STORAGE DEVELOPER CONFERENCE

SDC 23

BY Developers FOR Developers

New Cloud Workloads Implications for Storage Media Futures

Swapna Yasarapu and Shashidhar Joshi

Framing the Discussion

Framing the Discussion for Today

- Cloud Storage Services have been growing and the trend is expected to continue
- Storage Media Technologies (SSDs and HDDs) have supported these services well, until now
- However, with the growing scale, what has worked in the past, may not be good enough for the future
- We would like to bring some critical aspects of both these storage technologies that need improvement to
  - **Continue growth within their incumbent applications**
  - **Opportunity to grow into new applications**
- Through our understanding of workloads and applications, we will map requirements for these storage technologies to be able to meet the demands of the future

Cloud Applications and Workloads

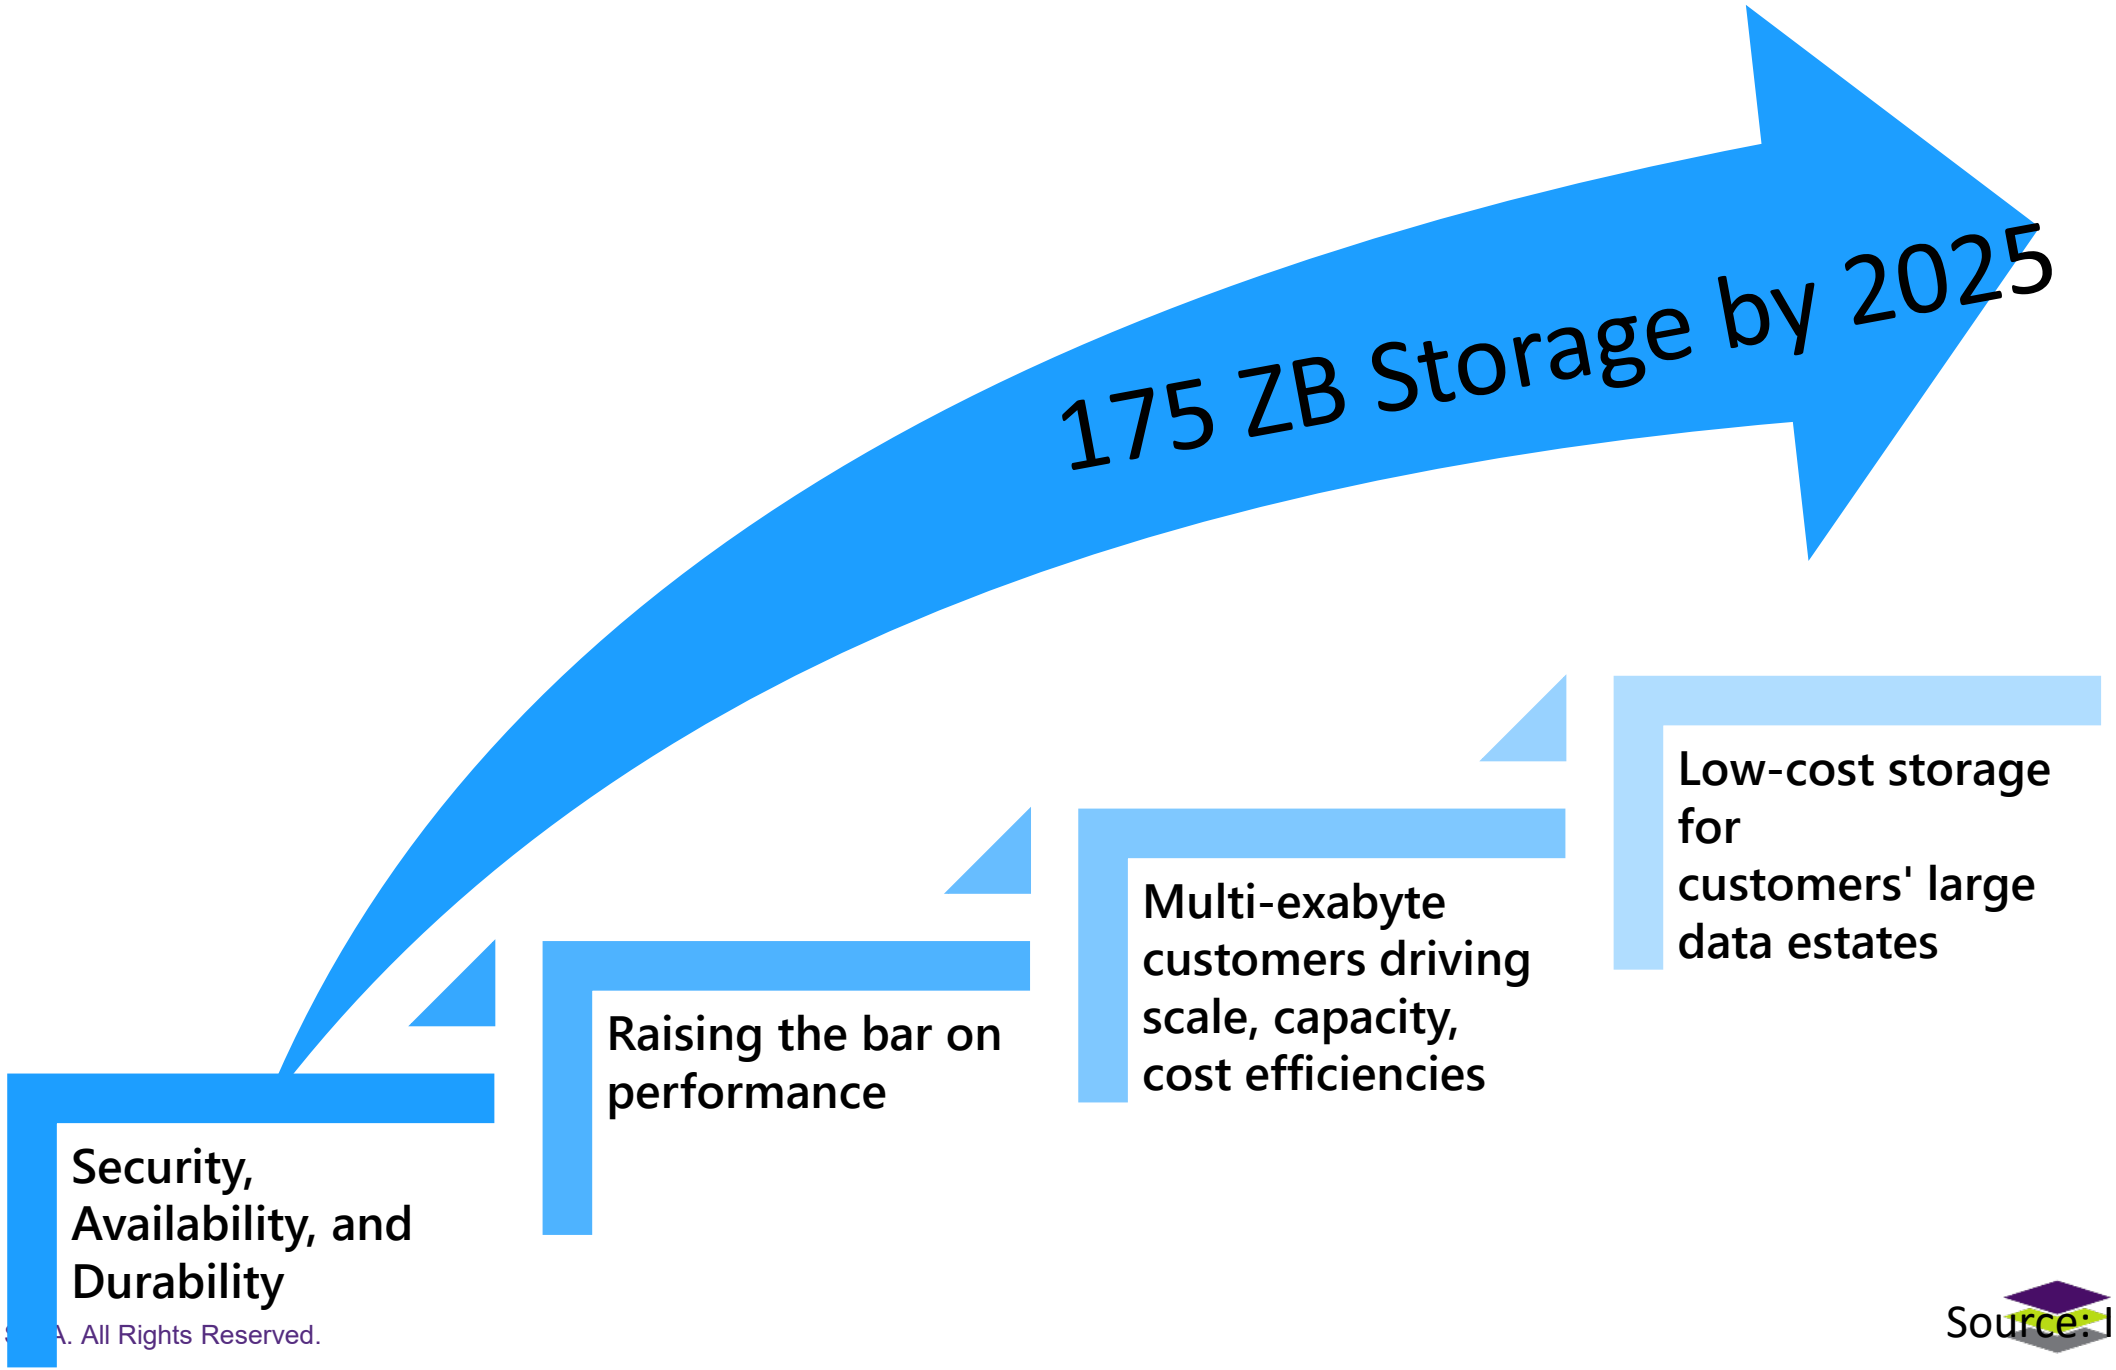

Source: IDC

Azure Storage portfolio

Durable, highly available, massively scalable

Block storage

Services

Azure Disk Storage

Azure Elastic SAN

Unique capabilities

Enterprise SAN Capabilities

Container Optimized Storage

Object storage

Services

Azure Blob

Data Lake Storage

Azure Managed Lustre

Unique capabilities

Premium Blob

Multi-protocol access (e.g. NFS, HDFS)

