



# INTEL<sup>®</sup> VIRTUAL RAID ON CPU (INTEL<sup>®</sup> VROC)

Public Training Guide

# Legal Notices and Disclaimers

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

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.

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 computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at [intel.com](https://www.intel.com).

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit [www.intel.com/benchmarks](https://www.intel.com/benchmarks).

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

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.

Intel, the Intel logo, Intel Optane, Xeon, and others are trademarks of Intel Corporation in the U.S. and/or other countries.

© Intel Corporation.

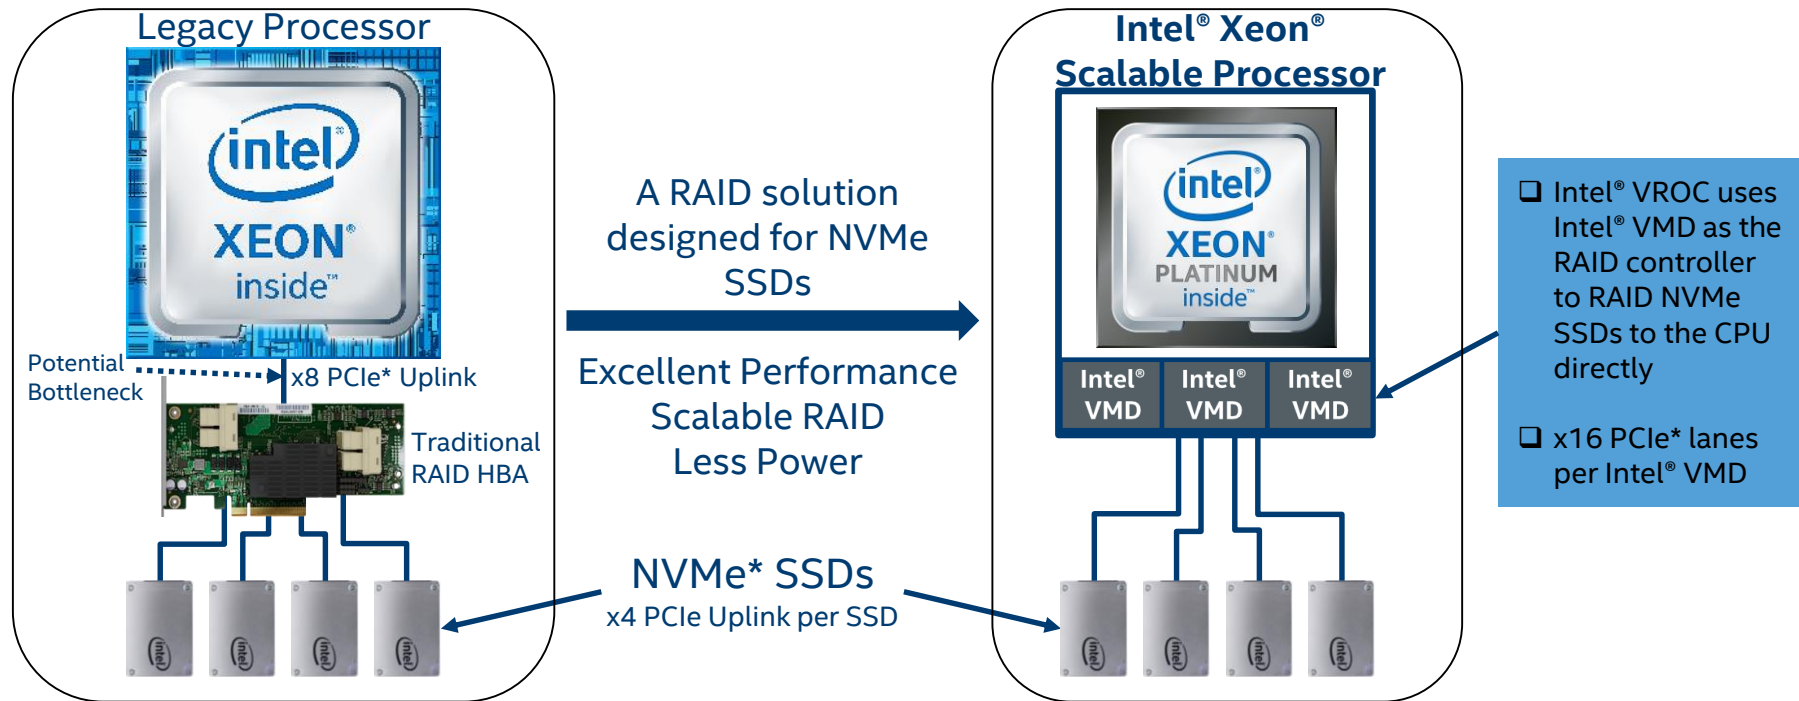
\*Other names and brands may be claimed as the property of others.

