

DUG'20

# Platform Performance Evolution - Learning from Reference Storage Platform

Sarika Mehta, Storage Performance Engineer – Intel Corporation

The Intel logo is located in the bottom left corner of the slide. It consists of the word "intel" in a white, lowercase, sans-serif font, with a registered trademark symbol (®) to its right. The logo is positioned over a dark blue background. Above the logo, there are several overlapping squares in shades of blue, creating a decorative graphic element.

intel®

# Legal Disclaimers

All product plans and roadmaps are subject to change without notice.

Intel optimizations, for Intel compilers or other products, may not optimize to the same degree for non-Intel products.

Intel technologies may require enabled hardware, software, or service activation.

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.

Performance varies by use, configuration, and other factors. Learn more at [www.intel.com/PerformanceIndex](http://www.intel.com/PerformanceIndex).

Performance results are based on testing as of dates shown in the configurations and may not reflect all publicly available security updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Results have been estimated or simulated.

© 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.

A graphic consisting of three squares of varying sizes and shades of blue, arranged in a descending staircase pattern from top-left to bottom-right.

# Agenda

- Reference Storage Platform
- Performance evolution over time
  - Bandwidth and IOPS
  - Read QoS(quality of service)
  - Read QoS in presence of writes


# Reference Storage Platform – A Learning Vehicle

## ■ Goal:


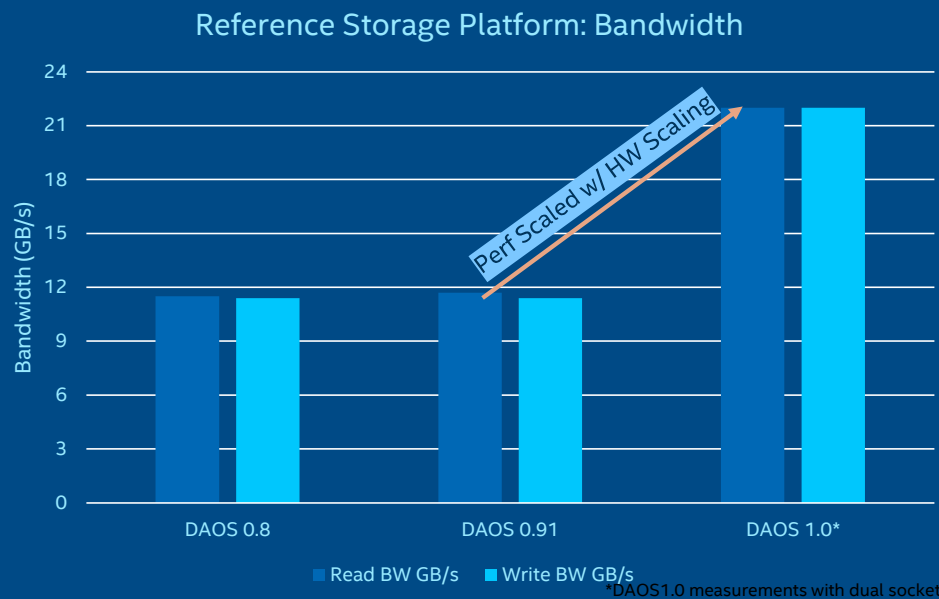
- Identify storage performance bottlenecks and optimizations
- Understand workloads with platform changes
- Verify hypotheses/assumptions of improved ingredients at the platform level (cost/performance)

## ■ Configuration:

- Distributed Asynchronous Object Storage (DAOS)
- 100GbE RoCE v2
- Intel® Xeon® Gold 6240 CPU @ 2.60GHz
- Intel® Optane™ Persistent Memory 768GB (6x128GB)
- Intel® SSD D5-P4326/ Intel's next generation QLC
- Test application - FIO



