Intel recommends a clean OS install on the machine(s) that will be used for testing/benchmarking.
- Recommended OS: Windows 10\* or later versions, x64.
- Install all chipset drivers for the motherboard and all peripheral drivers.
- Run Windows\* Update to ensure all drivers are current.



The following OS settings ensure stable, consistent, and repeatable results during benchmarking:

**Table 2: Recommended OS Settings**

Setting	Status
Power Scheme	High Performance mode
Power Options > Turn off hard disk after	0 Minutes (Never)
Power Options > Internet Explorer > JavaScript Timer Frequency	Maximum Performance
Power Options > Desktop background settings > Slide show	Paused
Power Options > Wireless Adapter Settings > Power Saving Mode	Maximum Performance
Power Options > Sleep > Sleep after	0 Minutes (Never)
Power Options > Sleep > Allow wake timers	Enabled
Power Options > USB settings > USB selective suspend setting	Disabled
Power Options > Intel Graphics Power Plan	Maximum Performance
Power Options > Power buttons and lid > Power button action	Do nothing
Power Options > Power buttons and lid > Sleep button action	Do nothing
Power Options > PCI Express > Link Power Management	Off
Power Options > Processor power management > Minimum processor state	100%
Power Options > Processor power management > System cooling policy	Active
Power Options > Processor power management > Maximum processor state	100%
Power Options > Display > Turn off display after	0 Minutes (Never)
Power Options > Display > Display Brightness	100%
Power Options > Display > Dimmed display brightness	100%
Power Options > Display > Enable adaptive brightness	Off
Indexing Service	Disabled
Scheduled Defragmentation	Disabled
System Protection (System Restore)	Disabled
Paging File (Swap File)	Disabled
Prefetch and Superfetch	Disabled
Hibernate and Sleep Mode	Disabled
Screen Saver	Disabled
Windows Defender	Disabled
Automatic Windows Updates	Disabled
Virus Protection	Disabled






## 2.3 Operating System Preparation Steps

These steps will detail the procedure necessary in order to prepare the OS so that it is in a true idle state before executing the measurements. Once the OS is installed and configured properly there are further steps necessary to ensure that the system is ready to execute performance measurements. The following steps should be followed to ensure a proper starting state:

- Close all applications that are not necessary for measurement.
- Issue ProcessIdleTask command to force OS to process any unfinished background tasks
  - `Start /wait Rundll32.exe advapi32.dll,ProcessIdleTasks`



- After the ProcessIdleTasks command has completed allow the system to idle for 15 minutes to ensure a true idle.

## 2.4 Device Preparation Steps

These steps will detail the procedures necessary in order to prepare the devices so that they are updated and in a steady state before executing the measurements.

- Verify that the Intel® Optane™ SSD has the latest firmware applied.
- Use the latest Intel NVMe driver:  
<https://downloadcenter.intel.com/download/27231/NVMe-Drivers-for-Intel-SSDs>  
The storage device should be partitioned and formatted in NTFS.
- Ensure a minimum power-on time of 3 hours before running performance evaluations to ensure that background data refresh is not abnormally active after an extended de-energized state.
- Common steps when testing NAND-based SSDs which are **not required** when testing SSDs based on Intel® Optane™ Technology:
  - Low Level Format the SSD in order to ensure a consistent starting point for measurements.
  - Contiguously fill and run with the first workload for a short amount of time.

## 3 Appendix A – Benchmarking Considerations

---

Benchmarks are designed to mimic a particular type of 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

### 3.1 Application-based Benchmarks

Application-based benchmarks run real-world programs on a system and measure the application performance as it uses the system.

Application-based benchmarks emulate end-user usage with the help of 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. Individual subsystem scores are created by grouping those applications that are common to a type of user and are then reported as subsystem scores.

SYSmark\* 2014 SE is an example of an application-based benchmarking tool. It provides results in the form of four sub scores and one overall score.

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.

Application-based benchmarks may also result in a single measure for system performance, which can be used as a product ranking index.

