



SEAGATE

WHITE PAPER

HYPERSCALE STORAGE TCO

Why HDDs Dominate Cloud Architecture

CONTENTS

- 1 INTRODUCTION
- 2 STORAGE ARCHITECTURE IS EVOLVING
- 3 THE DATA STORAGE IMPERATIVE
- 4 TOTAL COST OF OWNERSHIP MATTERS
- 5 UNDERSTANDING TCO
- 6 ELECTRICITY AS A COST VARIABLE
- 7 HDDS AND THE TCO SWEET SPOT
- 8 CONCLUSION



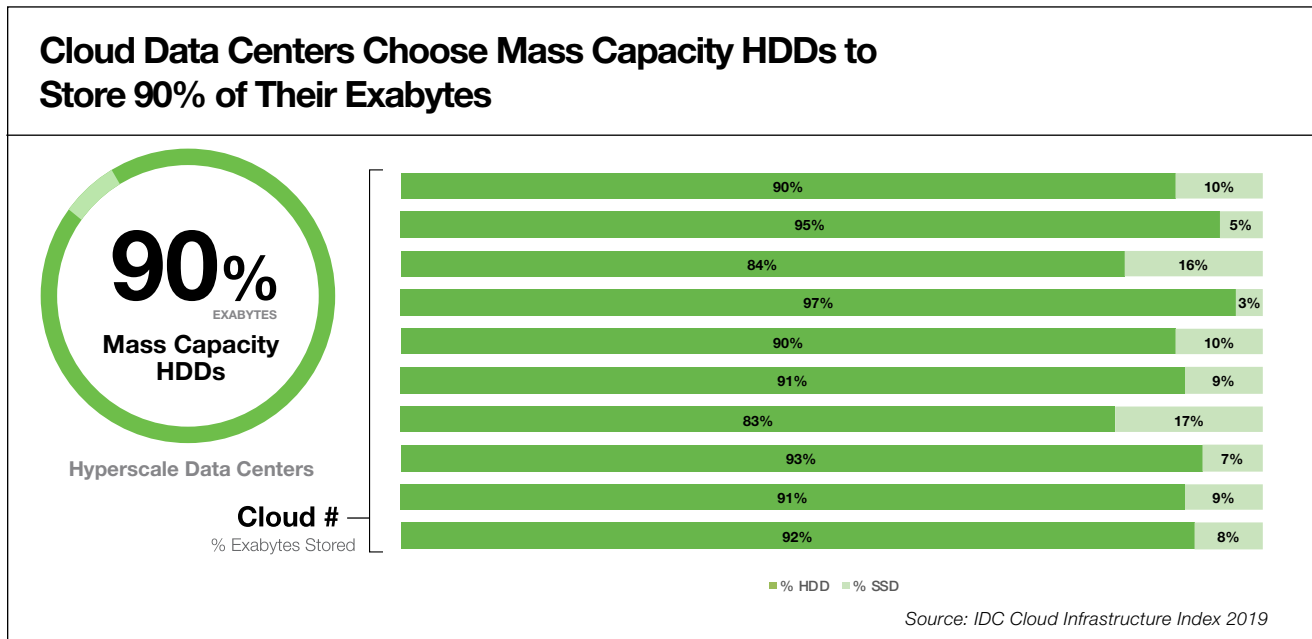
Introduction

Data is fueling digital transformation across industry sectors, and companies large and small are in a race to monetize precious information. Due to data’s growing value, and to the increasing quantities of structured and unstructured data, companies that want to remain competitive need to embrace a new digital alchemy that enables them to capture, store, and analyze massive amounts of data.

Mass storage is more important than ever, and today, IT architects who manage public, private, edge and hybrid cloud infrastructures need to deploy storage solutions that deliver optimal performance for the lowest possible cost. Reducing total cost of ownership (TCO) is a primary driver when it comes to a cost-benefit analysis of storage solutions and is the reason cloud architects continue to use HDDs as the primary workhorse for enterprise data centers.

The big public cloud service providers (CSPs) have dominated an era of mobile-cloud centralized architecture. As the leading supplier of exabytes to the world’s top public CSPs, Seagate has unique insights into hyperscale storage architectures and demand trajectories.

HDDs currently dominate the cloud exabyte market—offering the lowest cost per terabyte based on a combination of factors including price, cost, capacity, power, performance, reliability, and data retention. SSDs, with their performance and latency metrics, provide an appropriate value proposition for performance-sensitive, highly transactional workloads closer to compute nodes. HDDs represent the predominant storage for cloud data centers because they provide the best TCO for the vast majority of cloud workloads. More



than 90% of exabytes in cloud data centers are stored on HDDs, and the remaining 10% are stored on SSDs according to market intelligence firm IDC. Industry analyst TRENDFOCUS stated that more than 1ZB of HDD storage capacity shipped in just the last year.