File storage

Services

Azure Files

Azure NetApp Files

Unique capabilities

Native NetApp File Storage

Azure File Sync

Capacity

100s of trillions of objects across many exabytes of data

Throughput

>300 Tbps average

(>100 exabytes per month)

IOPS

>400M tps

(>1 quadrillion per month)

Flash/SSD (Premium Storage)

HDDs (Standard Storage)

Applications To Workloads

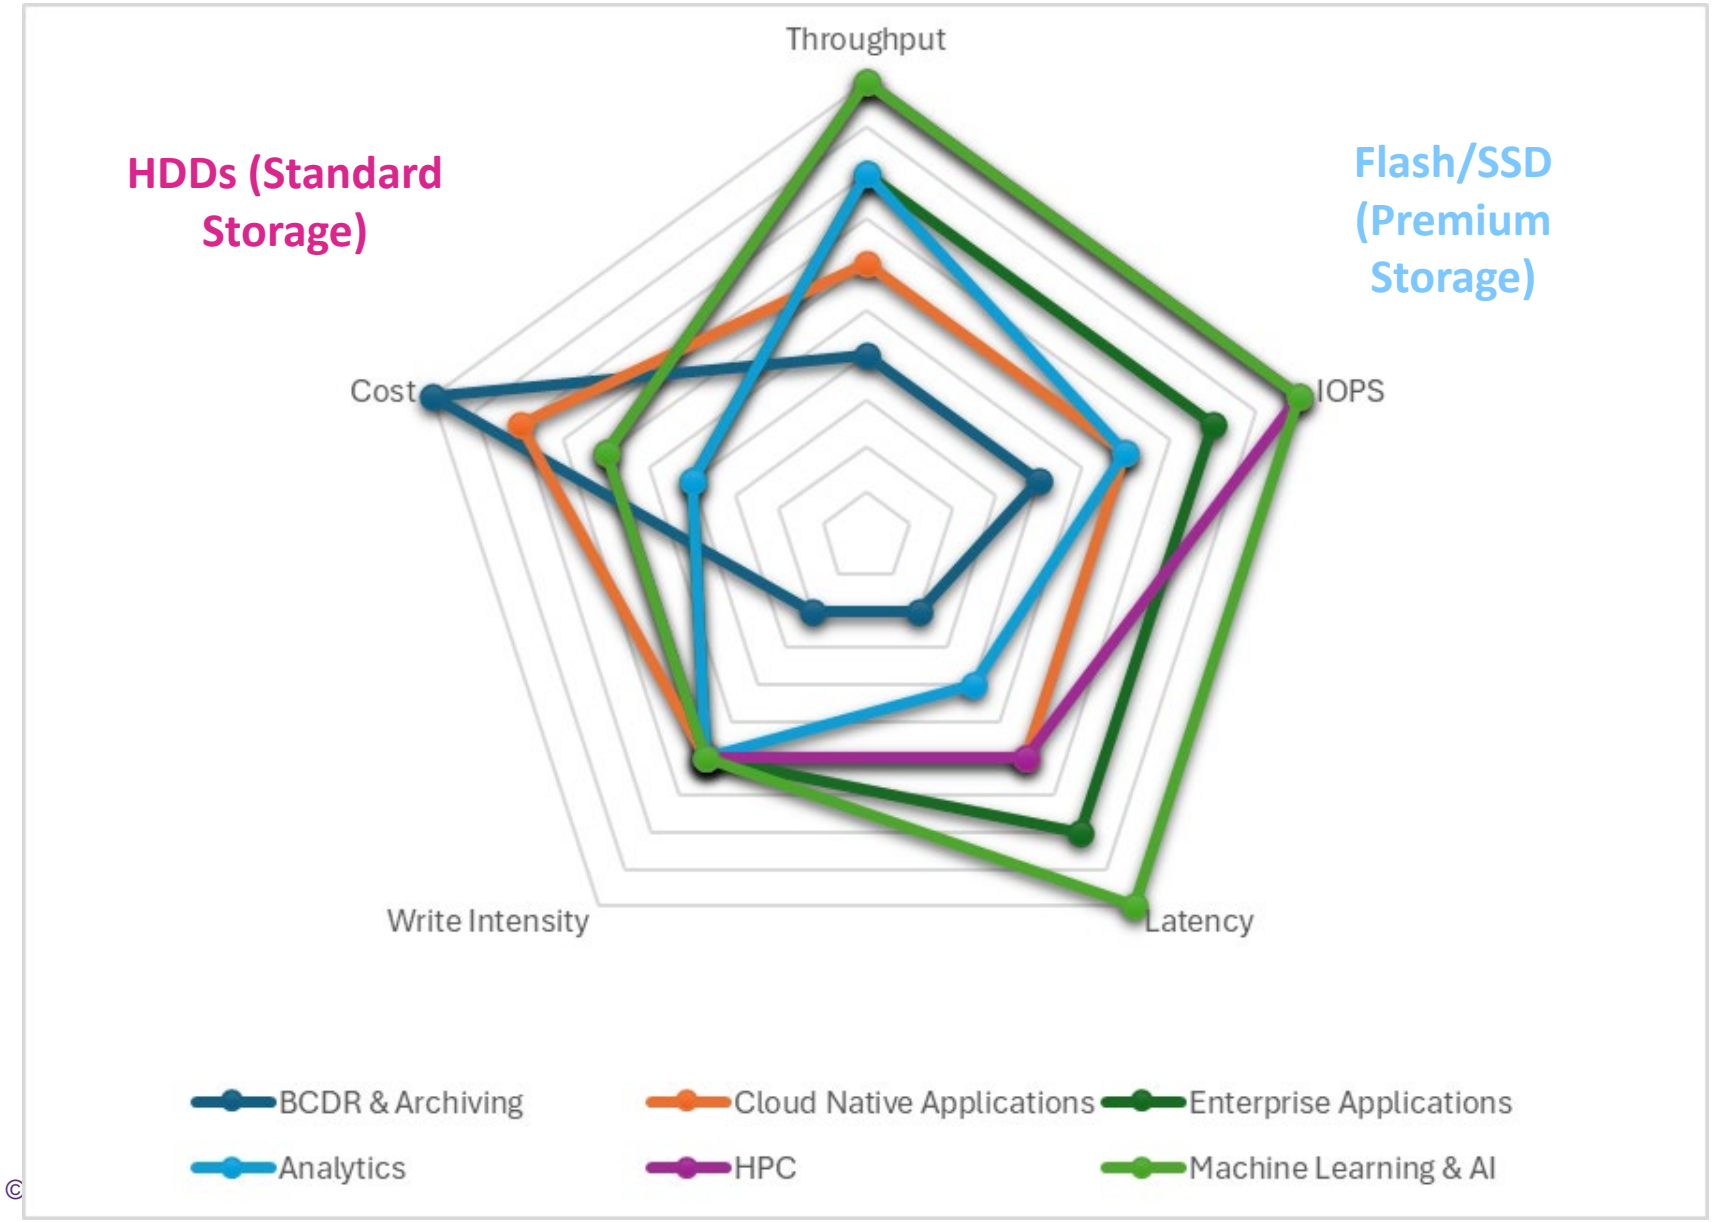

- BCDR and Archiving
- Cloud Native Applications
- Enterprise Apps
- Analytics
- HPC
- Machine Learning & AI