# Agenda

- Intel® VROC: Features and Benefits
- RAID Configurations
- Performance
- Intel® VROC Public Resources

# FEATURES AND BENEFITS

# Intel® Virtual RAID on CPU (Intel® VROC)



Intel® VROC provides compelling RAID solution for NVMe\* SSDs

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

# Intel® VROC Major Features

## Major Features

**Bootable RAID** – Build redundancy to protect your system volume

**Hot Insert and Surprise Removal** – Expand volume, replace bad drive without system reboot

**LED Management** – Visually display RAID status via LEDs

**Close RAID5 Write Hole** – Maintain data integrity with no battery backup unit

**RAID management** – Remote management via webpage or RESTful APIs, pre-OS management via UEFI HII and CLI, OS management via GUI and CLI

**3rd Party vendor SSD support** – Validation, WHQL and support for selected 3<sup>rd</sup> party SSDs

**OS support** – Windows\* and Linux\*. Intel® VROC for Linux is open-source, builds on MDRAID

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

# Intel® VROC is Hybrid RAID (HW + SW)

Major RAID Features	HW RAID	Intel® VROC	SW RAID	Intel® VROC Comment
Error Handling Isolation	✓	✓	X	Intel® VMD isolates SSD error/event handling from OS to reduce system crash or reboot due to error
Reliable data storage	✓	✓	X	Protect RAID 5 data even when power loss occur while volume is degraded
Boot support	✓	✓	X	Provide protection to system volume
Management	✓	✓	X	UEFI, GUI, CLI, RESTful API, remote web
Dedicated I/O processor for RAID	✓	1	X	1=Uses powerful Intel® Xeon® CPU to RAID the unprecedented fast NVMe* SSD
Protected write back cache	✓	2	X	2=Use the power loss protected feature inside enterprise SSDs
Easily upgraded	X	✓	✓	Software update vs new HW purchase
Less hardware required	X	✓	✓	No need for HBA, cable, battery backup unit. Save power. Save precious PCIe* lanes for others.

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

# Intel® VROC SKUs

		Pass-thru SKU	Standard SKU	Premium SKU	Intel® SSD SKU
	<b>MM#</b>	<b>NA</b>	<b>951605</b>	<b>951606</b>	<b>956822</b>
NVMe* RAID	CPU attached NVMe	✓	✓	✓	✓
	3 <sup>rd</sup> party SSD support	✓	✓	✓	-
	Hot-Plug/ Surprise Removal	✓	✓	✓	✓
	LED management	✓	✓	✓	✓
	Web UI management	✓	✓	✓	✓
	Bootable RAID	-	✓	✓	✓
	RAID 0/1/10	-	✓	✓	✓
	RAID 5	-	-	✓	✓
	RAID Write Hole closed (BBU replacement)	-	NA	✓	✓
	SATA RAID	Bootable RAID Volume	✓	✓	✓
RAID 0/1/5/10		✓	✓	✓	✓