These hyperscalers have enabled data from billions of endpoint devices to be stored in centralized IT infrastructure. But now, IT 4.0 is creating a new data-driven economy. The new era includes manufacturing automation and a proliferation of IoT-connected devices that can communicate, analyze, and use data for everything from controlling production to deriving actionable business insight.

As part of this new paradigm, companies are developing their own private cloud and edge storage solutions in addition to using public clouds and can look to the experiences and example of hyperscalers to understand how best to optimize storage architectures.

Mass capacity HDDs installed in private cloud and edge data centers enable businesses to implement the cloud applications needed for processing and analyzing massive amounts of unstructured data. According to the Seagate-sponsored IDC report *Data Age 2025: The Digitization of the World from Edge to Core*, there will be more than 175ZB of data per year coursing through the economy by 2025. To put the storage requirements in context, 50 million 20-terabyte hard drives are needed to store one zettabyte of data.

Storage Architecture Is Evolving

Of course, cloud storage doesn't mean that data is parked in thin air. On-demand data storage requires energy-efficient network infrastructure and hyperscale architecture that is capacity-optimized, composable, and software-defined.

A well-designed storage infrastructure delivers end-to-end management and real-time access required to help businesses extract maximum value from their data. Data center storage architectures must be optimized for massive capacity, fine-tuned for maximum resource utilization, and engineered for ultra-efficient data management.

As the IT 4.0 marketplace evolves, it will continue to generate massive quantities of unstructured data. This provides companies with significant opportunities—if they can capture, store, and analyze that data. IT architects and infrastructure planners have to facilitate even more data storage across multiple public, private, and edge clouds. According to Seagate's *Rethink Data* report—based on an IDC survey of 1500 global enterprise leaders—by 2025, 44% of data created in the core and at the edge will be driven by analytics, artificial intelligence, deep learning, and an increasing number of IoT devices capturing, creating, and feeding data to the edge.

The evolution of IT 4.0 will include more distributed storage architecture and an increase in data processing capacity at the edge. That means businesses will have to manage new data sources and highly complex unstructured data types. Per Seagate's *Rethink Data* report, total volume of enterprise data is expected to grow at the average annual growth rate of 42.2% between 2020 and 2022.

Monolithic, centralized storage architectures are being disrupted by a growing need for latency-sensitive applications running at the edge. New hybrid architectures—incorporating multiple public, private, and geo-distributed edge clouds—are emerging in order to capture and store data across multiple locations using a decentralized model.



The Data Storage Imperative

As the economic importance of data increases, demand for capturing every byte of valuable information will grow. But the disparity between data created and data stored is increasing, and that means large amounts of data—and its value—are lost.

According to the *Rethink Data* report, IDC projects that 175ZB of data a year will be produced by 2025, but less than 10% will be stored. That means 90% of all data produced ends up on the cutting room floor. In the IT 4.0 universe, that translates into lost opportunities.

Data is the currency of the digital economy. As companies realize the importance of data for competitiveness and profitability, demand for storage will skyrocket. By 2025, IDC predicts in the *Rethink Data* report that 12.6ZB of installed storage capacity will be managed by businesses, with cloud service providers managing 51% of this capacity.

Business cloud architects need to continue to expand public cloud capacity while investing in the infrastructure needed to support private and hybrid clouds. Developing that infrastructure is an important proactive step for cloud architects.

Business demand for software-defined storage and hybrid architecture is growing. Per the *Rethink Data* report, IDC predicts that companies will store 9ZB a year by 2025. Private and hybrid clouds offer improved interoperability, control, and security. In order for data centers to ingest, store, and service this data—and provide value from it—they need to deploy cost-effective storage technologies at scale.

TCO is crucial when it comes to storage deployment decisions. The gold standard for data storage in the cloud is HDD, which offers a combination of maximum capacity and optimal performance at the lowest overall cost.

Total Cost of Ownership Matters

HDD capacity continues to increase, with 18TB drives now widely available. Increased capacity is driving down the cost per terabyte. Moreover, new **HAMR** (heat assisted magnetic recording technology) drives were rolled out in 2020, with an initial capacity of 20TB and are forecasted to increase to 50TB by 2026. Higher capacity HDDs change the dynamics of data center storage. **MACH.2 HDDs**, for example, use multi-actuator technology to double hard drive IOPS performance, enabling cloud architects to maintain performance as capacity scales.