Read- Write IOPS Mix for SSDs & HDDs

SSD Read-Write Mix

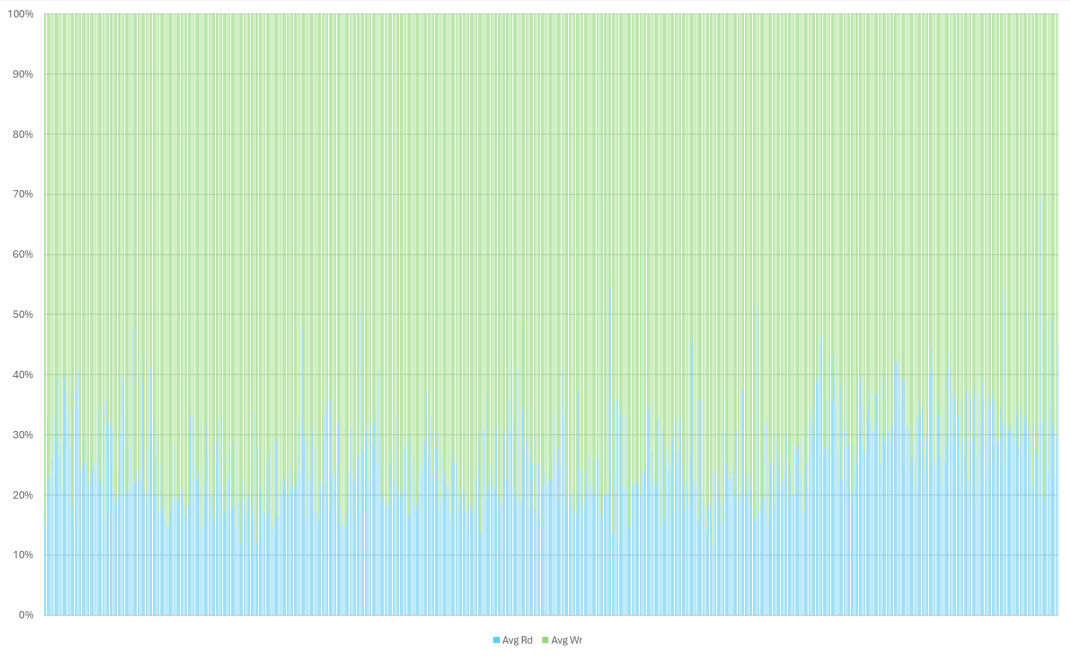

HDD Read-Write Mix

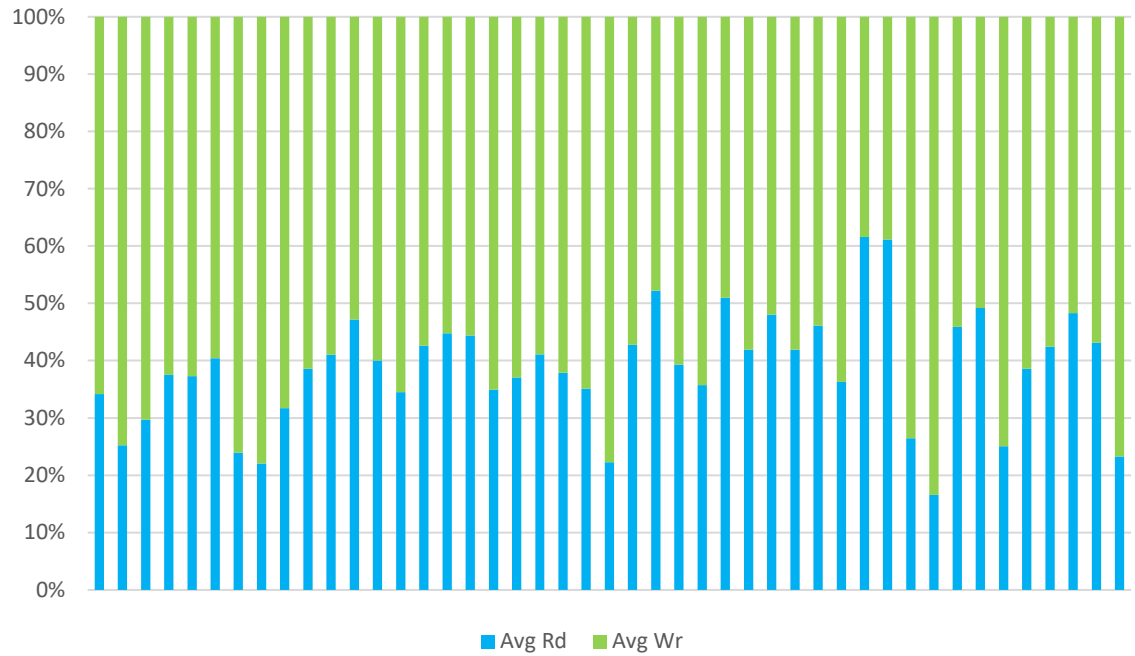

- Writes are 60% of the Workload, which includes Replication

TBW Continues to be a very important metric for both storage classes

Read- Write Throughput Mix for SSDs & HDDs

SSD Read-Write Mix

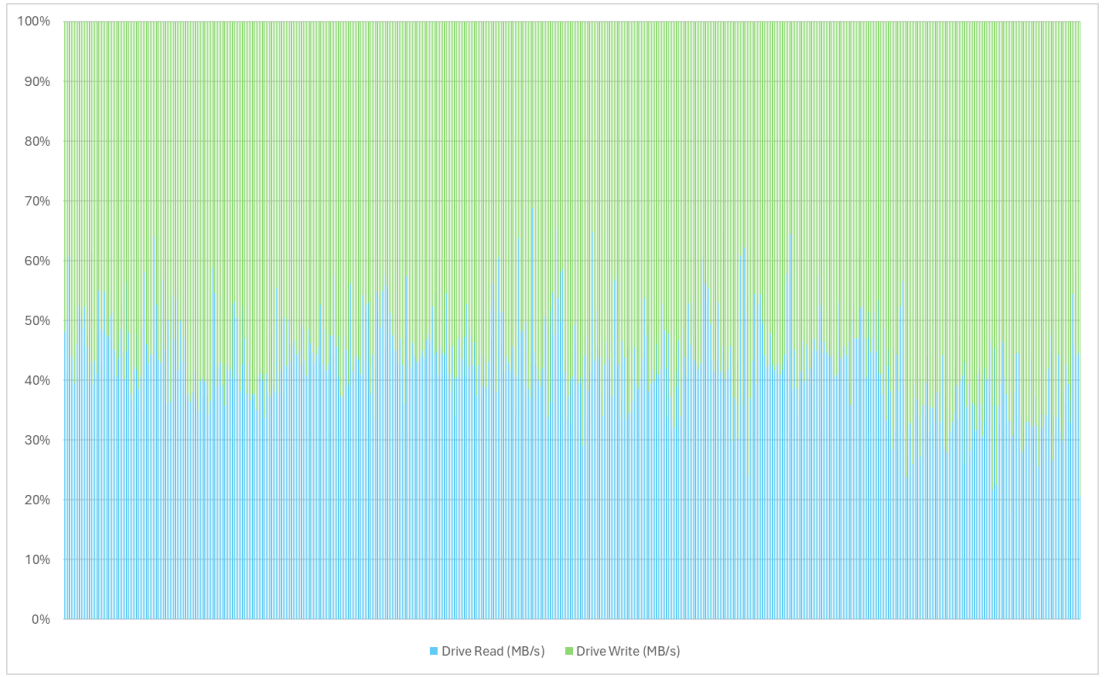

HDD Read-Write Mix

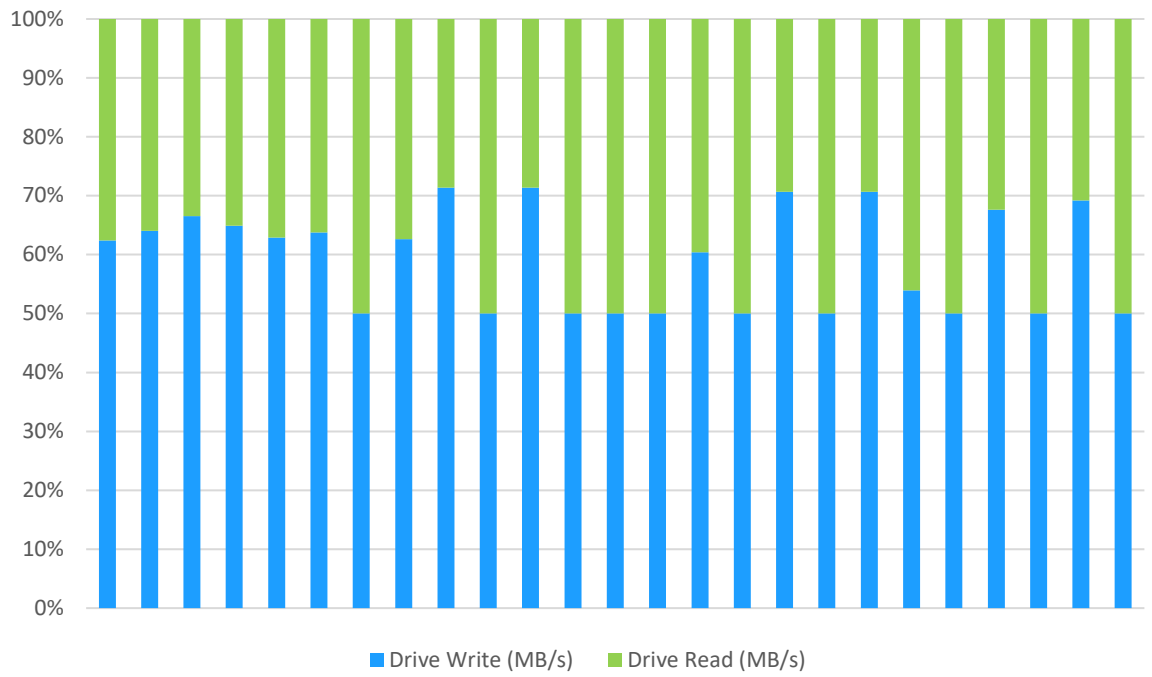

- Writes and Reads are equally important for workloads

Read-Write Mix consistent across both storage tiers; albeit at different overall performance capability

Average Read Latency

- SSD Latency is multiple orders better than HDD
- SSD Latency relatively consistent across capacities
- HDD Latency is degrading as a function of higher capacity