One disadvantage of application-based benchmarks when trying to determine storage device speed is that they traditionally focus on CPU, memory, and graphics performance and may not properly weigh storage sub-system speed in the final results. They also may not take into account 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 that 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, as a practical basis, these benchmarks will not be representative of storage device usage over time.

### 3.2 Trace-based Benchmarks

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

PCMark Vantage\* HDD test is an example of a trace-based benchmarking tool. Trace-based benchmarking has many of the advantages of application-based benchmarking, when 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. It is also possible to run long traces, so performance over time can be measured.

One disadvantage of trace-based benchmarking is that third-party tools required for testing are not as readily available as those for synthetic and application-based benchmarking. Additionally, 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 storage 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 taking into account 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.

Examples of synthetic benchmarking tools include IOMeter\*, CrystalDiskMark\* and HD Tune.

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 and HD Tune are considered fixed, short duration synthetic benchmarks.
  - If using CrystalDiskMark, see the “Queues & Threads” options within the “Settings” menu to explore the impact of various queue depths and thread counts on test results. A suggested starting point to explore would be Queue Depths = 1-4, Thread Count = 8 (with a Core i7 processor with Hyper-threading enabled). See [Appendix D](#) for more information.



## 4 Appendix B – Why use Intel NVMe Driver for Windows OS?

---

There are two primary storage driver choices to use with Intel® Optane™ SSDs; the Windows in-box NVMe storage drivers and the Intel NVMe storage drivers. We recommend using the Intel NVMe storage drivers.

There are three distinct areas that allow Intel to provide a better experience with our Intel NVMe storage drivers for the Windows operating system.


1. Intel's device drivers are architected, designed and implemented to integrate seamlessly with other Intel SSD software products such as Intel® SSD Toolbox and Intel® Data Center Tool which are unique products for managing and monitoring Intel SSD's.
2. Intel device driver debug efforts will be more responsive, meaning that customer or OEM specific support requests will be handled in time critical manners, since Intel's device drivers are not tied to any operating system release cadence.
3. New Intel® SSD features are developed jointly across the I/O stack from hardware, firmware, and host software (device driver and tools) which allows Intel to release device drivers with focused development and validation efforts and on a cadence that is specific to Intel's SSD's, instead of infrequent operating system releases.

## 5 Appendix C – Operating System Optimization


This appendix provides tips on how and where to optimizations the operating system settings listed in [Section 2.2](#).

### 5.1 Power Scheme Settings

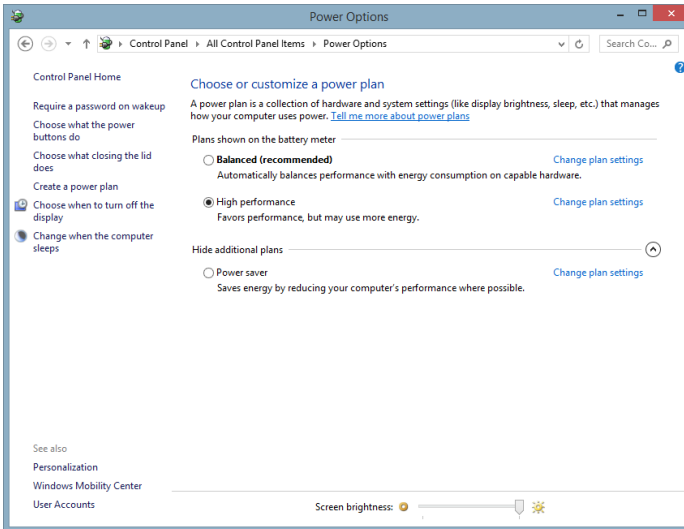
1. Right click on the **Start Button** and select **Control Panel**