Rapid, continuous adoption of the highest capacity hard drives by hyperscalers has given them a sustained cost advantage over comparable business data centers. But now, data center architects can take advantage of high-capacity HDDs to reduce both CapEx and OpEx.



Understanding TCO

Seagate has done an extensive internal analysis of the TCO of storage architectures, based on data from a large set of cloud data centers.

Higher-capacity drives dramatically reduce the TCO for storage infrastructure. Total cost of ownership improves about 10% for every 2TB of storage added to an HDD. Factors contributing to this TCO reduction are \$/TB, Watt/TB, slot cost reduction, densification, and exabyte availability.

Looking ahead, HDDs are expected to maintain market share, thanks in part to advances that continue to increase capacities while maintaining performance. A balance between managing price while ensuring required performance is key for cloud architects that need to deliver on performance benchmarks as capacity scales.

While cloud applications vary by use case, IT and cloud architects have the goal of lowering cost by providing the densest storage possible. Even though system configurations may be different, a data center's overall IT CapEx includes JBOD enclosure racks, servers, network cards, rack switches, routers, and other hardware required to operate. From an OpEx perspective, electricity and staffing are the most significant costs. The total dollar amount required to deploy and utilize a drive is the slot cost. TCO represents that slot cost plus the cost of acquiring the drive.

TCO			
Included in this TCO calculation			
	CAPEX	OPEX	
Land Acquisition	Networking Cost	Rack/Power Dist. Cost	Spares
Building	Server Acquisition Cost	Server Power Cost	Resources
	JBOD Acquisition Cost	JBOD Power Cost	Taxes
	HDD Acquisition Cost	HDD Power Cost	Repairs
			Software and Maintenance

Understanding TCO requires an assessment of HDD acquisition and slot costs. Using publicly available data from large cloud service providers, it is possible to calculate TCO based on IT hardware costs. The calculation used here does not include land acquisition, brick and mortar, depreciation, resources, and maintenance.



Electricity as a Cost Variable

Power consumption and the cost of electricity is typically a key variable when assessing how OpEx impacts TCO. **Electricity costs** vary based on volume of consumption and location in the marketplace, but the cost typically ranges from 5 to 20 cents per kilowatt hour. **Power Utilization Efficiency (PUE)**¹ also varies, 1.2 to 2, influenced heavily by data center power conversion and cooling efficiency. This metric is important because it reflects the energy efficiency of a data center.

Energy efficiency is important for both power budgets and power cost. Hyperscale customers have been able to drive down the cost of energy significantly, which has substantially diminished the proportionate contribution of power to the overall TCO. As a result, the component cost of the IT hardware becomes the overwhelming contributor to TCO.

The development of renewable energy resources is helping companies reduce electricity costs, as well as respecting the environment. In addition to a lower cost for power, the availability of more energy-efficient HDDs promises to drive down power consumption over the next decade. Moreover, power balancing features provide cloud architects the ability to adjust drive settings to align to performance and power savings tradeoffs to their specific needs.

The most significant impact on TCO, however, comes from a reduction in the cost of terabytes per watts. As each new generation of hard drive increases the data density per disk, higher-capacity drives need the same amount of power as the lower-capacity drives of previous generations. That means cloud architects can deploy higher capacity HDDs without increasing energy costs.