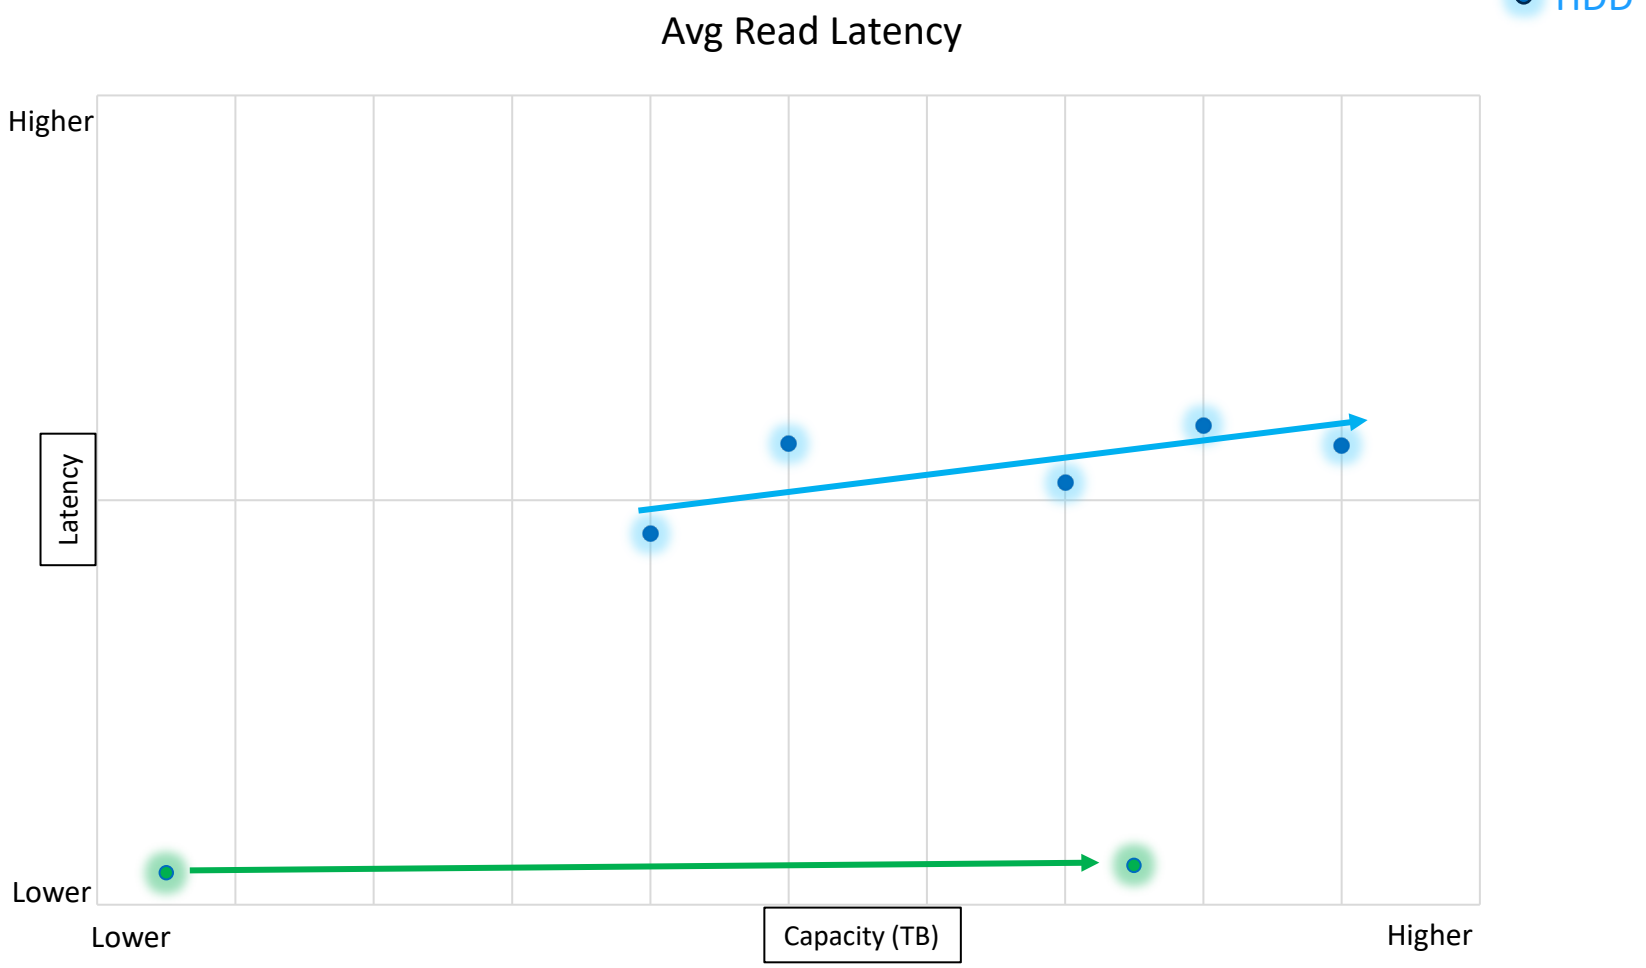

Latency Management is crucial to meet SLAs

Flash/SSD Implications

What's needed from Flash/SSDs Going Forward

Grow within incumbent applications

Power Efficiency	•Density w/power efficiency, i.e., Lower TiB/W Lower Device Idle Power Better Power Management for in-field Power/Performance tuning
On Par or better MBps Per TiB	•Scale throughput while maintaining latency consistency (Mixed Rd-Wr)
Workload Shaping	•Reduce E2E WAF (E.g. ZNS) for density scaling
Maintain PBW /TiB	•Maintain minimum endurance bar per segment

New Opportunities

Density	• 5-8x HDD
Cost	• 2-3x HDD Per TB
Thruput per TB	• 2-4x of HDD
Endurance in PBW	• Track to Writes