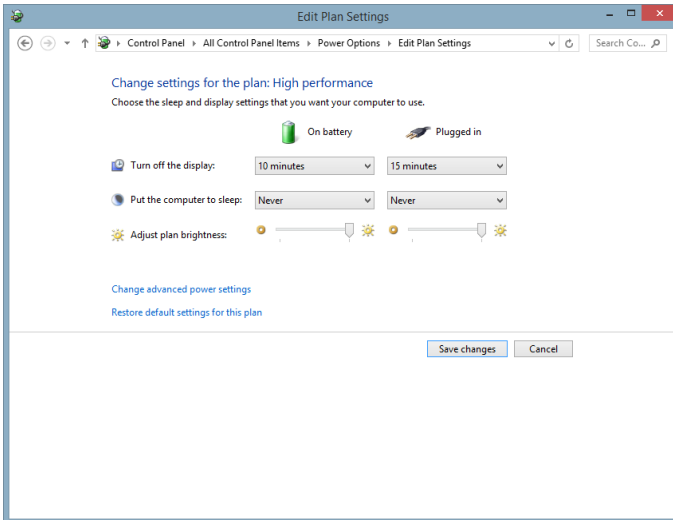
2. Click on **View by** and select **Small Icons** then select **Power Options**



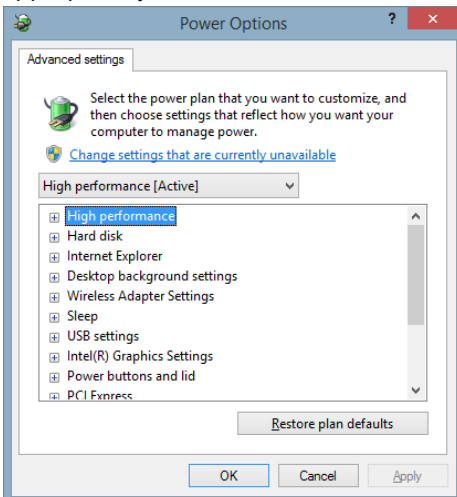
3. Once in Power Options select **High Performance** power plan



4. Click on **Change plan settings**, set both **Turn off the display** and **Put the computer to sleep** to **Never**. (Typical test system will not have On Battery sections)

