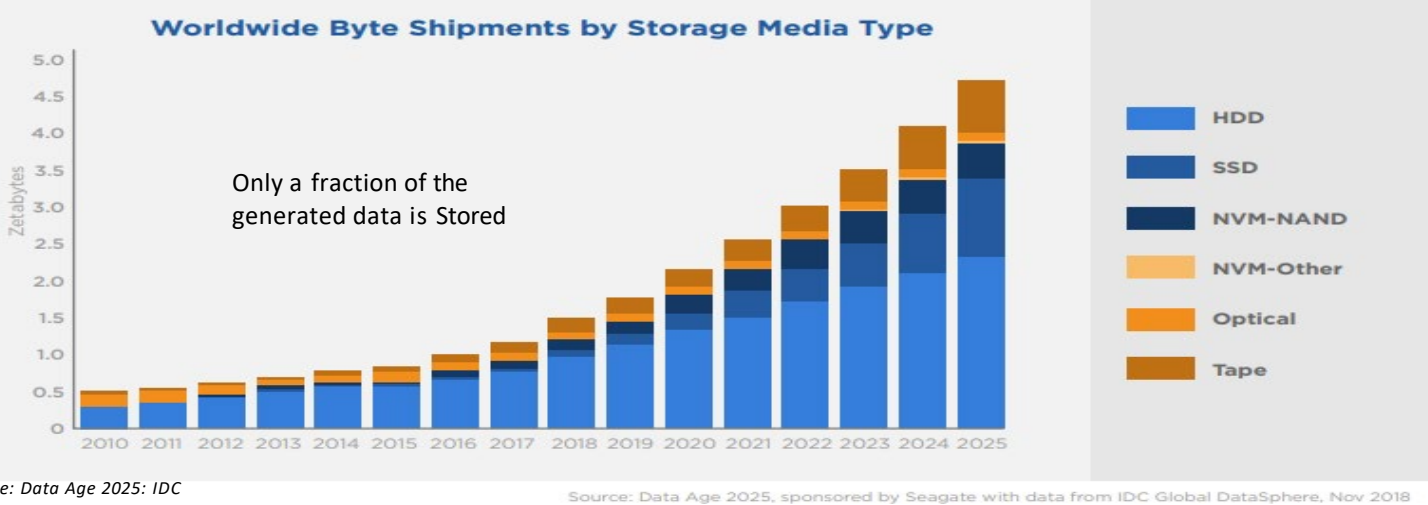




Design and Operation of Exascale Archives in Azure

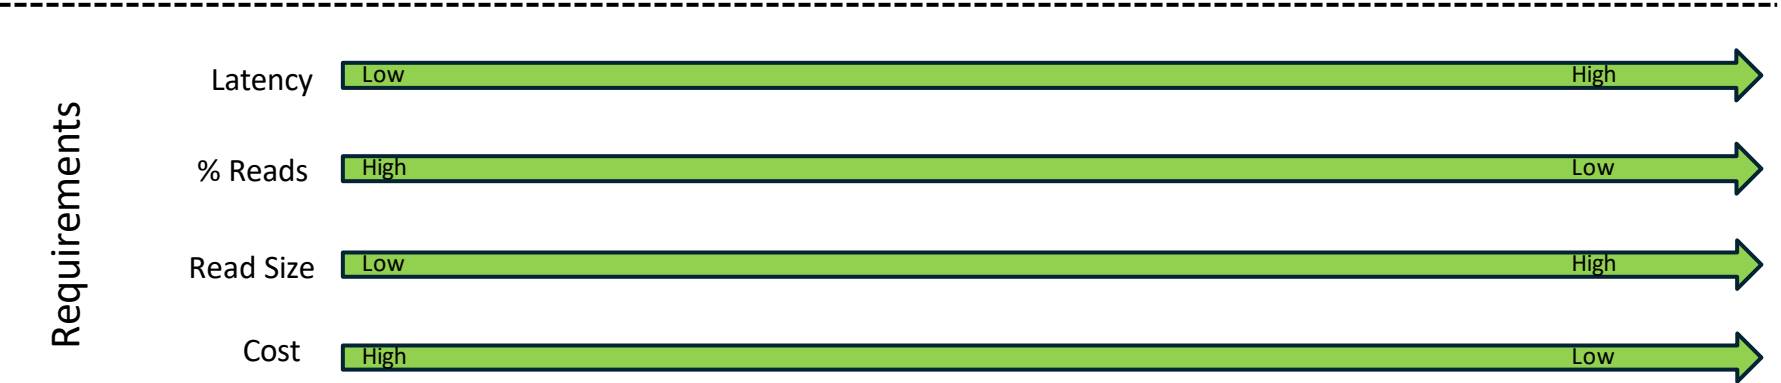
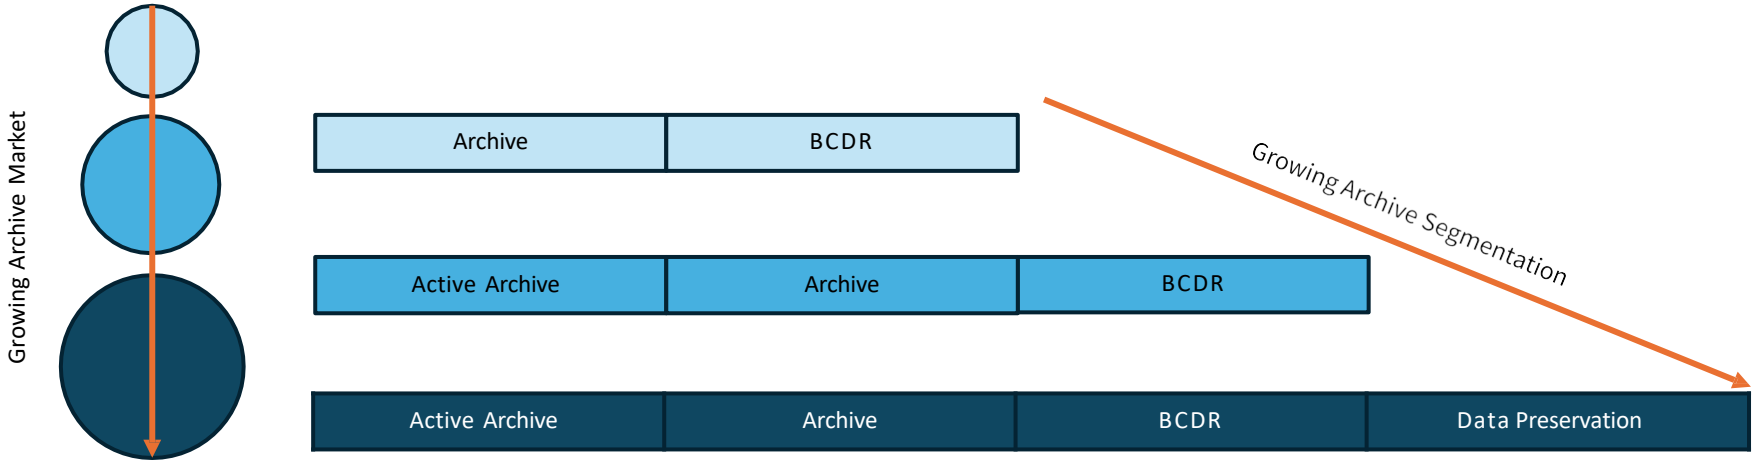
Aaron Ogus
Shashidhar Joshi

Data Growth and Storage



- Archive Storage is part of this Growth
- Just within Azure, Archive Storage operates at massive scale
 - EBs of Data
 - Trillions of Objects
 - Spread across 100's of Data Center across the Globe
 - Billions of Requests (Monthly)

Archive Landscape and Growth



Data Storage in the Cloud

Benefits with Cloud Storage

- Choose storage based on your need
- End to End Data Protection
- Geo, Regional and Zonal redundancies
- Management and Monitoring Tools
- Scale up and down as needed
- High Availability and Resiliency

Type of storage account	Supported redundancy configurations	Supported values for the kind parameter	Supported values for the sku or SkuName parameter
Standard general-purpose v2	LRS / GRS / RA-GRS / ZRS / GZRS / RA-GZRS	StorageV2	Standard_LRS / Standard_GRS / Standard_RAGRS / Standard_ZRS / Standard_GZRS / Standard_RAGZRS
Premium block blobs	LRS / ZRS	BlockBlobStorage	Premium_LRS / Premium_ZRS
Premium file shares	LRS / ZRS	FileStorage	Premium_LRS / Premium_ZRS
Premium page blobs	LRS	StorageV2	Premium_LRS
Legacy standard general-purpose v1	LRS / GRS / RA-GRS	Storage	Standard_LRS / Standard_GRS / Standard_RAGRS
Legacy blob storage	LRS / GRS / RA-GRS	BlobStorage	Standard_LRS / Standard_GRS / Standard_RAGRS