Landing Zone for SSDs – New Opportunity

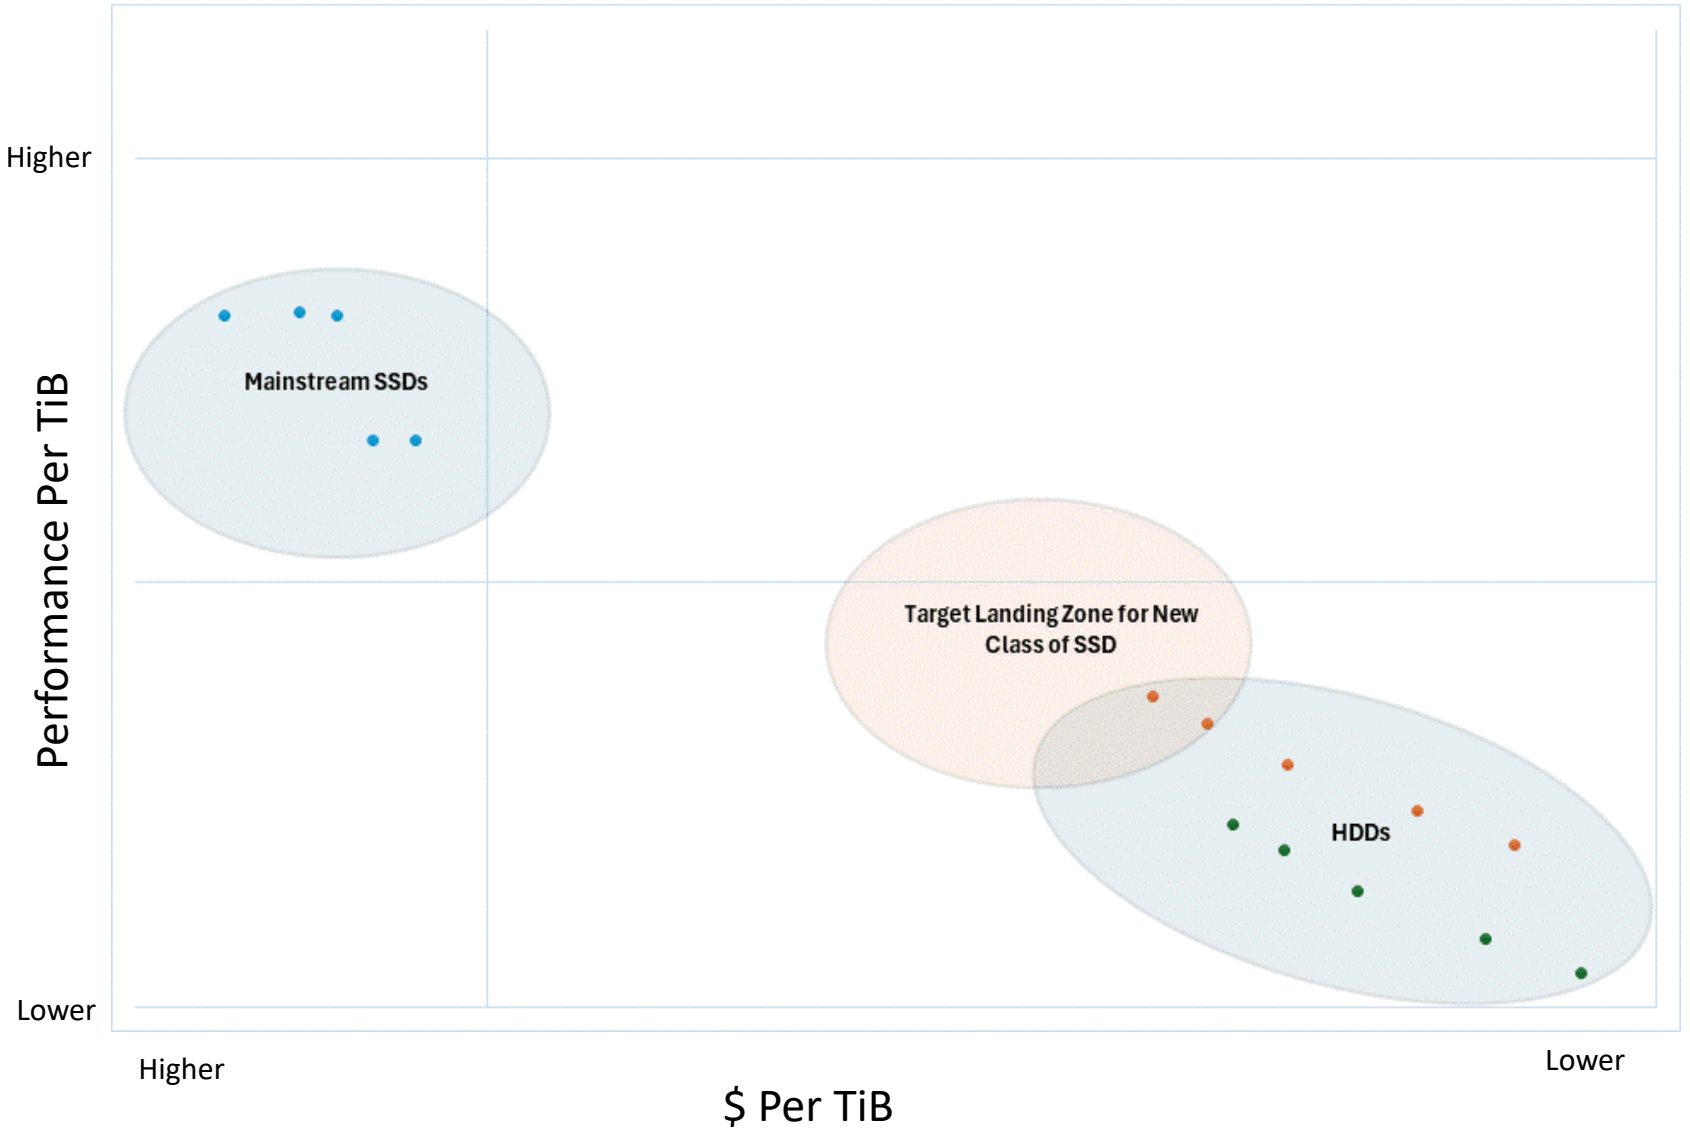

HDD Implications on Standard Storage

HDD Value for Standard Storage

Metric	Until Now	Projecting Forward
Capacity and Cost	• 2-4x Improvement in Capacity and Cost	HAMR and SMR Technologies
Performance Per TB	• Declining, but acceptable with Software Improvements	Decline is untenable
Workload Rating	• Stagnant but acceptable	Not sufficient
Power Efficiency	• Helium was the last improvement	Needs to Improve

HDD Performance Issue

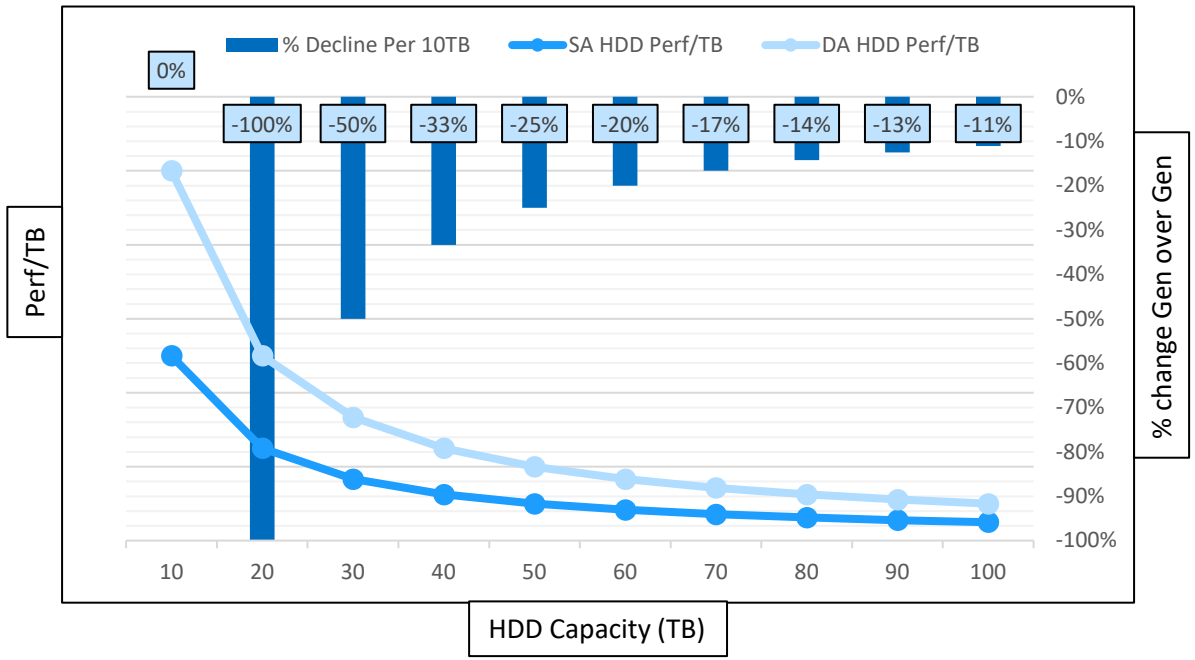

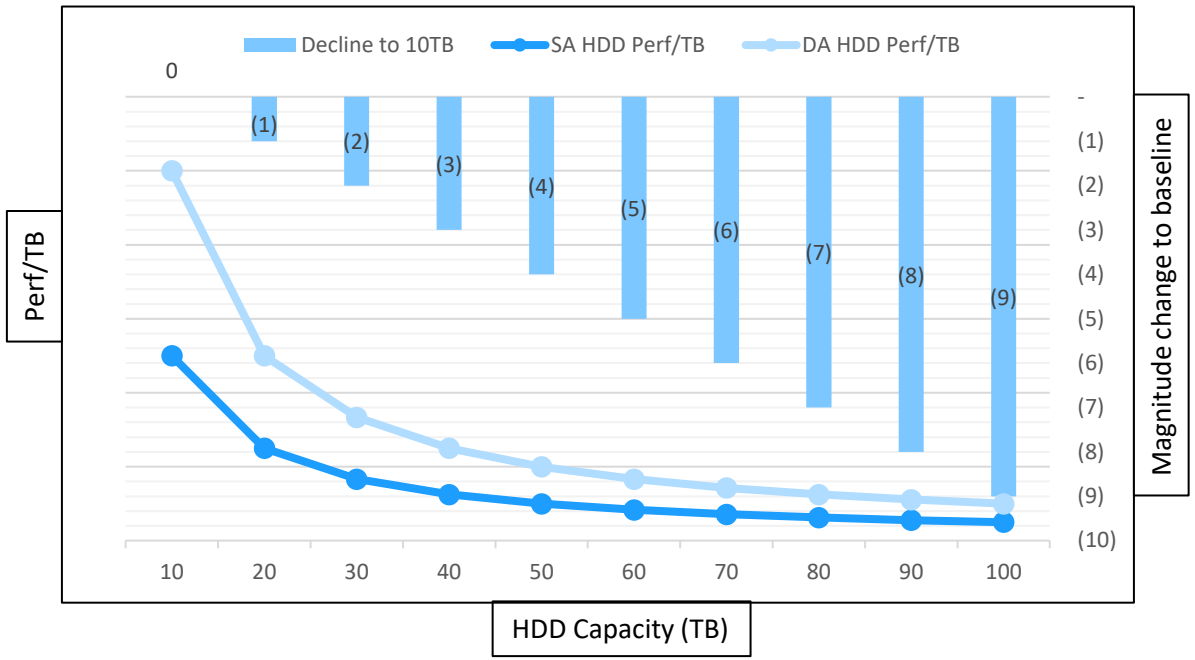

- HDD Performance metrics Per TB have been degrading
- Software improvements have compensated for HDD performance degradation at the Service level
- DA HDDs certainly help, but that benefit is also over a finite range of capacities
- The Industry needs to identify additional improvements to improve IOPS/TB and/or MBps/TB metrics to be able to provide the same level of service

Impact of HDD Performance on HDD Utilization

- 20TB HDD with host more customers with more data and get more requests compared to a 10TB HDD
- More Capacity = Busier Drive = More Utilization
- Approaching 100% busy rate = Not being able to use HDD capacity (Stranded Capacity)
- Hence the need for
  - **Better Performance**
  - **Higher Workload Rating**