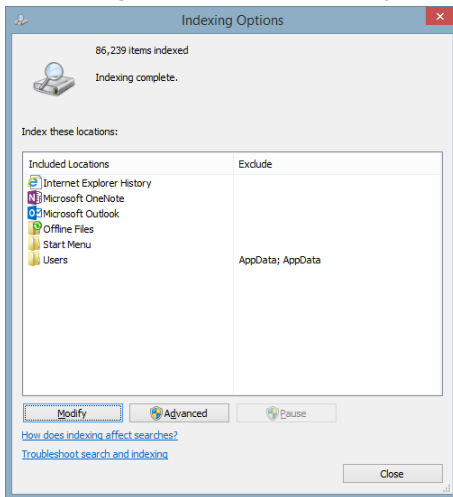


5. Click on **Change advanced power settings** and use the table at the beginning to change settings appropriately.

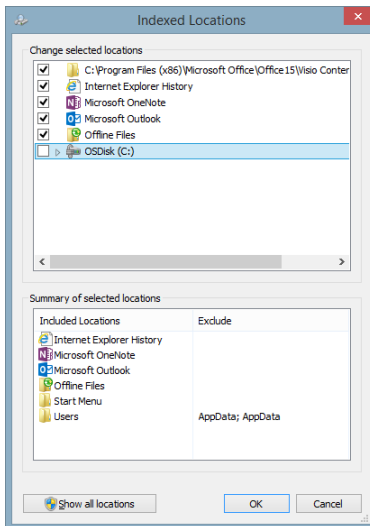


## 5.2 Indexing Service

1. Open Windows Search and search for **Indexing Options**
2. In **Indexing Options** click on **Modify**

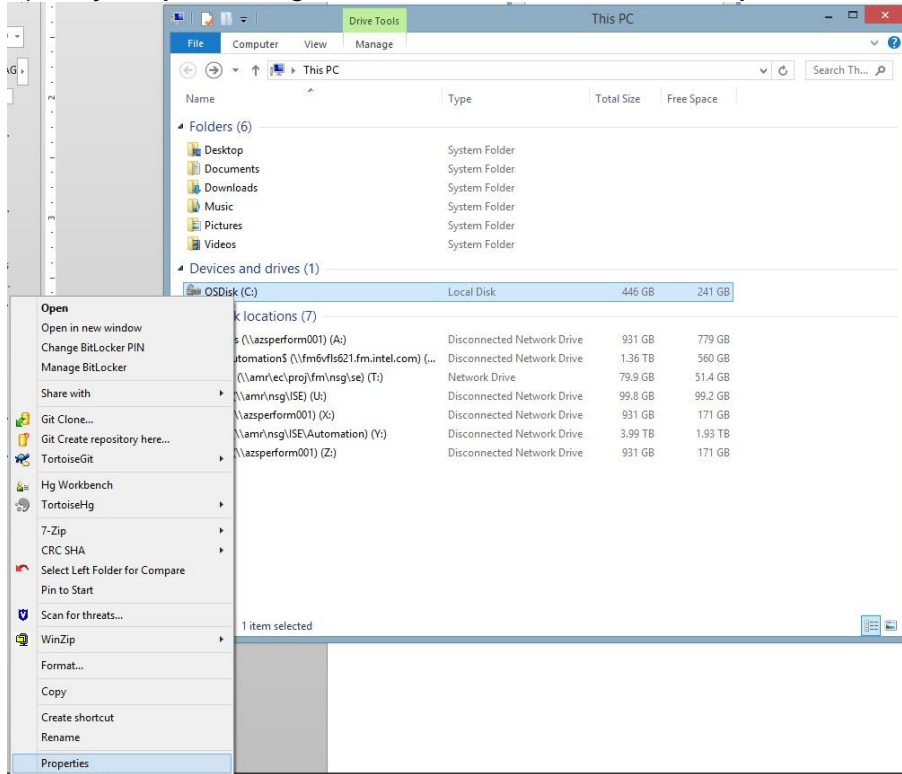


3. In the **Indexing Locations** window make sure to **uncheck** the drive that is being tested



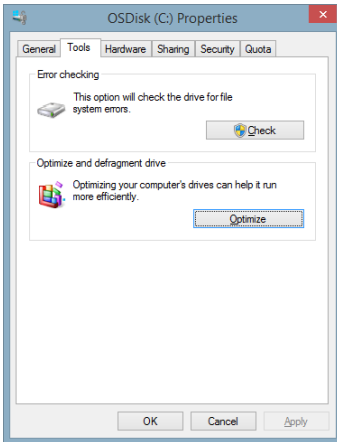
## 5.3 Scheduled Defragmentation

1. Open **My Computer** and right click on the test drive, then select **Properties**

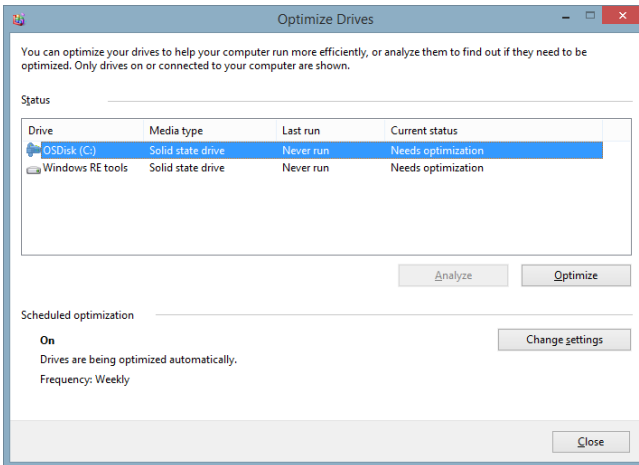




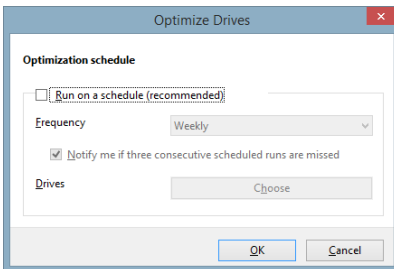
- In the **Properties** window select the **Tools** tab and click on **Optimize**



- In the **Optimize Drives** window select the test drive and click on **Change Settings**.