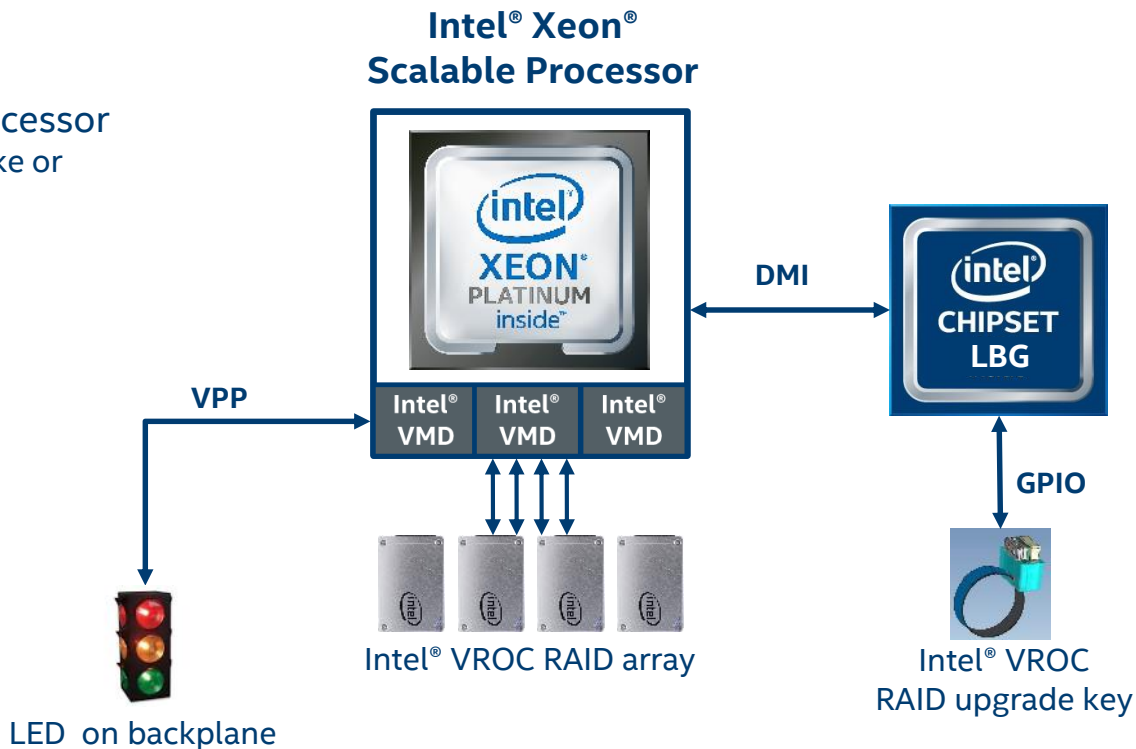
\*Other names and brands may be claimed as the property of others.