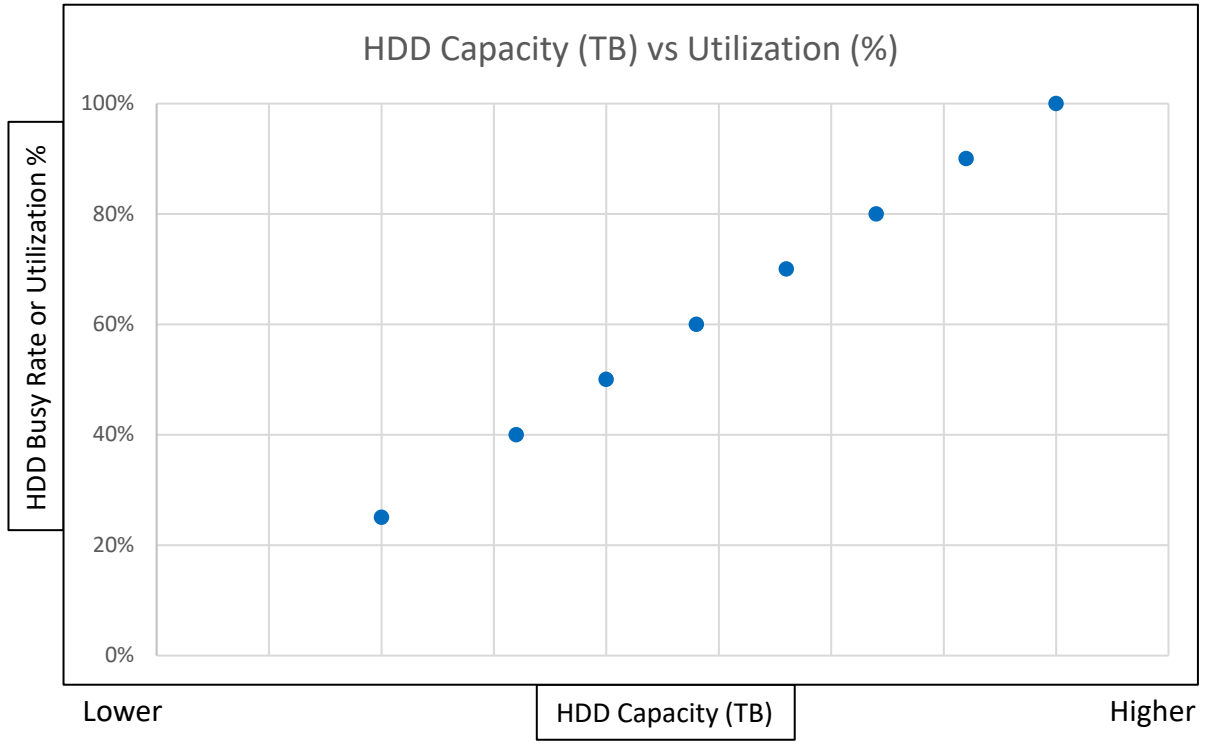

HDDs Going Forward – Deliver Better TCO

Grow within incumbent applications

Capacity and Cost	• Need Consistent ADC Growth • Historic $/TB Decline needs to Continue
Performance and Workload Rating	• Need Performance Per TB Metrics to Improve • Increase workload Rating
Quality of Service	• Latency Management • Blast Radius Management (Per Head Capacity)
Power Efficiency	• Reduce Unit Power Consumption • Power Management Features

New Opportunities – Expand HDD Archive

TCO (Capacity, Cost, Power)	• 3-4x better than Standard HDD
Performance	• Negotiable
Latency	• ms, but Negotiable
Workload Rating	• Negotiable