- In the Optimize Drive window uncheck **Run on a schedule**



## 5.4 System Protection (System Restore)

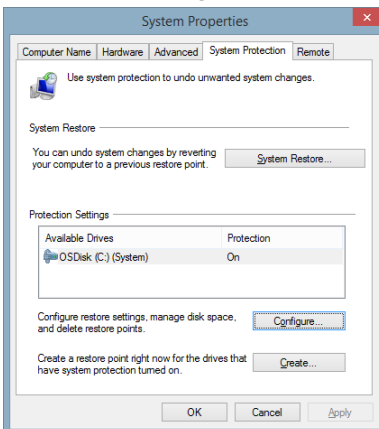
1. Right Click on **Start** button and select System



2. In **System** window click on **System protection**

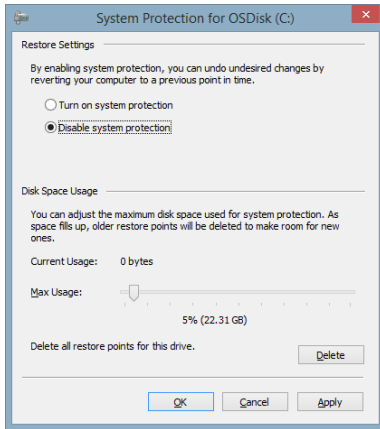


3. Click on the **Configure** button






4. In the **System Protection** prompt select **Disable system protection**

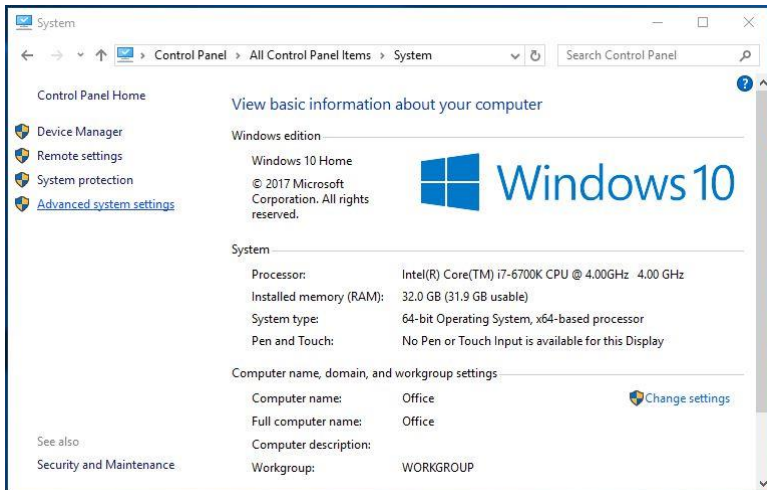


## 5.5 Paging File (Swap File)

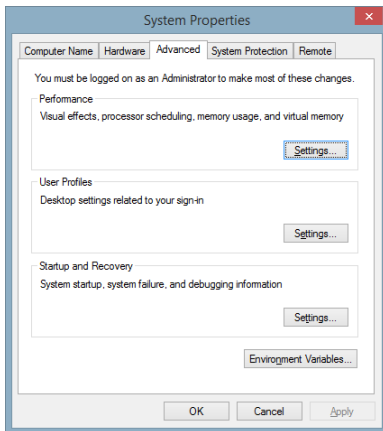
1. Right click on **Start** button and select **System**



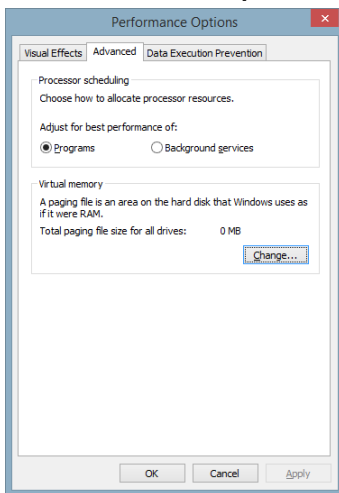
2. In **System** window click on **Advanced system settings**



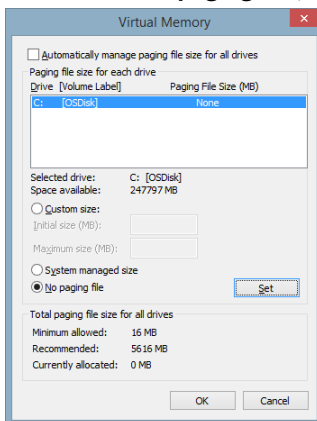
3. In the **System Properties** window click on the **Settings** button under the **Performance** section.