# CONFIGURATION

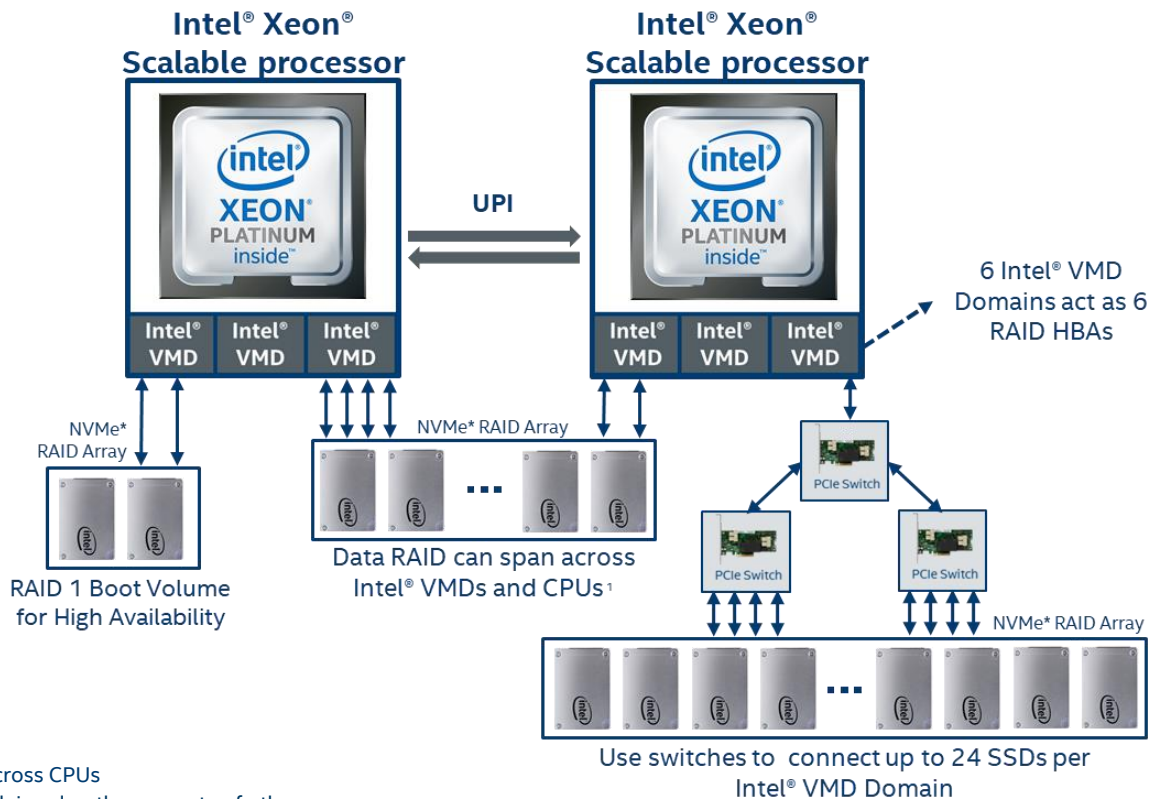
# Intel® VROC Checklist

- Intel® Xeon® Scalable Processor
  - Gen1 or Gen2 (i.e. Skylake or Cascade Lake)
  - -SP, -D, or -X versions
- VMD BIOS enablement
- Virtual Pin Port (for LED)
- VROC HW key socket
- Supported OS
  - RHEL 7.3+
  - Suse 12 SP3+
  - Win2012 R2+
  - Win10+



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

# Supporting RAID Configurations



<sup>1</sup> May incur performance penalty across CPUs


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

**PERFORMANCE**

# Performance – RAID vs Pass-thru

## RHEL7.4 with Intel® SSD DC P4510<sup>1</sup>

### (4k Random)



- Pass-thru raw data:
  - 4k Rand Write: 84k IOPS
  - 4k Rand Mixed: 183k IOPS
  - 4k Rand Read: 645k IOPS
- 4-Disk RAID0 Read: 2.5M IOPS
- Physical CPU Cores Used:
  - 4-Disk RAID0 Read: 4.7 Cores
  - 4-Disk RAID5 Write: 1.2 Cores

**52 total physical cores on this 2 socket,  
Intel® Xeon® 8170 based system**

See appendix for footnotes


Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit [www.intel.com/benchmarks](http://www.intel.com/benchmarks).

Performance results are based on testing as of October 5, 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

# Performance – RAID vs Pass-thru

## RHEL7.4 with Intel® SSD DC P4510<sup>2</sup>

(128k Seq., 1 Worker)



- Pass-thru raw data:
  - 128k Seq. Write: 1.7GB/s
  - 128k Seq. Read: 2.7 GB/s
- 4-Disk RAID 0 Read: 6.8 GB/s
- Physical CPU Cores Used:
  - 4-Disk RAID0 Read: 0.3 Cores
  - 4-Disk RAID5 Write: 1.0 Cores

**52 total physical cores on this 2 socket,  
Intel® Xeon® 8170 based system**

See appendix for footnotes.

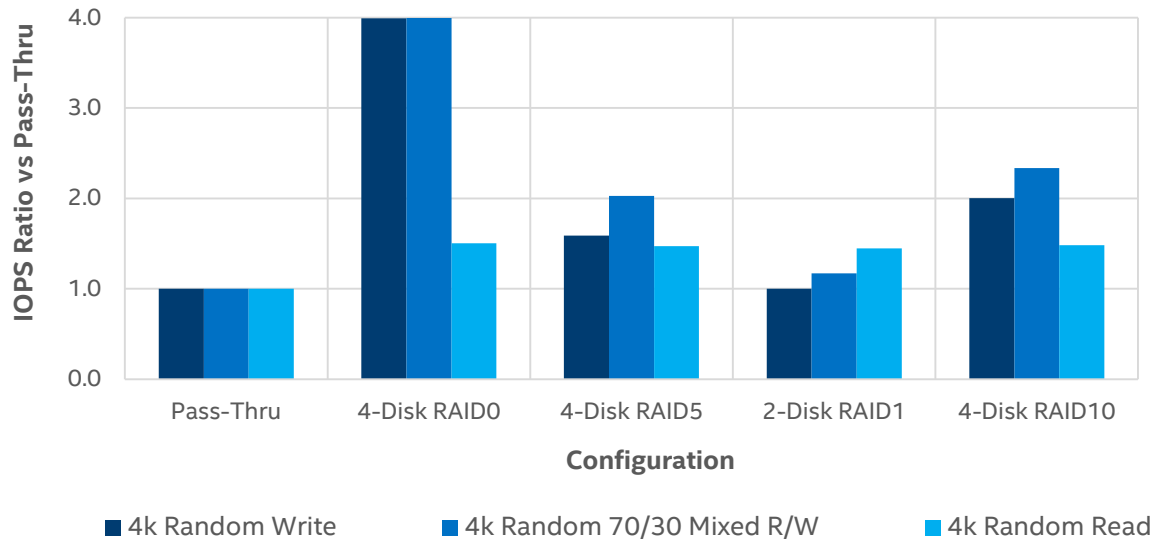
Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit [www.intel.com/benchmarks](http://www.intel.com/benchmarks).