Landing Zone for HDDs

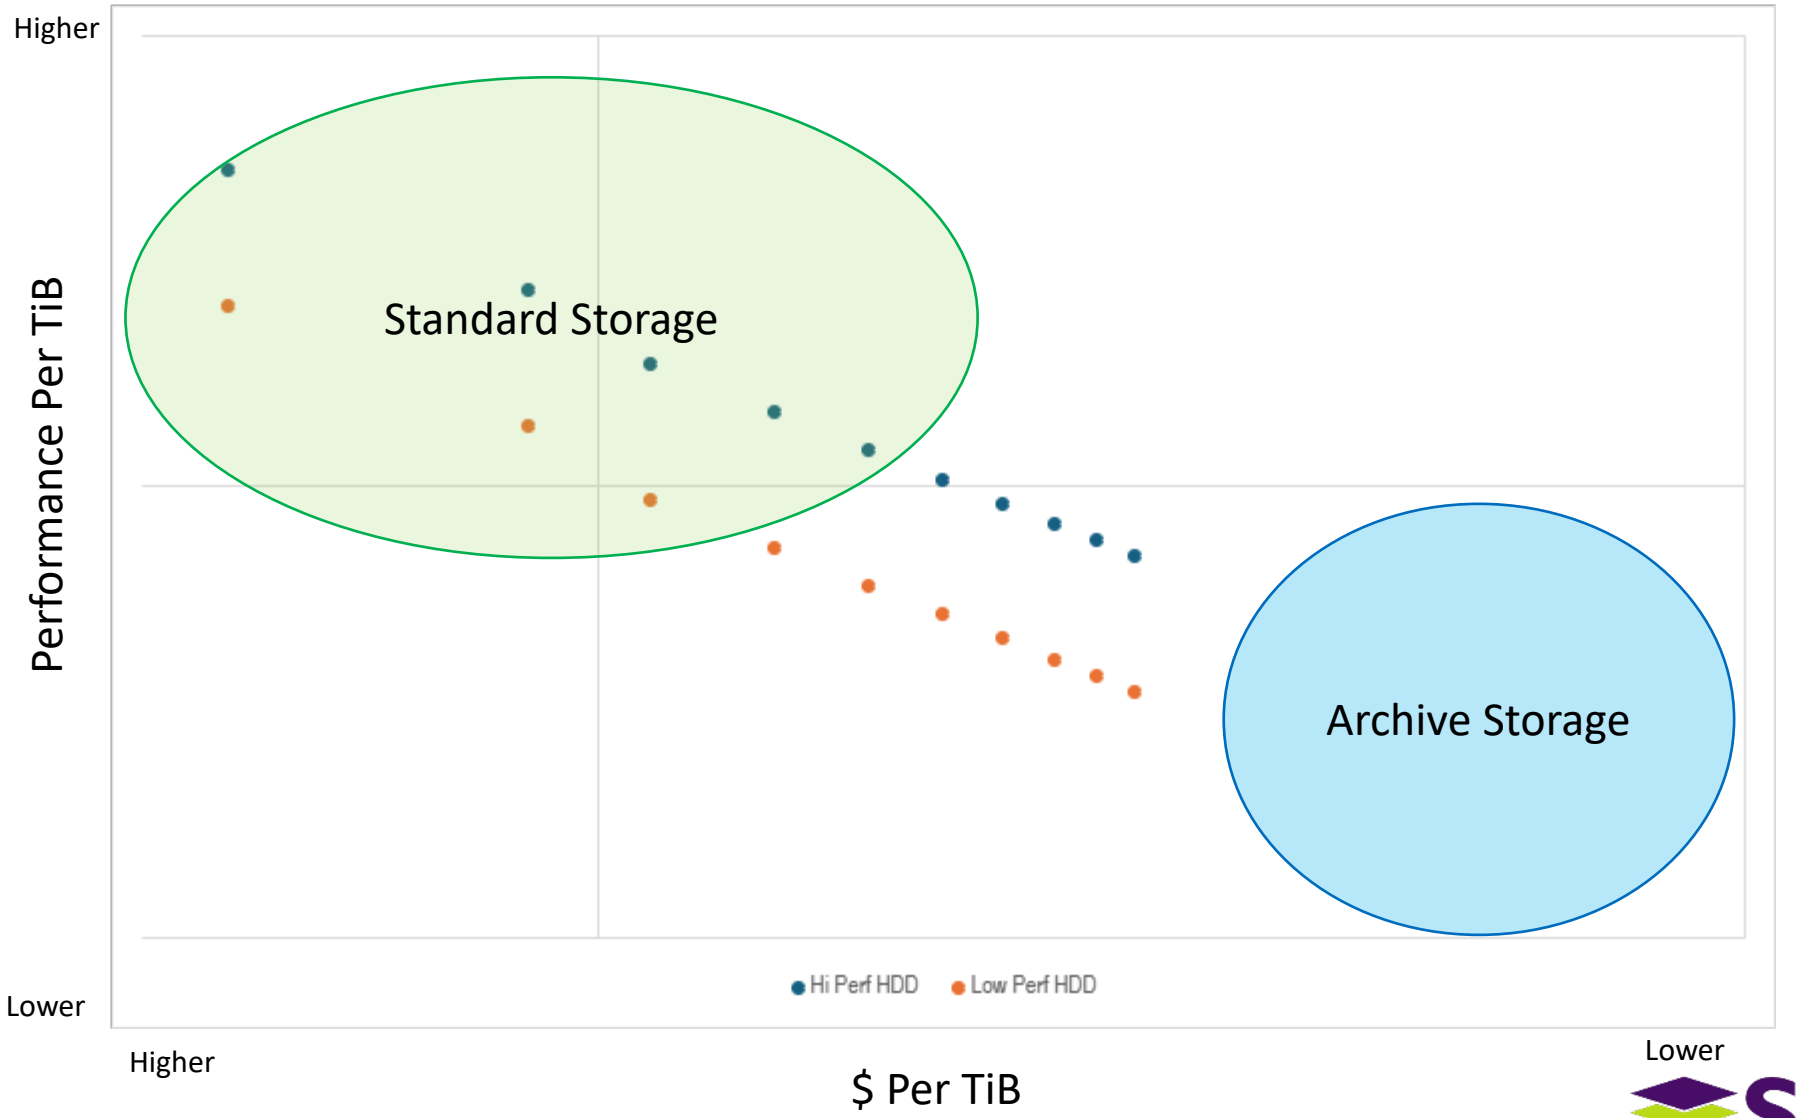

Storage Media Mapping for Future

Storage Media Landscape for Future

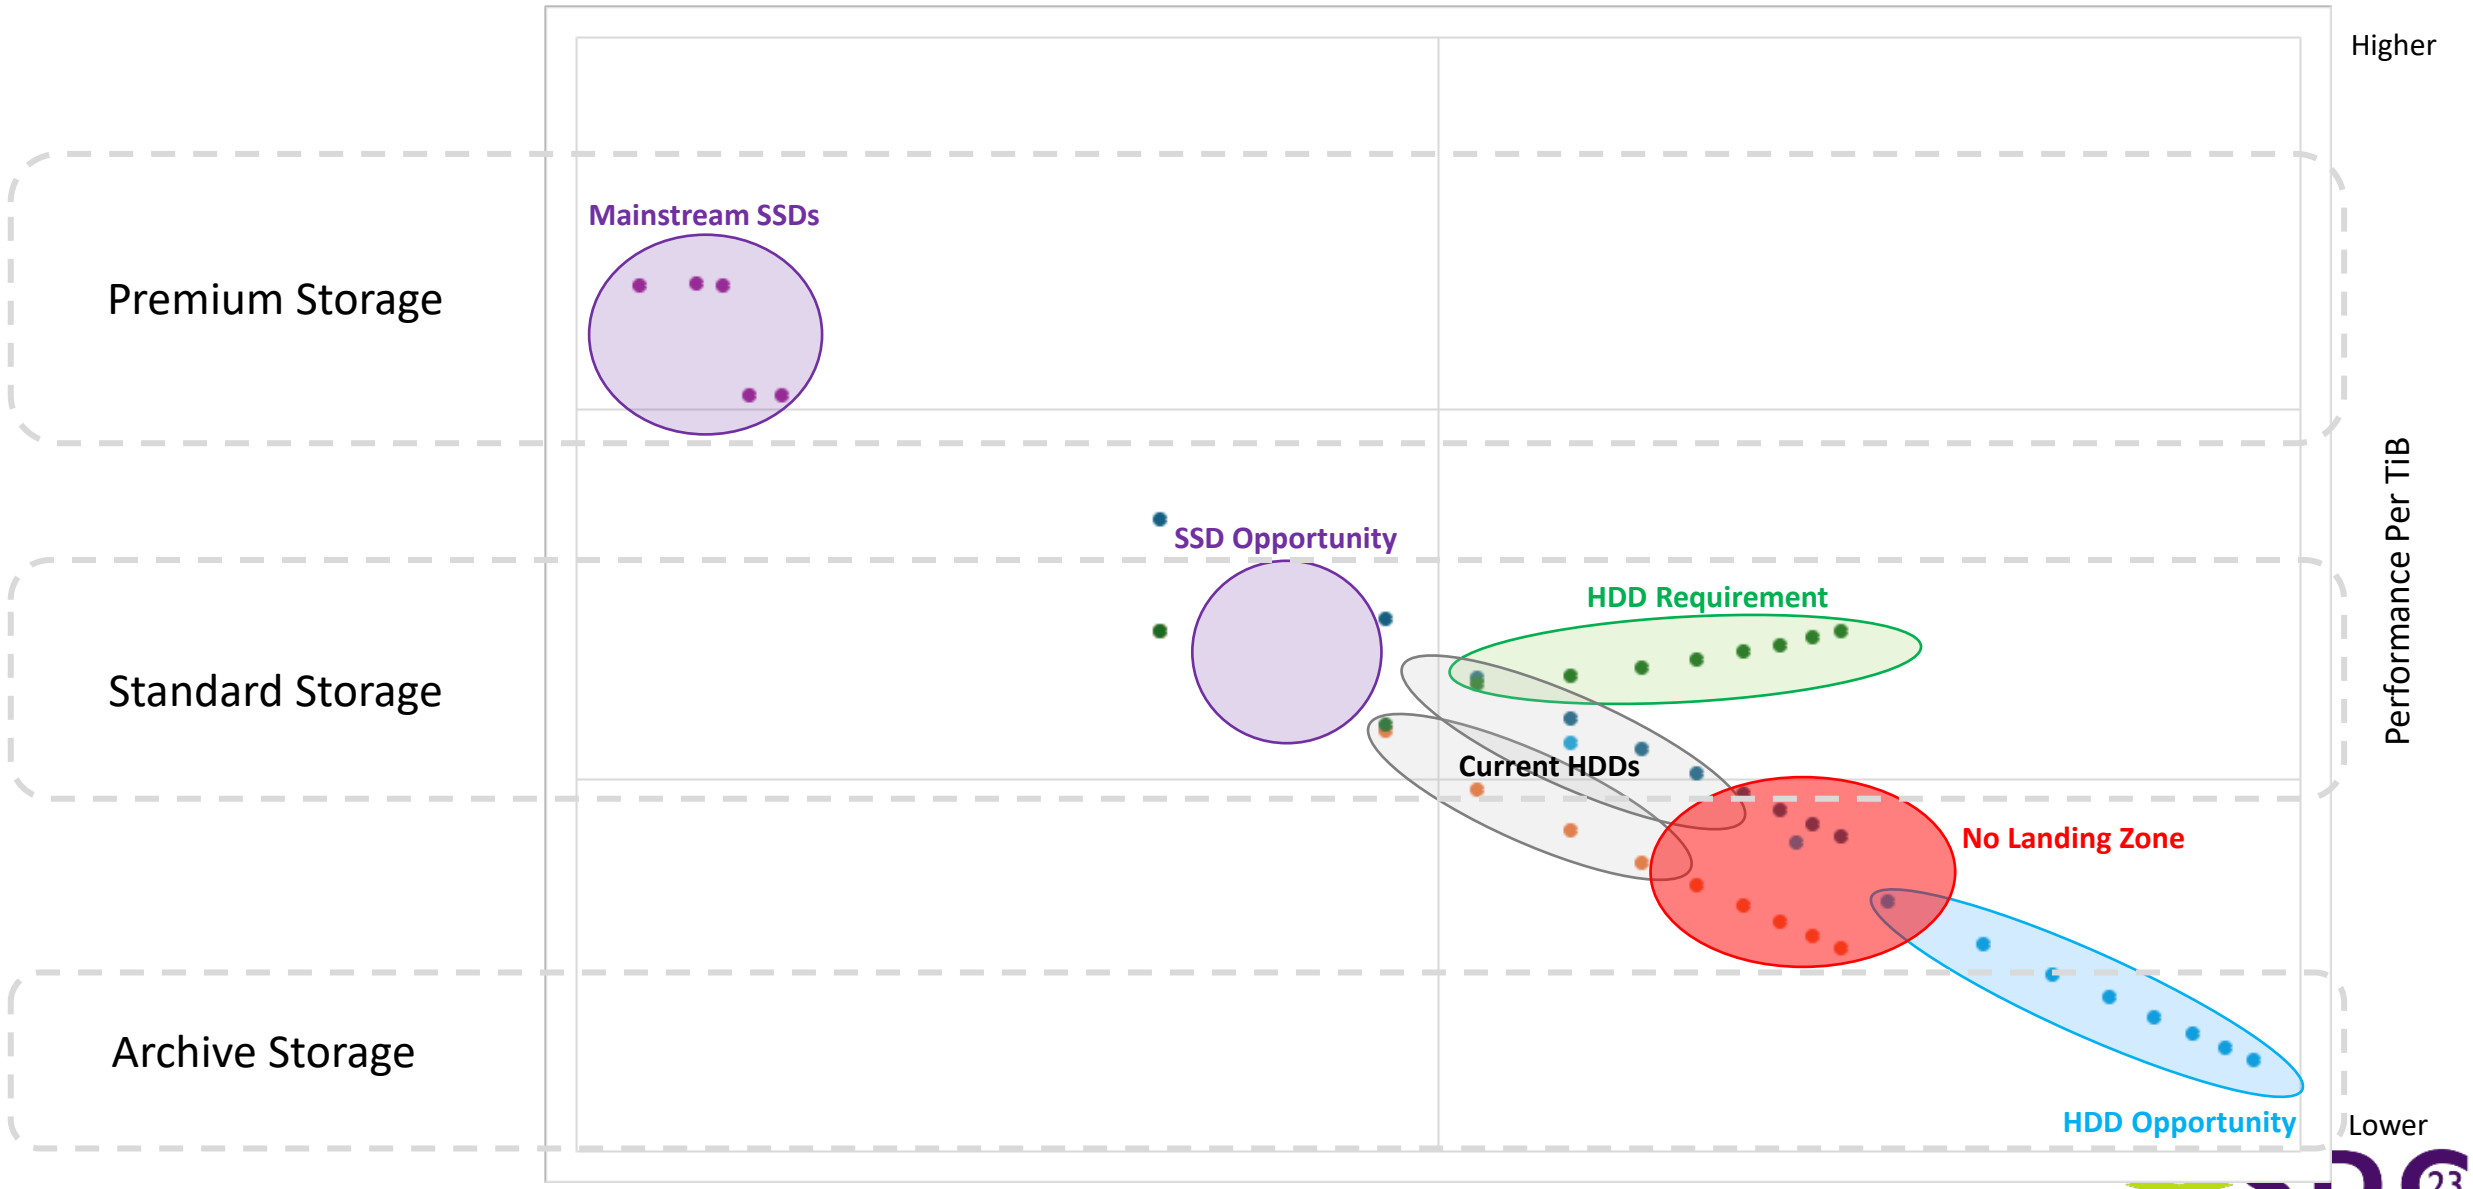