4. In the **Performance Options** window click on the *Advanced* tab




5. Click the **Change...** button.  
In the resulting window uncheck **Automatically manage paging file size for all drives**, select the test drive, select **No paging file**, and click the **Set** button.




## 5.6 Prefetch and Superfetch

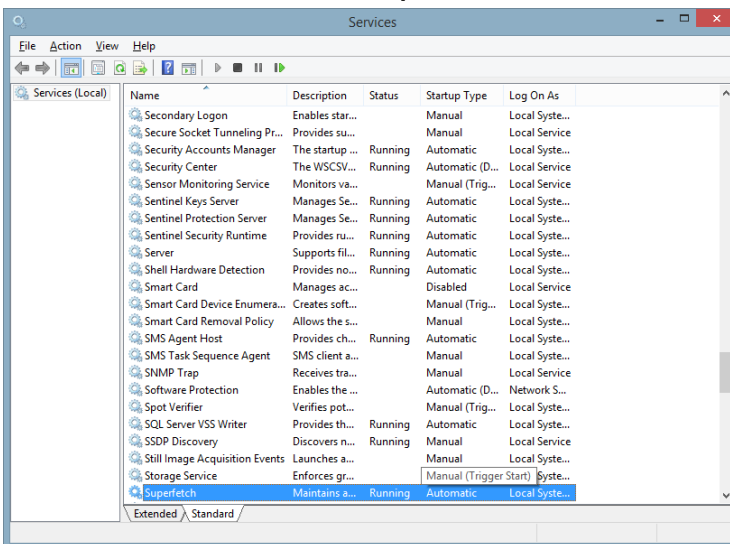
1. Right click on **Start** button and select **Run**



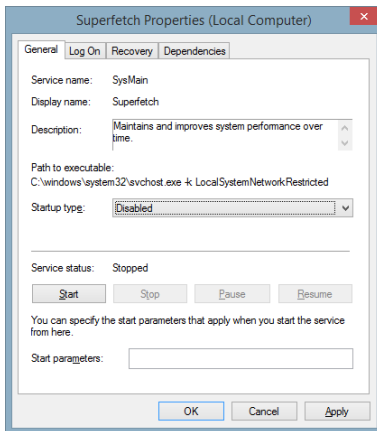
2. In the **Run** prompt enter **services.msc** command and click **OK**




3. In the **Services** window select the **Superfetch** service and hit **Enter**




- In **Superfetch Properties** click on **Stop** and change **Startup type** to **Disabled**, then click **OK**.



- To disable prefetch open the **Run** prompt again and type in the **regedit** command and click **OK**



- In the **Registry Editor** window navigate to the following key:  
**HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management\PrefetchParameters**

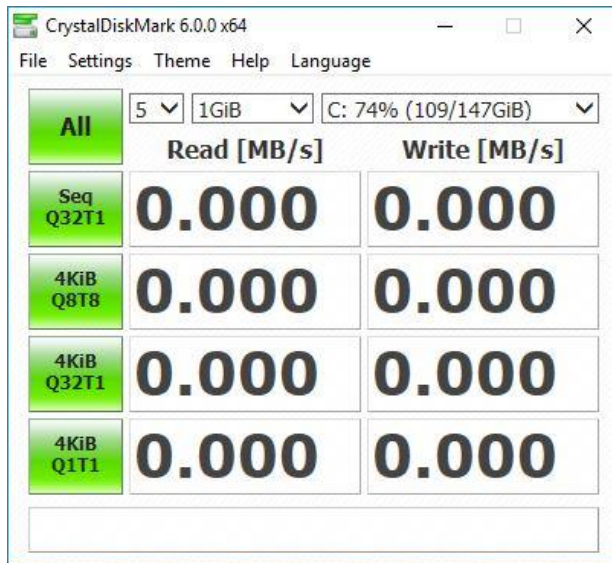


Double-click on **EnablePrefetcher** and change the value to **0** (zero) to disable it.