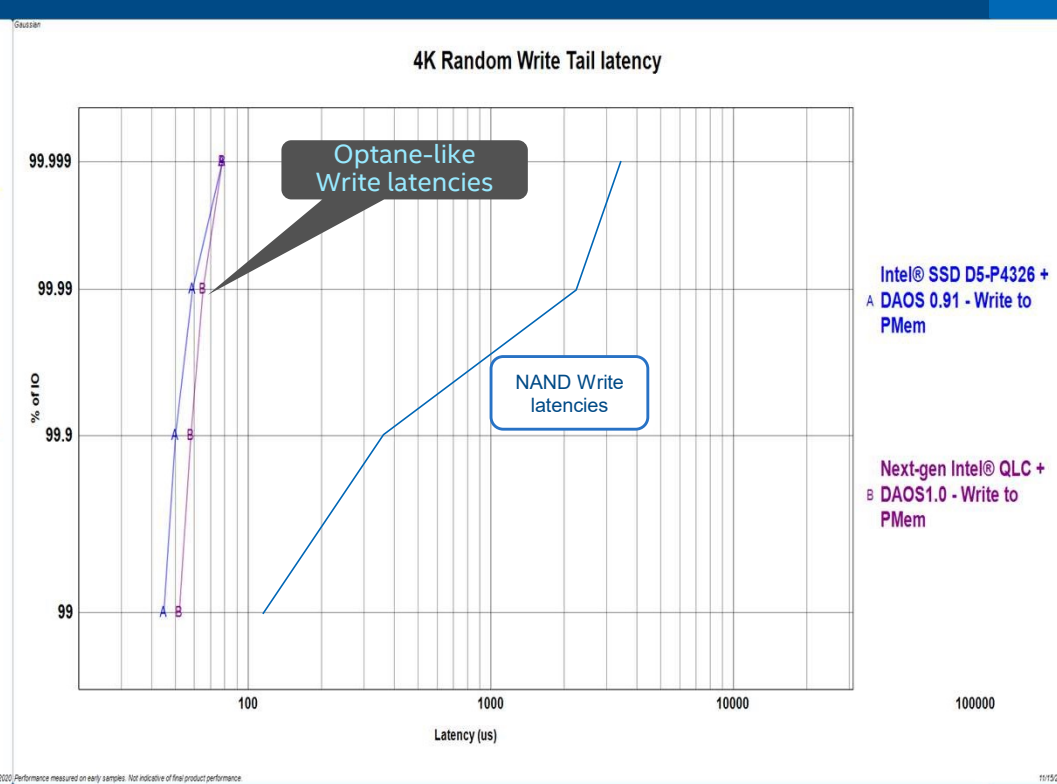
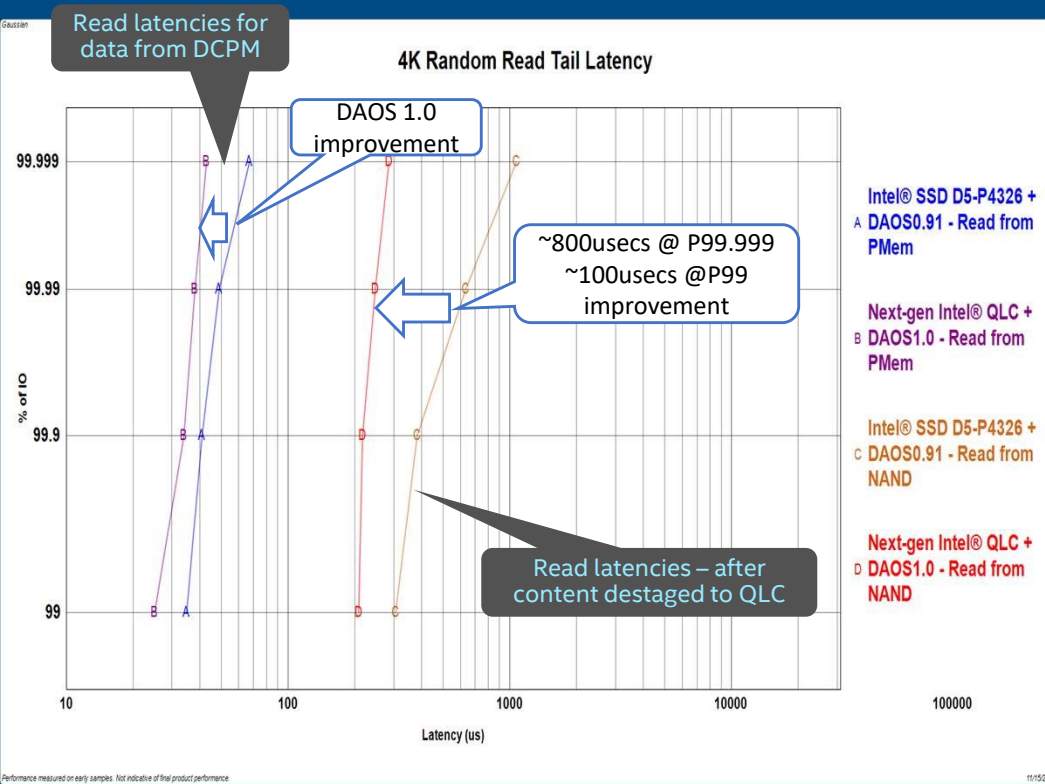
# Performance Improvements Over Time



100GbE Network link saturated for bandwidth and 4K Read IOPS.  
Write stack improvements in progress to improve 4K Write IOPS.

