Workload Guide

Intel® QLC 3D NAND SSDs
Data Center Storage

QLC NAND SSDs Are Optimal for Modern Workloads

Accelerate access to vast datasets with read-performance optimized, high density PCIe QLC NAND

What are QLC NAND SSDs?

At four bits per cell, quad-level cell (QLC) NAND technology expands data capacity beyond single-level cell (SLC), multi-level cell (MLC), and triple-level cell (TLC) NAND solid state drives (SSDs). The greater density provided by QLC drives enables more capacity in the same space, for a lower cost per gigabyte.

Why are QLC NAND SSDs needed now?

More data is being created, stored, and analyzed than ever before. High-growth segments in computing today—such as artificial intelligence (AI) and machine learning (ML), big data/analytics, hyperconverged infrastructure (HCI), high-performance computing (HPC), content-delivery networks (CDNs) cloud digital video recorder (cDVR), and cloud storage—all require access to ever-expanding data volumes. And because these workloads need rapid, high bandwidth access to data, they require read-optimized performance with low latency.

While hard disk drives (HDDs) have historically been the standard for warm storage, they are struggling to keep up with the demands for these read-intensive workloads. Additionally, HDDs require a significant footprint in the data center, which adds to space, power, and cooling costs.

TLC NAND SSDs, on the other hand, are a good fit for mixed and write-heavy workloads, but they are typically not as cost-effective and capacity-optimized for large-scale, read-centric data needs. With 33 percent more bits per cell than TLC NAND SSDs, and efficient form factors providing up to 30.72TB of storage, Intel® QLC NAND SSDs enable both accelerated data access and high capacities (see Table 1).

Compared to legacy HDD arrays, Intel QLC 3D NAND SSDs can:

- Consolidate warm storage footprints up to 20x[1]
- Reduce the total cost of ownership (TCO) of a typical hybrid array up to 42 percent[2]
- Accelerate access to stored data by up to 25x[3]

Table 1. QLC NAND SSDs: The sweet spot for read-intensive workloads.

HDD	Intel® QLC NAND SSD	TLC NAND SSD
Performance: Very slow	**Performance:** Very fast reads/fast writes	**Performance:** Very fast reads/very fast writes
Capacity: High	**Capacity:** Very high	**Capacity:** High
Space efficiency: Poor	**Space efficiency:** Excellent	**Space efficiency:** Very good
Operating expenses (OpEx): High	**OpEx:** Low	**OpEx:** Low
Best for: Cold storage	**Best for:** Warm storage	**Best for:** Performance storage

Intel QLC 3D NAND SSDs: Unlocking the value of stored data

Like TLC NAND SSDs, QLC NAND SSDs can saturate the PCIe 4.0 bus on the read side and have near-TLC-like latency and quality of service (QoS). This makes the drive orders of magnitude more responsive than HDDs (see Table 2).

Table 2. How QLC SSDs compare to HDDs and TLC SSDs.

	HDD Western Digital Ultrastar DC HC650[4]	HDD Seagate Exos X18[5]	**QLC SSD Intel SSD D5-P5316[6]**	TLC SSD Intel SSD D7-P5510[7]
Sustained transfer rate (HDD)/sequential read (SSD)	250 MB/s	270 MB/s	**7,000 MB/s**	7,000 MB/s
Random read (input/output operations per second [IOPS])	Not specified	170 (queue depth 16)	**800K**	930K
Latency (random read/write)	4.16 ms	4.1 6ms	**120 µs/180 µs**	82 µs/15 µs

Modern workloads: characteristics and storage needs

QLC NAND SSDs are optimized for read-intensive workloads needing rapid access to vast datasets. Figure 1 shows examples of a range of workloads[8] that QLC NAND is well suited for based on I/O patterns.