## 6 Appendix D – CrystalDiskMark Benchmark Example

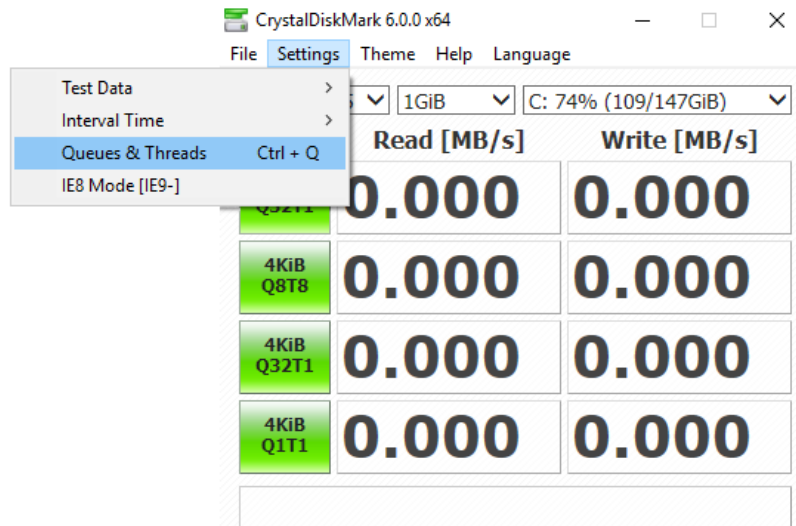
CrystalDiskMark 6.0 has default settings which are a good starting point for testing with Intel® Optane™ SSDs. Just click “All” to run with default settings.




### 6.1 Queue Depths and Thread Count

No significant variations in results are expected by changing the **Test Data** or **Interval Time** parameters. Select the queue depth and number of threads you would like to test. It is suggested to start with Thread Count to equal the amount of physical plus virtual cores available on the processor.

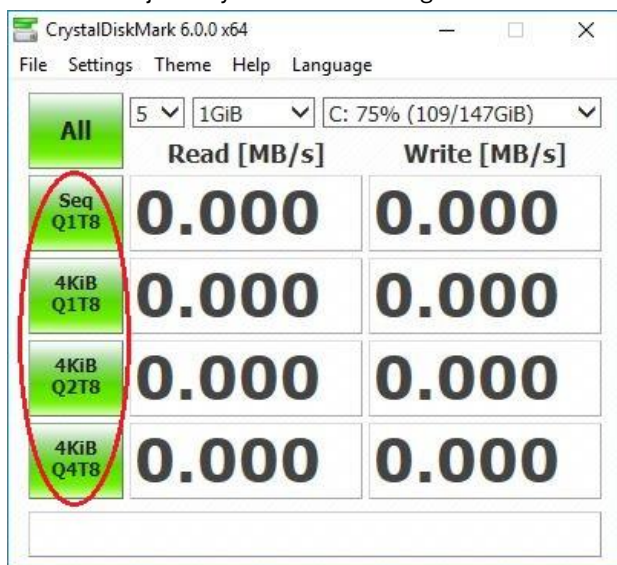
1. Click **Settings** and select **Queues & Threads**



- In the below example, different queue depths were chosen, all paired with 8 threads (Intel® Core™ i7-7700: 4 physical cores, 4 virtual cores).



- Close out of **Queues & Threads** by clicking on the 'X' (not by clicking on the 'Default' button), the test labels will adjust to your chosen settings.



- Click **All** to run through all 4 tests.  
A tip on reading "IOPS" – the results are presented in MB/s. You can see the random performance in IOPS by saving the results (Ctrl-S) and reviewing the text file, or by hovering your mouse over the MB/s number.