Performance results are based on testing as of October 5, 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

# Performance – RAID vs Pass-thru

## Windows\* 2016 with Intel® SSD DC P4510<sup>3</sup>

### (4k Random)



- Pass-thru raw data:
  - 4k Rand Write: 80k IOPS
  - 4k Rand Mixed: 179k IOPS
  - 4k Rand Read: 634k IOPS
- 4-Disk RAID0 Read: 952k IOPS
- Physical CPU Cores Used:
  - 4-Disk RAID0 Read: 17 Cores
  - 4-Disk RAID5 Write: 6.3 Cores

**48 total physical cores on this 2 socket,  
Intel® Xeon® 8160T based system**

See appendix for footnotes.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit [www.intel.com/benchmarks](http://www.intel.com/benchmarks).


Performance results are based on testing as of August 26, 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

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

# Performance – RAID vs Pass-thru

Windows\* 2016 with Intel® SSD DC P4510<sup>4</sup>

(128 Seq., 1 Worker)



- Pass-thru raw data:
  - 128k Seq. Write: 2.1GB/s
  - 128k Seq. Read: 2.9 GB/s
- 4-Disk RAID 0 Read: 10.3 GB/s
- Physical CPU Cores Used:
  - 4-Disk RAID0 Read: 1.9 Cores
  - 4-Disk RAID5 Write: 1.6 Cores

**48 total physical cores on this 2 socket,  
Intel® Xeon® 8160T based system**

See appendix for footnotes.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit [www.intel.com/benchmarks](http://www.intel.com/benchmarks).

Performance results are based on testing as of August 26, 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

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



# Other Intel® VROC Performance Resources

## Intel® VROC Scaling Assessment:

- [https://www.intel.com/content/www/us/en/support/articles/000032616/memory-and-storage/ssd-software.html?productId=122484&localeCode=us\\_e](https://www.intel.com/content/www/us/en/support/articles/000032616/memory-and-storage/ssd-software.html?productId=122484&localeCode=us_e)

## Intel® VROC Performance Comparison to RAID HBA:

- [https://www.intel.com/content/www/us/en/support/articles/000033123/memory-and-storage/ssd-software.html?productId=122484&localeCode=us\\_en](https://www.intel.com/content/www/us/en/support/articles/000033123/memory-and-storage/ssd-software.html?productId=122484&localeCode=us_en)

# **INTEL<sup>®</sup> VIRTUAL RAID ON CPU PUBLIC RESOURCES**

# Intel® VROC Resources and Access


Intel.com/VROC for:



- Product Brief
- FAQ
- Supported OS/HW Information
- Related Links

Intel® VROC Support Page for:

- User Guides
- Tech Briefs
- Performance Documentation



### Features at a Glance

Supported Platform	Platforms with Intel® Xeon® processor Scalable family processors, Intel® Xeon® W processors, or Intel® Xeon® D processors
Supported Configurations	For current list of supported NVMe SSDs, Operating Systems, and configurations, please refer to <a href="#">the Intel® VROC Support page</a>
SKUs Available	Intel® VROC Standard: RAID 0/1/10; 3rd Party SSD Support Intel® VROC Premium: RAID 0/1/5/10; 3rd Party SSD Support Intel® VROC Intel® SSD Only: RAID 0/1/5/10; No 3rd Party SSD Support
Key Features	<ul style="list-style-type: none"><li>• Bootable RAID</li><li>• RAID controller spanning for data volumes</li><li>• Management Tools (UEFI CLI, UEFI HII, OS CLI, GUI, Remote Webpage)</li><li>• Surprise Hot-plug</li><li>• Status LED Indication</li><li>• Hot Spare and Auto-rebuild</li><li>• Email Notifications for RAID events</li><li>• RAID 5 Power Loss Protection for Degraded Volume (Double Fault Protection)</li><li>• Bad Block Management</li><li>• Various Strip Sizes (4k, 8k, 16k, 32k, 64k, 128k)</li></ul>