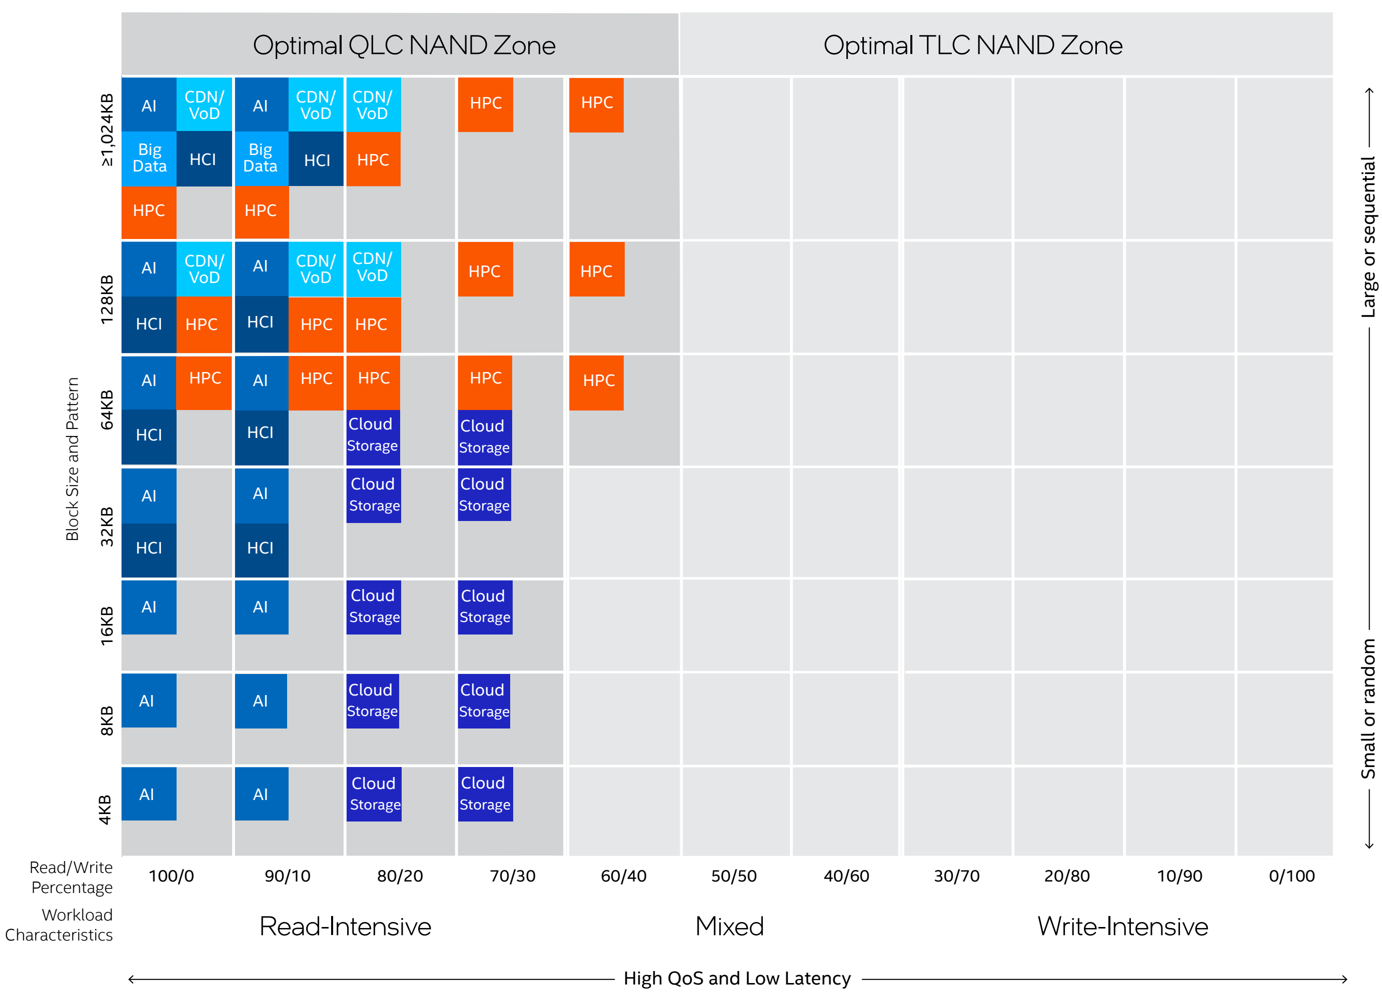

Figure 1. QLC NAND SSDs are optimized for read-intensive workloads.

The unique characteristics of these workloads are detailed in Table 3.

Table 3. Best-fit QLC SSD segments and workloads.