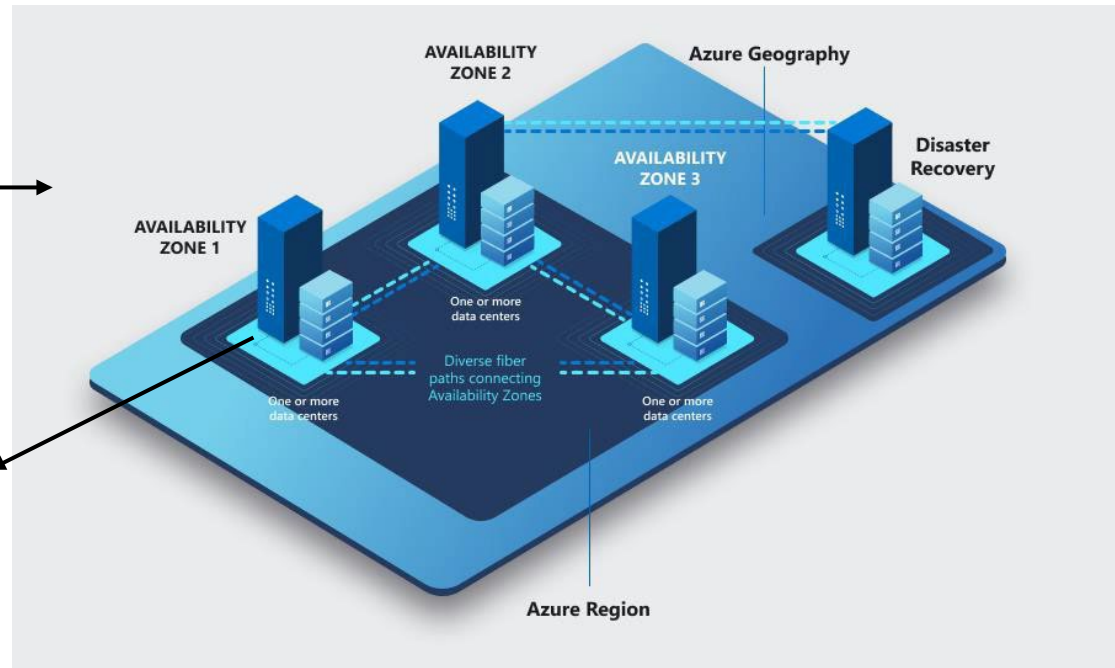
Data Storage Tiering in the Cloud

Customer can

- Set up Access Tier based on need
- Change Blob Tier Settings
- Lifecycle Management Policies
- Geo, Regional and Zonal redundancies

	Hot tier	Cool tier	Cold tier	Archive tier
Availability	99.9%	99%	99%	99%
Availability (RA-GRS reads)	99.99%	99.9%	99.9%	99.9%
Usage charges	Higher storage costs, but lower access and transaction costs	Lower storage costs, but higher access and transaction costs	Lower storage costs, but higher access and transaction costs	Lowest storage costs, but highest access, and transaction costs
Minimum recommended data retention period	N/A	30 days ¹	90 days ¹	180 days
Latency (Time to first byte)	Milliseconds	Milliseconds	Milliseconds	Hours ²
Supported redundancy configurations	All	All	All	LRS, GRS, and RA-GRS ³ only

Storing EBs of Data in the Cloud




Racks consist of different storage technologies

- SSDs
- HDDs
- TAPE



Replicate Data Across 3 Storage Clusters within a Single Region

Cloud Storage Flow and Features



- Erasure coded for data protection
- Stored on Magnetic Media
- Rehydrate data to online tier when requested
- Periodic data scrubbing to correct errors, ensure data accuracy and reliability
- End to End Data Protection
- Migrate to new media at EOL

Archive Workloads

Observations from Workload

- Writes dominate
- Reads are infrequent
- Small reads require low latency random access
- Large reads require good throughput
- Archive Storage system needs dynamic provisioning to account for workload variances

Challenges with current storage technologies

- Mechanical overheads lead to latencies
- Environmental conditions limit deployment capabilities
- Uncertainty with roadmaps, capacities and costs
- Need for media migrations at EOL
- Opportunity for new storage technologies

Silica for Archive Storage Technology

Storage System Disaggregation

Scale Media, Writers, Readers, Library independently

Glass as Media, Optical Writers and Readers, Library and Robotics for Media Handling

Unique Media Properties

Abundant supply

Inexpensive

Passive Storage – No Power


Resilient to Environmental Conditions, EMP Proof

Recyclable


Immutable, No Bit Rot or Data Corruption

Capacity and Performance


- Meet Storage Growth capacity demand via better storage density
- Performance (Random IO) per TB metric better than currently available Archive technologies
- Better TCO
- Sustainable



Glass Media



Writer



Reader



Media Storage and Handling

Meeting the Industry Growth

- The Industry will need more storage capacity to support future growth
- Magnetic Media Storage growth continues which is very promising and much needed for the foreseeable future
- In addition, the industry continues to look for new emerging technologies that are better suited for scale, economics and sustainability
 - DNA Storage
 - Folio Photonics
 - Cerabyte