HDDs and the TCO Sweet Spot

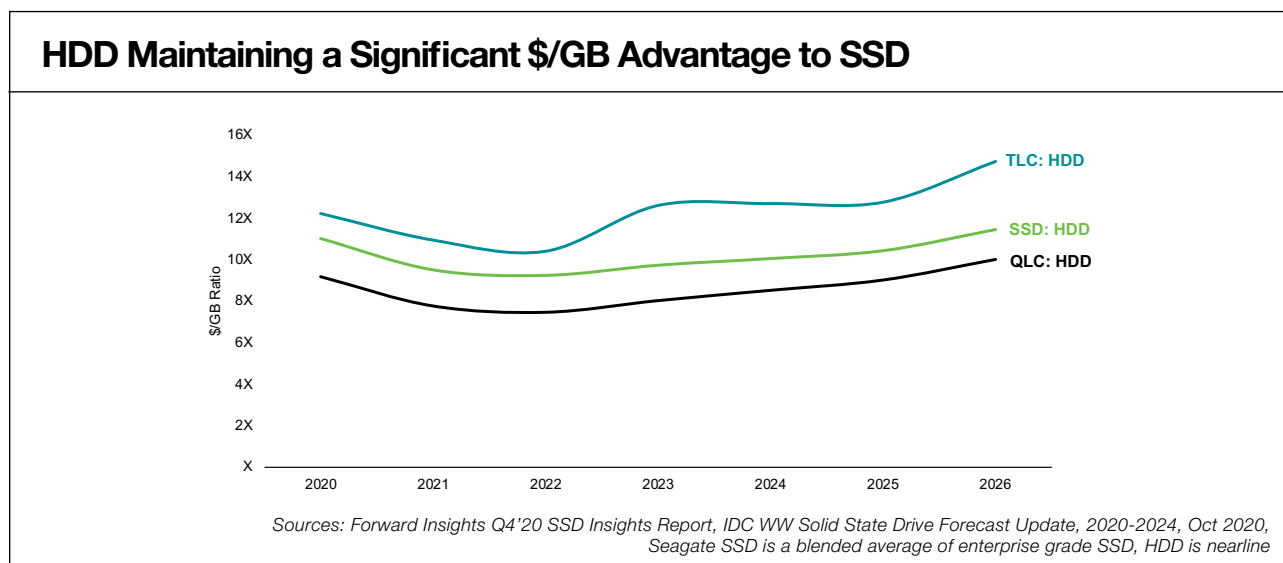
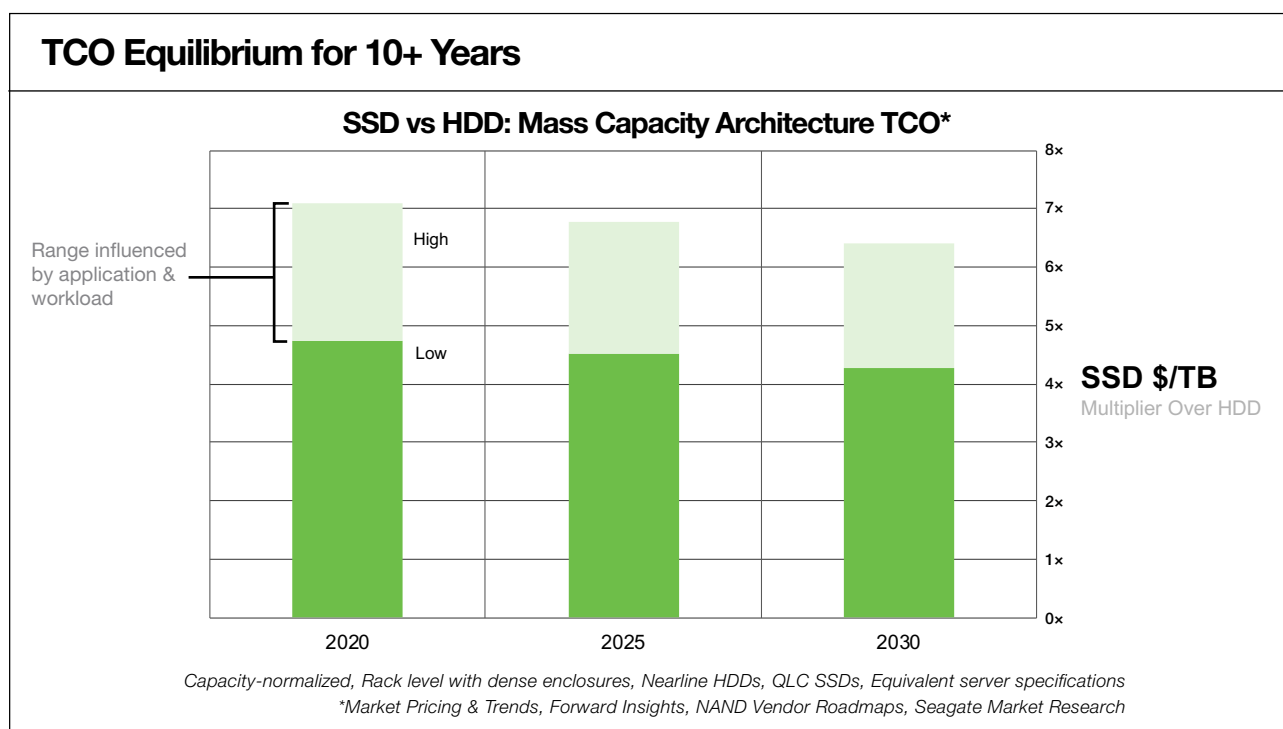
Total cost of ownership measured in \$/TB for a storage system is a critical metric for private data centers. In terms of TCO, HDDs continue to offer a better combination of price and performance than SSDs for everything from cloud-based gaming and surveillance applications to personal and enterprise-level computing. Based on enterprise-level application needs and TCO requirements, HDDs are well-positioned to dominate data centers for the next 10 years. In particular, HDDs built for cloud and hyperscale applications provide the right combination of capacity and performance for public, private, and edge environments.