### Related Materials

- > [Download the Solution Brief](#)
- > [Intel® Virtual RAID on CPU \(Intel® VROC\) FAQs](#)
- > [Download the Technology Brief](#)
- > [Intel® Volume Management Device \(Intel® VMD\)](#)
- > [Support Documentation and User Information](#)

# Intel® VROC RAID Performance Analysis

**WHITE PAPER**  
Data Center  
Memory and Storage

## RAID Performance Analysis on Intel® VROC

Exploring the mathematics behind RAID performance benefits and how they relate to Intel® VROC.

**Table of Contents**

Purpose .....	1
Scope .....	1
Target Audience .....	1
Terms and Symbols .....	1
RAID Mathematic Performance Model .....	2
OS Performance Impacts on Intel VROC .....	6

**Purpose**  
This white paper provides an analysis of RAID performance and describes the mathematics behind RAID performance, as well as the specific RAID implementations of Intel® Virtual RAID on CPU (Intel® VROC) and how operating systems, storage devices, and other components impact performance output.

**Scope**  
RAID storage solutions aim to provide users with an appropriate combination of data protection and performance acceleration. Each RAID level processes storage I/O in a different manner and stores data in a specific pattern across a set of RAID member disks. Our goal is to highlight those storage patterns for RAID levels 0/1/10/5 and explain how each pattern affects the performance of the storage solution. By understanding the underlying I/O process, one can understand the theoretical performance maximum for each RAID level. We will explore how Intel VROC implements these RAID levels in an actual product and how Intel® VROC's performance compares to the theoretical maximums laid out. In the future, as this document continues to evolve, we will explain why any performance deltas exist between Intel VROC and the theoretical maximums caused by the Intel VROC RAID engine, OS limitations, SSD architecture, and more.

**Target Audience**  
Data Center Administrators, Architects, and Managers; Storage Performance Testers; Server Platform Manufacturers

**Terms and Symbols**  
Throughout the rest of this section, the following terms and symbols will be used:

- N: Total number of drives in a RAID volume.

Public guide to understand RAID principles and set expectations for Intel® VROC performance results:

- Calculations for RAID level performance maximums
- Intel® VROC Performance considerations like OS impacts
- Continually updated with new findings (this is version 1.0)

# Intel® VROC Performance Evaluation Guide

Step by step instructions to run an effective performance test in Linux:

- Allow end-users to get results in line with our published performance documentation
- Code base for direct CLI copy paste for easy replication (.txt file)
- Intel approved BKMs for repeatable testing



Intel® Solid State Drive DC P4510 Series using Linux\*  
Intel® Virtual RAID on CPU (Intel® VROC) with FIO\*

*Performance Evaluation Guide*

*January 2019*

**Disclaimer:** This document is not intended for the general public and should be considered as guidance on testing Intel® VROC under Linux using FIO.

# What Intel® VROC Means to the Storage Industry

More **Value** derived from NVMe\* SSDs and Intel® Optane™ storage media!

Make NVMe\*  
RAID Economical

Better cost structure to promote NVMe through the supply chain

Improve System  
Performance

Unleash NVMe bandwidth for full market benefit

Leverage Intel®  
Innovations

Enable Intel® VMD for robust NVMe SSD connections

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

# Appendix

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit [www.intel.com/benchmarks](http://www.intel.com/benchmarks).