See backup for workloads and configurations. Results may vary.

# Intel® Optane™ technology-like write latencies, NAND like read latencies




Performance measured on early samples. Not indicative of final product performance. 10/15/2020

Significant improvement in Read tail latencies with DAOS 1.0 and Next-gen Intel® QLC SSD

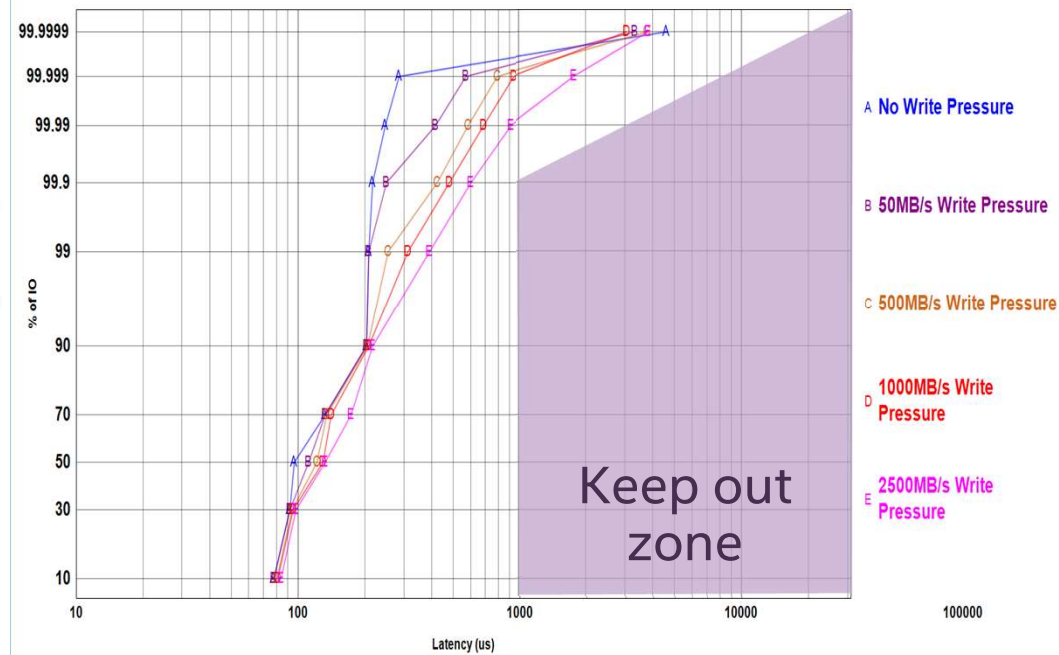
See backup for workloads and configurations. Results may vary.

# Responsivity is key for apps to consume IOPS/TPS

4K Random Read Tail Latencies in presence of Writes (Intel®SSD D5-P4326, QD=1, DAOS 1.0)



4K Random Read Tail Latencies in presence of Writes (Next-gen Intel®QLC, QD=1, DAOS 1.0 + Patches)



Improved transactions per second from 500MB/s to 2500+ MB/s

See backup for workloads and configurations. Results may vary.

intel®



# Test configuration and workloads

## ■ Test configurations:

- Distributed Asynchronous Object Storage (DAOS)
- 100GbE RoCE v2
- Intel® Xeon® Gold 6240 CPU @ 2.60GHz
- Intel® Optane™ Persistent Memory 768GB (6x128GB)
- Intel® SSD D5-P4326/ Intel's next generation QLC(Early samples)

## ■ Workloads

- All data was collected using FIO.
- For Bandwidth: IO Transfer size = 256K Queue Depth=16,128
- For IOPS: IO Transfer size = 4K Queue Depth=16,128
- For Latency: IO Transfer size = 4K Queue Depth=1

- Results were collected by Intel® Corporation from August through October of 2020.