Questions

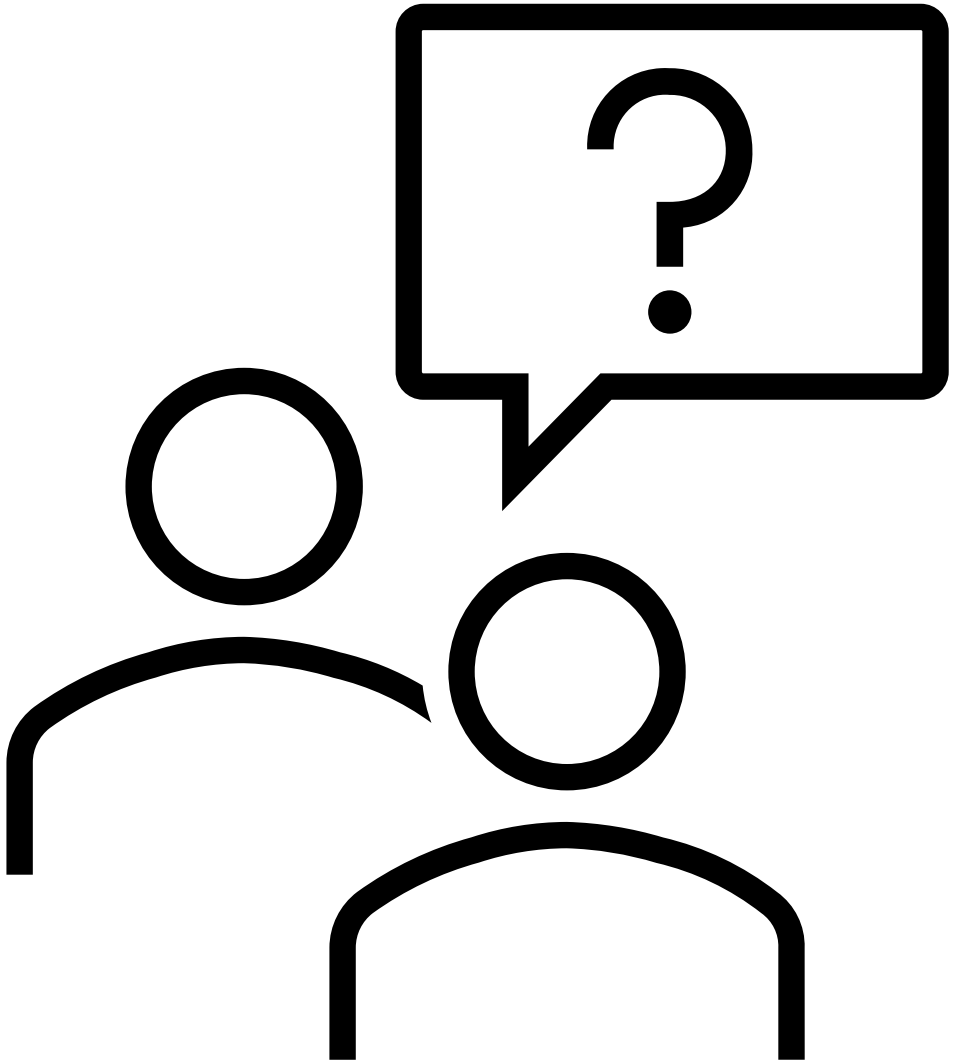

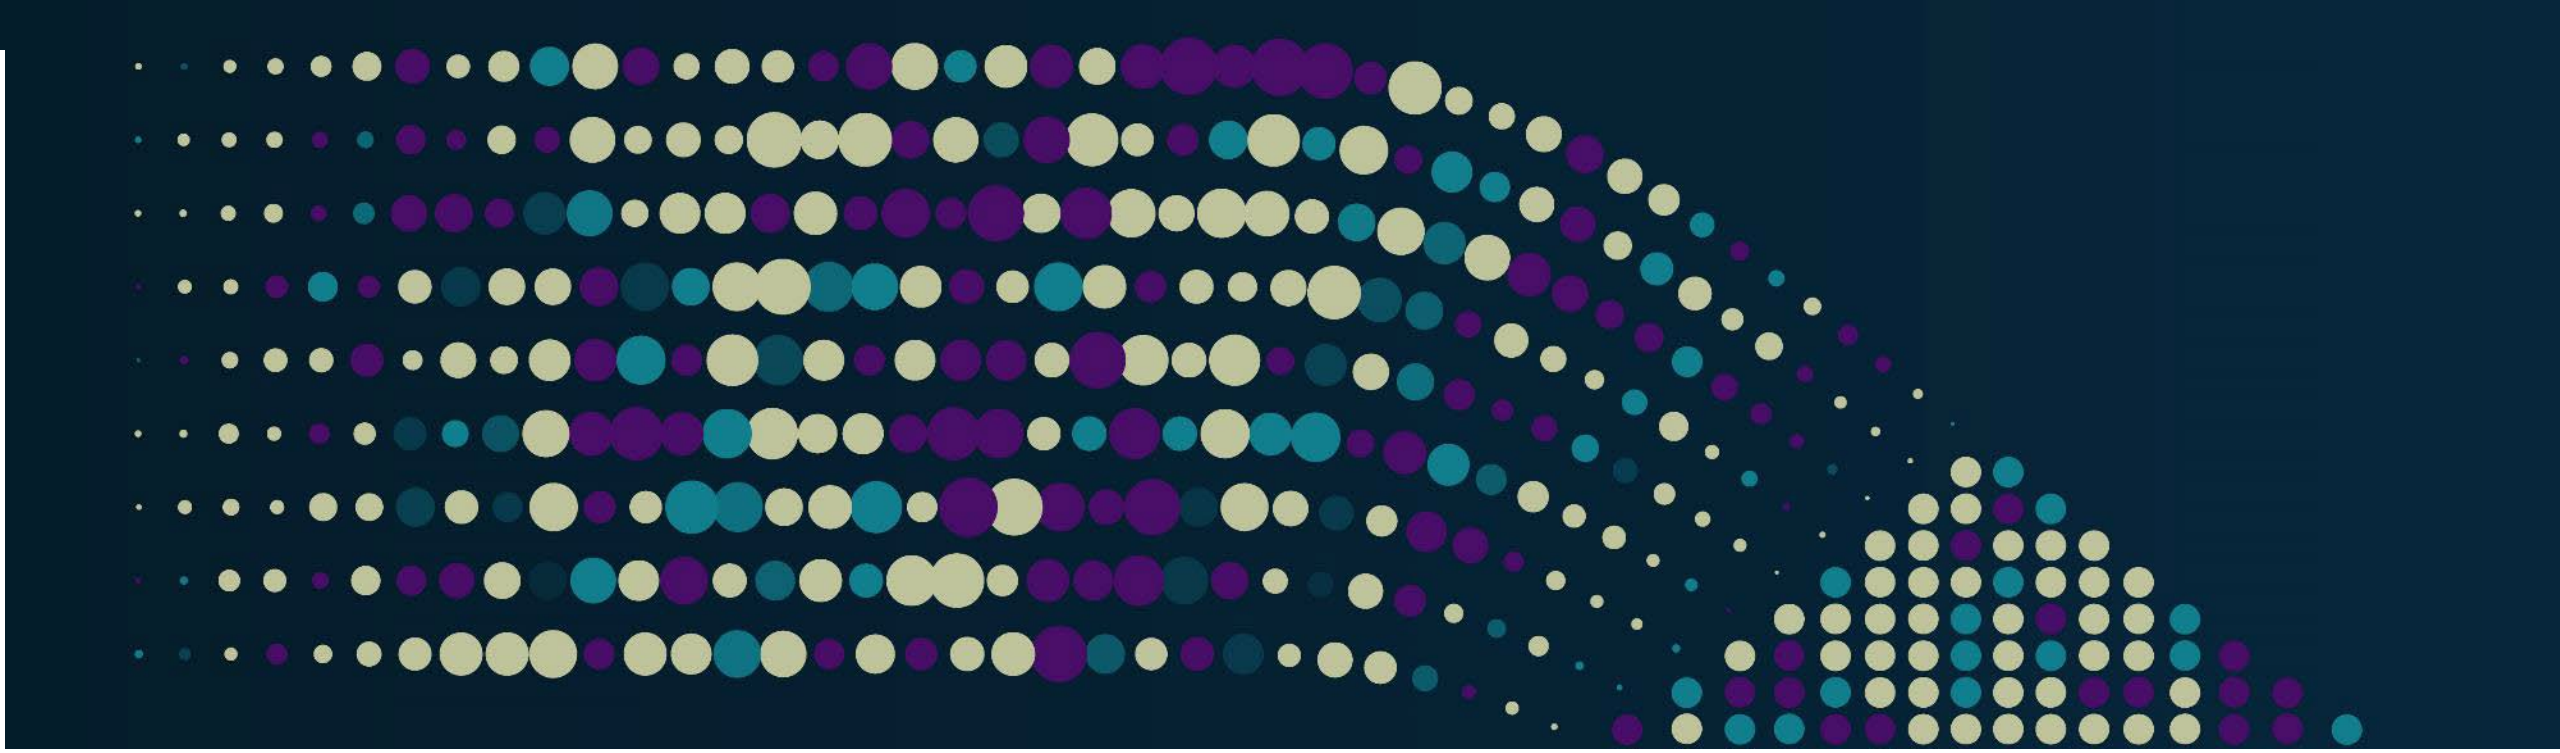

Please take a moment to rate this session.

Your feedback is important to us.