**1. System configuration:** Intel Tested: Intel® Server Board S2600WFT family, Intel® Xeon® 8170 Series Processors, 26cores@ 2.1GHz, RAM 192GB , BIOS Release 7/09/2018, BIOS Version: SE5C620.86B.00.01.0014.070920180847

**OS:** RedHat\* Linux 7.4, kernel- 3.10.0-693.33.1.el7.x86\_64, mdadm - v4.0 - 2018-01-26 Intel build: RSTe\_5.4\_WW4.5, Intel® VROC Pre-OS version 5.3.0.1039, 4x Intel® SSD DC P4510 Series 2TB drive firmware: VDV10131, Retimer

**BIOS setting:** Hyper-threading enabled, Package C-State set to C6(non retention state) and Processor C6 set to enabled, P-States set to default and SpeedStep and Turbo are enabled

**Workload Generator:** FIO 3.3, RANDOM: Workers-24, IOdepth- 256, No Filesystem, CPU Affinitized

**Pass Thru Baseline:** 1x Intel® SSD DC P4510 Series, 2 TB, Firmware: VDV10120, SSDPE2KX020T8

Performance results are based on testing as of October 5, 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

**2. System configuration:** Intel Tested: Intel® Server Board S2600WFT family, Intel® Xeon® 8170 Series Processors, 26cores@ 2.1GHz, RAM 192GB , BIOS Release 7/09/2018, BIOS Version: SE5C620.86B.00.01.0014.070920180847

**OS:** RedHat\* Linux 7.4, kernel- 3.10.0-693.33.1.el7.x86\_64, mdadm - v4.0 - 2018-01-26 Intel build: RSTe\_5.4\_WW4.5, Intel® VROC Pre-OS version 5.3.0.1039, 4x Intel® SSD DC P4510 Series 2TB drive firmware: VDV10131, Retimer

**BIOS setting:** Hyper-threading enabled, Package C-State set to C6(non retention state) and Processor C6 set to enabled, P-States set to default and SpeedStep and Turbo are enabled

**Workload Generator:** FIO 3.3, SEQUENTIAL: Workers-1, IOdepth- 128, No Filesystem, CPU Affinitized

**Pass Thru Baseline:** 1x Intel® SSD DC P4510 Series, 2 TB, Firmware: VDV10120, SSDPE2KX020T8

Performance results are based on testing as of October 5, 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

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

# Appendix

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit [www.intel.com/benchmarks](http://www.intel.com/benchmarks).

**3. System configuration:** Intel Tested: Intel® Server Board S2600WFT family, Intel® Xeon® 8160T Series Processors, 24cores@ 2.1GHz, RAM 192GB , BIOS Release 07/09/2018, BIOS Version: SE5C620.86B.00.01.0014.070920180847

**OS:** Windows Server 2016, Version 10.0.14393 Build 14393, Intel RSTe UI version: 5.4.0.1464, Intel® VROC Pre-OS version 5.4.0.1039, 4x Intel® SSD DC P4510 Series 2TB drive firmware: VDV10131, Retimer

**BIOS setting:** Hyper-threading enabled, Package C-State set to C6(non retention state) and Processor C6 set to enabled, P-States set to default and SpeedStep and Turbo are enabled

**Workload Generator:** IOMeter version: 1.1.0 Dyn buckets 1.2, RANDOM: Workers-16, IOdepth- 32, No Filesystem, CPU Affinitized

**Pass Thru Baseline:** 1x Intel® SSD DC P4510 Series, 2TB, Firmware: VDV10131, SSDPE2KX020T8))

Performance results are based on testing as of August 26, 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

**4. System configuration:** Intel Tested: Intel® Server Board S2600WFT family, Intel® Xeon® 8160T Series Processors, 24cores@ 2.1GHz, RAM 192GB , BIOS Release 07/09/2018, BIOS Version: SE5C620.86B.00.01.0014.070920180847

**OS:** Windows\* Server 2016, Version 10.0.14393 Build 14393, Intel RSTe UI version: 5.4.0.1464, Intel® VROC Pre-OS version 5.4.0.1039, 4x Intel® SSD DC P4510 Series 2TB drive firmware: VDV10131, Retimer

**BIOS setting:** Hyper-threading enabled, Package C-State set to C6(non retention state) and Processor C6 set to enabled, P-States set to default and SpeedStep and Turbo are enabled

**Workload Generator:** IOMeter version: 1.1.0 Dyn buckets 1.2, SEQUENTIAL: Workers-1, IOdepth- 128, No Filesystem, CPU Affinitized

**Pass Thru Baseline:** 1x Intel® SSD DC P4510 Series, 2TB, Firmware: VDV10131, SSDPE2KX020T8))

Performance results are based on testing as of August 26, 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

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