Some vendors have suggested QLC NAND as an alternative technology for storage in data centers. Although NAND price, density, supply, and other dense-storage technical benefits don't come close to those of HDDs, NAND has been considered in this TCO analysis. Device comparisons here are based on raw \$/TB, not on a so-called effective \$/TB figure which would consider the impact of compression and deduplication data reduction capabilities; HDD and SSD technologies both provide equivalent data reduction capabilities.



To achieve the comparable density, 1x4U JBOD of HDD space would be replaced with 4x1U SSD JBOFs. This allows a direct 4U to 4U comparison within a rack infrastructure. Although additional networking and compute costs are required to utilize the performance boost from an SSD, this TCO analysis is focused on SSDs vs HDDs only, therefore it does not include these additional overhead costs for SSDs and contemplates only the network/compute infrastructure spend as it applies to HDDs.

The TCO of SSDs today is about six times that of HDDs (ranging from just under 5x to over 7x, depending on application and workload variables), while at the device level SSDs continue to be about eight times the cost of HDDs. Over the next decade, increased HDD capacity is expected to offset any decline in SSD costs. As a result, by 2030, the TCO of HDD infrastructure will continue to be about one-sixth the cost of equivalent capacity SSD deployments.



Continued lower electricity costs also mean that the TCO component of power consumption has a smaller impact for data center users. Even though SSDs require about one-third the power of HDDs, the difference does not significantly impact the overall TCO calculation based on cost per terabyte.

In order to fully realize the TCO benefits of the new HDD technology, and to remain competitive, IT and cloud architects will also need to use software for optimizing mass capacity and data-intensive workloads.

For most data center storage needs, HDDs are the most cost-effective option. Dual actuator technology, which nearly doubles the performance of HDDs, provides data centers an attractive value proposition for getting the performance they need at the lowest cost per terabyte.

Conclusion

Understanding TCO is essential for IT and cloud architects. TCO is also important for any company serving large business customers with private, on-prem data centers or private cloud hosting services. Companies offering infrastructure-as-a-service, platform-as-a-service, or storage-as-a-service face the same challenges as cloud storage architects.

Data lakes are overflowing, and in order to harness and analyze all that data, businesses need a place to store it. The options are public clouds, private clouds, or hybrid clouds. A hybrid might, for example, include an on-premises data center to reduce latency for edge applications and a backend in the public cloud for less time-sensitive storage. While the public clouds have been a catalyst for growth, the data-driven IT 4.0 economy will require hybrid clouds that offer the benefits of both public and private clouds. These cloud solutions must also seamlessly integrate with on-premises and edge data centers.

As performance-intense computing demands increase—fueled by data-hungry deep learning applications and increasing numbers of IoT connected devices—the symbiotic relationship between HDDs and SSDs will grow. As more SSDs are deployed to support high compute demands, even more HDDs will be required to store the data used for and generated by performance-hungry applications.

Understanding storage costs, improving capacity, and balancing performance and price will continue to drive decision-making when it comes to HDD deployment. TCO is the metric that will help ensure that cloud architects and companies of all sizes can get the most value out of their data.



References

Endnotes

¹ Source: United States Data Center Energy Usage Report, Lawrence Berkeley National Laboratory, 2016

Resources

Seagate Private Cloud Solutions:

<https://www.seagate.com/solutions/cloud/private-cloud/>

Lyve Rack – Open, affordable, and durable object storage solution:

<https://www.seagate.com/products/storage/object-storage-solutions/lyve-drive-rack/>

CORTX – Open source mass capacity object storage:

<https://www.seagate.com/products/storage/object-storage-software/>

Ready to Learn More?

Visit us at [seagate.com](https://www.seagate.com)



© 2021 Seagate Technology LLC. All rights reserved. Seagate, Seagate Technology, and the Spiral logo are registered trademarks of Seagate Technology LLC in the United States and/or other countries. MACH.2 is either a trademark or registered trademark of Seagate Technology LLC or one of its affiliated companies in the United States and/or other countries. All other trademarks or registered trademarks are the property of their respective owners. When referring to drive capacity, one gigabyte, or GB, equals one billion bytes and one terabyte, or TB, equals one trillion bytes. Your computer's operating system may use a different standard of measurement and report a lower capacity. In addition, some of the listed capacity is used for formatting and other functions, and thus will not be available for data storage. Seagate reserves the right to change, without notice, product offerings or specifications. TP724.1-2102US, February 2021