Segment	Workload Description	Storage Input/Output (I/O) Characteristics[9]
AI/ML data pipelines	Data pipelines store and move data efficiently through all AI/ML stages while also delivering performance and scalability to support other workloads.	• **Bandwidth:** High-bandwidth reads (training and inference phases) and writes (ingest phase) • **Latency:** Very low, with high QoS • **Endurance:** Low to moderate total bytes written (TBW) required because performance storage absorbs most writes • **Pattern:** Sequential reads (training and inference phases) and writes (ingest and preparation phases) • **Block size:** Highly variable
Big data/analytics	Collect and manage massive datasets to enable data mining and analysis for new insights.	• **Bandwidth:** Very high-bandwidth reads from capacity storage • **Latency:** Low, with high QoS • **Endurance:** Low TBW required because writes are to performance storage • **Pattern:** Sequential reads from capacity storage, multiple read requests present as random, and sequential writes to performance storage • **Block size:** Very large block reads and writes
CDN/video on demand (VoD)	Stream video content to end users. Mid-tier and edge servers are used to move content close to users to optimize the customer experience and reduce network traffic.	• **Bandwidth:** High read bandwidth • **Latency:** Low, with high QoS • **Endurance:** Low TBW required because writes are typically less than 5 percent[10] • **Pattern:** Highly read-intensive (up to 95 percent[10]) with sequential writes scheduled at low-use periods; reads most often present as random due to multi-user and caching distribution methods • **Block size:** Very large (most reads and writes are 128KB or more[10])
cDVR	Enable content to be saved in edge appliances in place of traditional DVRs. This creates a high-capacity, high-bandwidth storage workload with locality constraints on space, power, and cooling.	• **Bandwidth:** Sufficient bandwidth per capacity to saturate the interface • **Latency:** Low, with high QoS • **Endurance:** Low to moderate TBW because content is overwritten infrequently and performance storage is used to manage wear • **Pattern:** ~10/90, because less than 10 percent of recorded content is viewed;[11] reads are from the performance tier and mostly sequential, but they can present as random due to multi-user and caching distribution methods • **Block size:** Very large reads and writes
Cloud storage	Cloud service providers (CSPs) offer a range of storage services to enterprises based on capacity, throughput, and access time.	• **Bandwidth:** Highly variable (CSP offerings range from 250 MB/s to 4,000 MB/s per drive[12,13]) • **Latency:** Broad range (CSP services can be as low as sub-millisecond to single-digit-millisecond) • **Endurance:** Low TBW required because the capacity tier is paired with a performance tier • **Pattern:** Block storage is random or sequential; object storage is sequential, but multiple simultaneous small objects can present as random • **Block size:** Block storage default is 4KB; object storage is highly variable (ML data is 2–4KB; rich content is 64KB or higher)

HCI	Modern HCI storage subsystems are architected with performance and capacity storage tiers. Performance storage absorbs most writes with de-staging to capacity only when needed, which helps extend the life of the capacity tier.	• **Bandwidth:** High to efficiently saturate the network • **Latency:** Average (100–200 ms) • **Endurance:** Low TBW required because performance storage absorbs writes • **Pattern:** ~90 percent reads from capacity storage with more than 90 percent sequential;[14] might be presented as random when requests come from multiple hosts; occasional sequential writes from the performance tier during de-staging • **Block size:** Usually large block (Microsoft Azure Stack HCI is 256KB,[15] VMware vSAN I/Os are normalized to 32KB and 64KB[16])
HPC	Read directly from capacity storage at the start of, or during, long compute runs, when massive amounts of data are required. Also sequentially read data from capacity storage in cases of compute recovery. Frequent snapshots are taken during compute runs, with data temporarily held in memory and subsequently written as sequential blocks to capacity storage.	• **Bandwidth:** High • **Latency:** Very low (<1ms), with high QoS • **Endurance:** Low TBW required because writes are coalesced in memory or SSD cache drives • **Pattern:** Mix of sequential reads (during compute recovery) and writes (during memory de-staging) • **Block size:** Wide range (4–16KB to 64KB and higher)

Extract more value from your data with Intel QLC 3D NAND SSDs

Built on proven technology, Intel QLC 3D NAND SSDs provide exceptional high-bandwidth, low-latency read performance to support modern business-critical workloads, including ML, AI, CDNs, analytics, and big data while helping to lower TCO

Learn more

Intel QLC 3D NAND SSDs

Intel® SSD D5-P5316 product brief

QLC NAND Technology Is Ready for Mainstream Use in the Data Center

Replace Legacy Storage in CDNs with Efficient, Cost-Effective SSDs

compared to HDDs and TLC SSDs by consolidating the storage footprint and reducing operational costs.

[1] "Up to 20x reduction of warm storage footprint" claim is based on comparing 4TB HDDs, which require 10 (2U) of rack space to fill up 1PB or storage, against 30.72TB Intel® SSD D5-5316 E1.L or U.2 drives, which take 1U of rack space to fill up 1PB of storage.

[2] Using a CDN mid-tier server opportunity. Baseline server assumptions of 20K active users, minimum capacity of 480TB, target throughput of 190Gbps, and 85 percent cache hit ratio. Hybrid array per server configuration of 2 x Intel® Xeon® Scalable 6330 processor, 512GB memory, 12 x 8TB Seagate Exos HDD with 2.088Gbps throughput, 2 x 7.68TB Intel® SSD D7-P5510 with 56Gbps throughput. Hybrid array cost (CPU + memory + RAID controller + NVMe expander + chassis + power supply) net of storage $10,530. All-QLC per server configuration of 2 x Intel Xeon Scalable 6338N processor, 512GB memory, 20 x 30.72TB Intel® SSD D5-P5316 with 54.4Gbps throughput. All-QLC server cost (CPU + memory + RAID controller + NVMe expander + chassis + power supply) net of storage $11,530. HDD pricing as of March 15, 2021. SSD pricing is Intel budgetary pricing. Actual price can vary and may not reflect the pricing used in the TCO model. Operating cost assumptions of 1,100watts/server $0.12/kWh power and cooling and, $75.76 annual rack cost per RU.

[3] Sequential read performance based on Intel® SSD D5-P5316 (intel.com/content/www/us/en/products/memory-storage/solid-state-drives/data-center-ssds/d5-series/d5-p5316-series.html) compared to Seagate Exos X18 (seagate.com/files/www-content/datasheets/pdfs/exos-x18-channel-DS2045-1-2007GB-en_SG.pdf).

[4] Western Digital Data Sheet: Ultrastar DC HC650. https://documents.westerndigital.com/content/dam/doc-library/en_us/assets/public/western-digital/product/data-center-drives/ultrastar-dc-hc600-series/data-sheet-ultrastar-dc-hc650.pdf.

[5] Seagate. "X18 SATA Product Manual." July 2020. seagate.com/www-content/product-content/enterprise-hdd-fam/exos-x18/_shared/en-us/docs/100865854a.pdf.

[6] Intel. "Intel® SSD D5-5316 provides industry-leading endurance at 0.41 DWPD." intel.com/content/www/us/en/products/memory-storage/solid-state-drives/data-center-ssds/d5-series/d5-p5316-series.html.

[7] Intel. Intel® SSD D7-P5510 Series. https://ark.intel.com/content/www/us/en/ark/products/205382/intel-ssd-d7-p5510-series-7-68tb-2-5in-pcie-4-0-x4-3d4-tlc.html.

[8] HPC workload characterization assumes I/O optimization through Intel® Optane™ persistent memory using DAOS (Distributed Asynchronous Object Storage).

[9] Storage workload characterizations are, unless noted otherwise, Intel estimates based on subject-matter expert deep knowledge of storage I/O patterns and behavior for a given segment.

[10] Percentage is an estimate based on Intel internal analysis across top tier CDNs.

[11] Percentage is an estimate based on Intel internal analysis across top tier cDVRs.

[12] Amazon AWS. "Amazon EBS Volume Types." https://docs.amazonaws.cn/en_us/AWSEC2/latest/UserGuide/ebs-volume-types.html.

[13] Tencent. "An Overview of Cloud Hard Drive Prices." https://cloud.tencent.com/document/product/362/2413.

[14] Percentage is an estimate based on Intel internal analysis across top tier HCI solutions.

[15] Microsoft. Azure/Virtual Machines/Disks. https://docs.microsoft.com/en-us/azure/virtual-machines/disks-types#premium-ssd.

[16] VMware. "Performance Metrics when using IOPS Limits with vSAN (vmware.com)" https://core.vmware.com/resource/vmware-vsan-design-guide#section1

Printed in USA 0321/RG/PRW/PDF Please Recycle 346489-001US