Quantifying the value of NVMe SSD in all stages of AI Datapipeline

Solidigm

Kiran Bhat

the Future of Memory and Storage

Agenda

- Overview of AI/ML Data pipeline
- Data movement in AI/ML Cluster
- FIO Based workload performance
- SSD Range of usages
- Summary

©2024 Conference Concepts, Inc.
All Rights Reserved

Typical AI Data pipeline

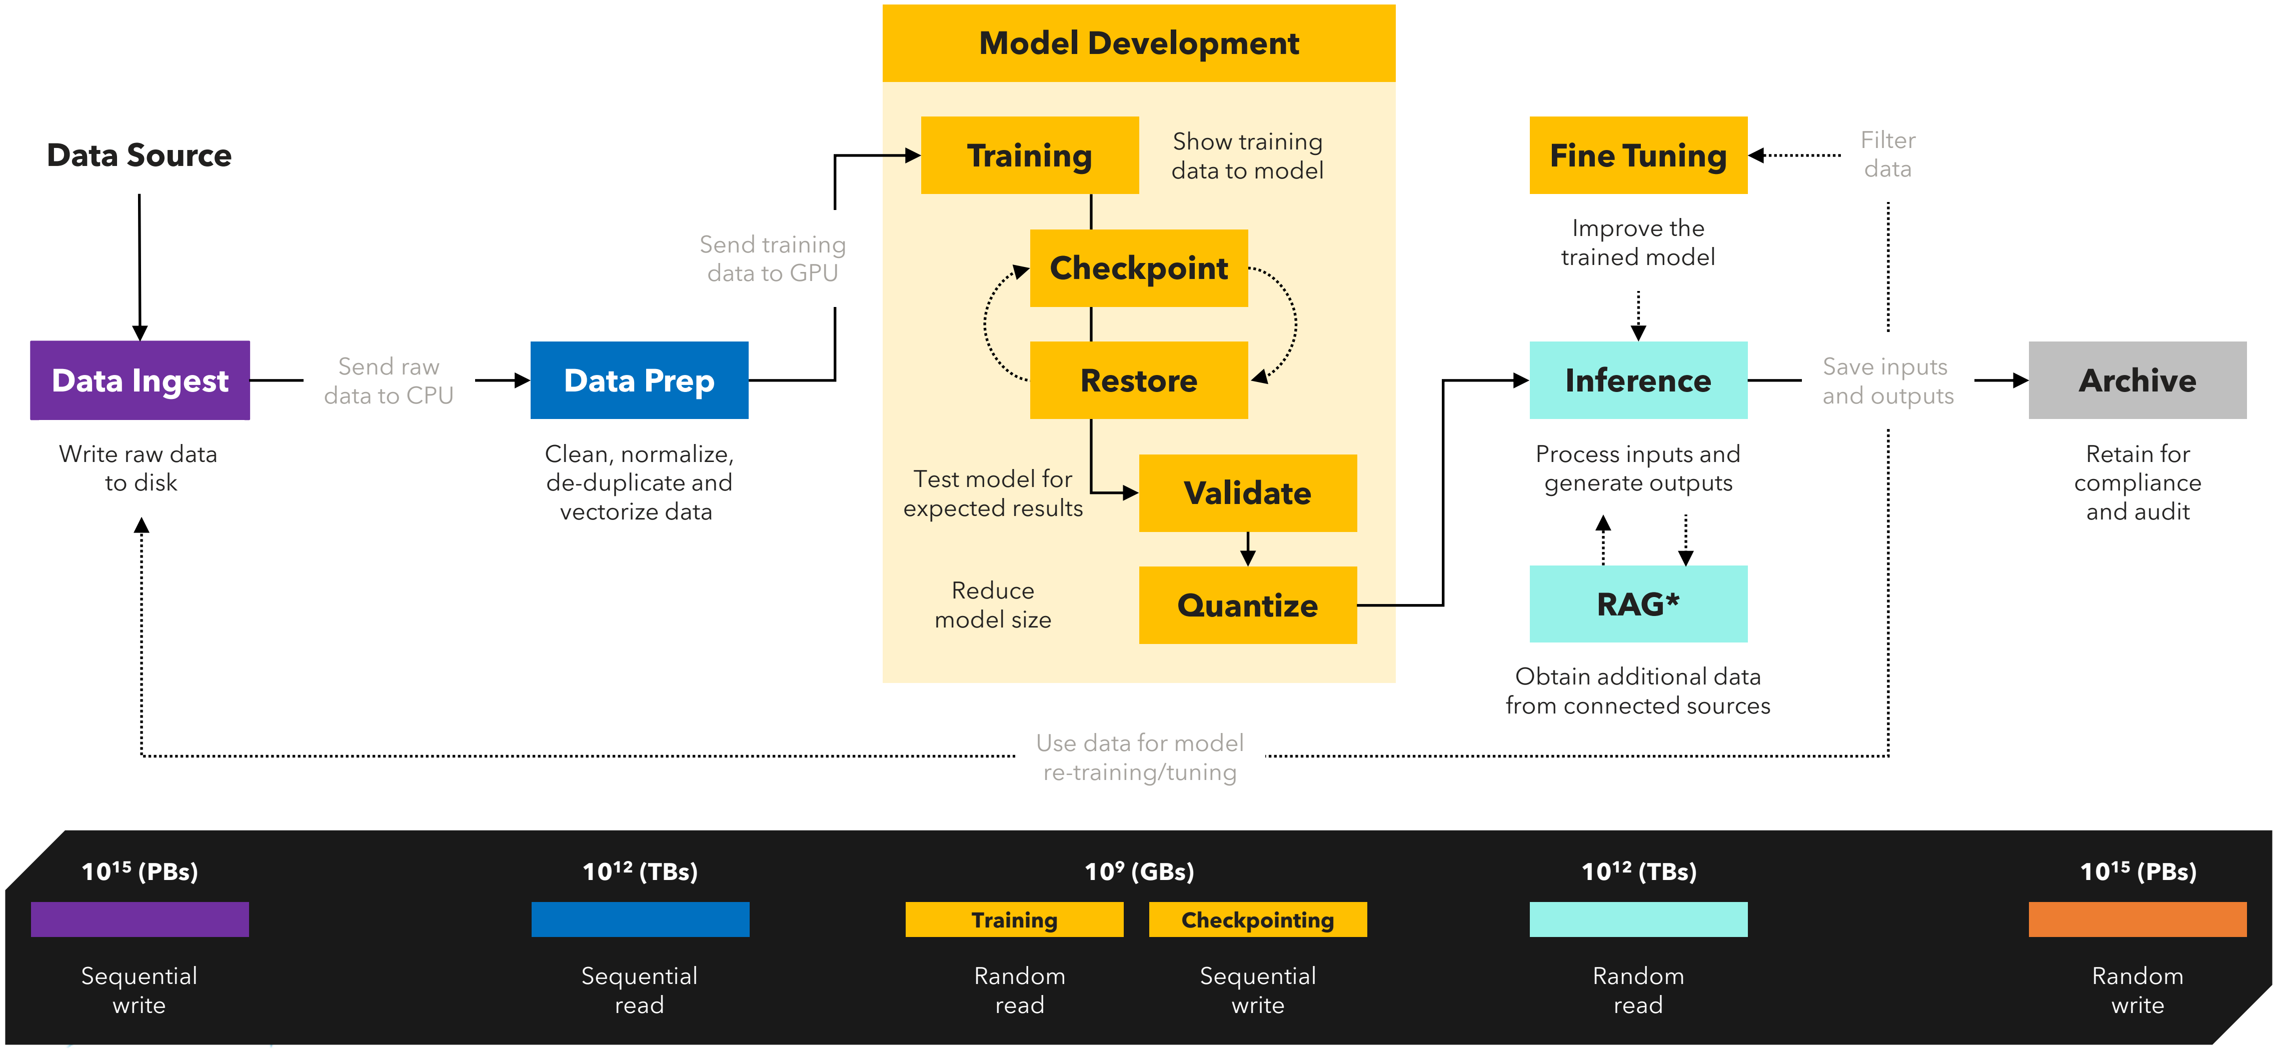

©2024 Conference Concepts, Inc.
All Rights Reserved

Data Movement in an AI Cluster

Compute Server
NAND Cache Tier
Object Storage Tier
Networking

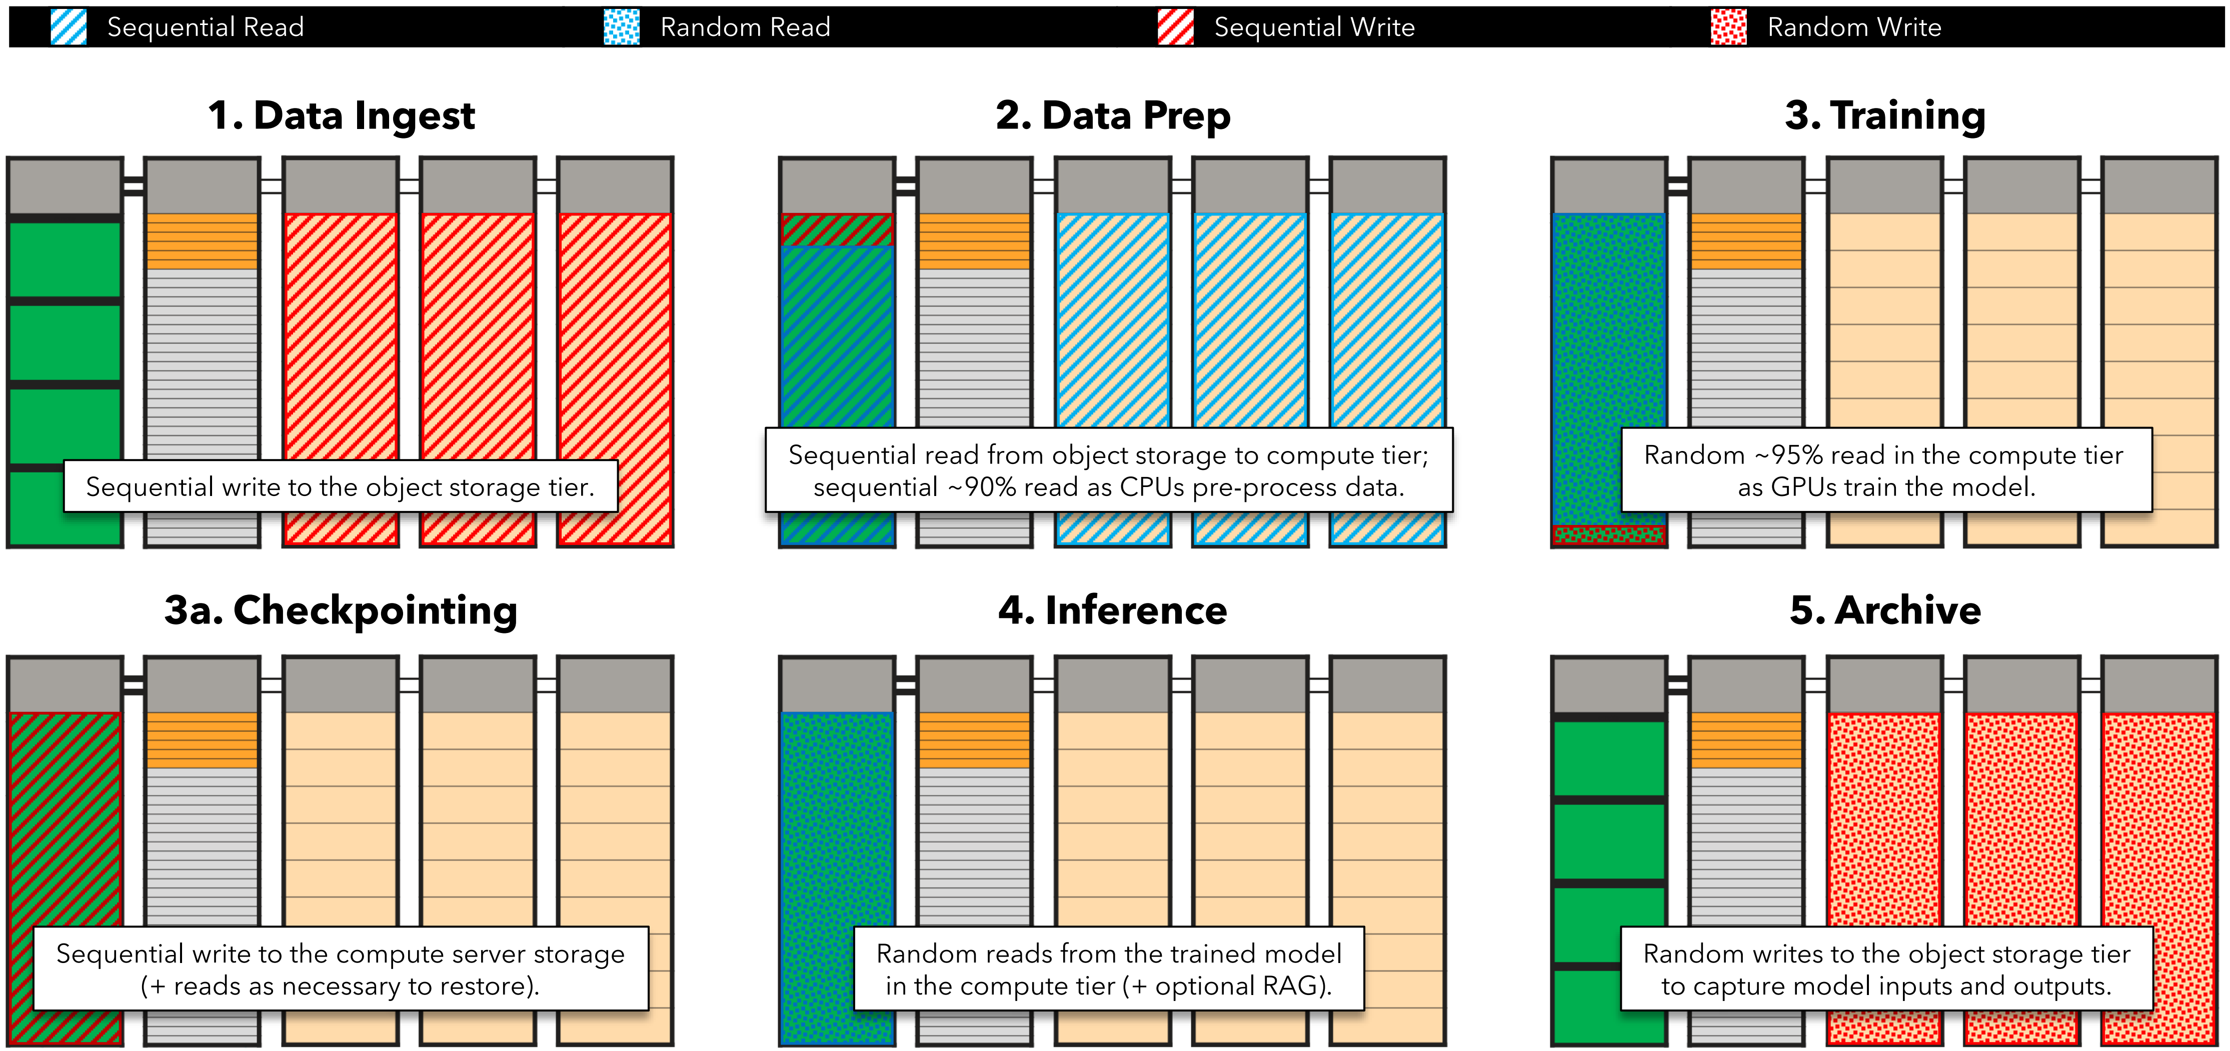

©2024 Conference Concepts, Inc.
All Rights Reserved

FIO based Workload for AI/ML Performance Analysis

	SSD Product:	HDD	Solidigm P5336	Solidigm P5520	Solidigm PS1010
	PCIe Gen:	SATA	PCIe Gen4	PCIe Gen4	PCIe Gen5
	Capacity:	24TB	30TB	7.68TB	7.68TB
		GB/s	GB/s	GB/s	GB/s
Ingest	32KB Sequential Write QD32	0.17	3.51	4.30	9.03
	128KB Sequential Write QD32	0.28	3.51	4.46	9.17
Preparation	4KB Random Read 90% QD32 / 128KB Sequential Write 10% QD32	0.015	2.45	3.01	6.72
Training	4KB Random Read 95% QD32 / 128KB Sequential Write 5% QD32	0.013	2.60	3.10	7.01
Checkpointing	128KB Sequential Read QD32	0.29	7.01	7.16	14.80
	128KB Sequential Write QD32	0.28	3.51	4.46	9.17
Inference (Decision)	4KB Random Read QD512	0.001	4.21	4.54	11.40

©2024 Conference Concepts, Inc.
All Rights Reserved

SSD **range of usages** in AI data pipelines

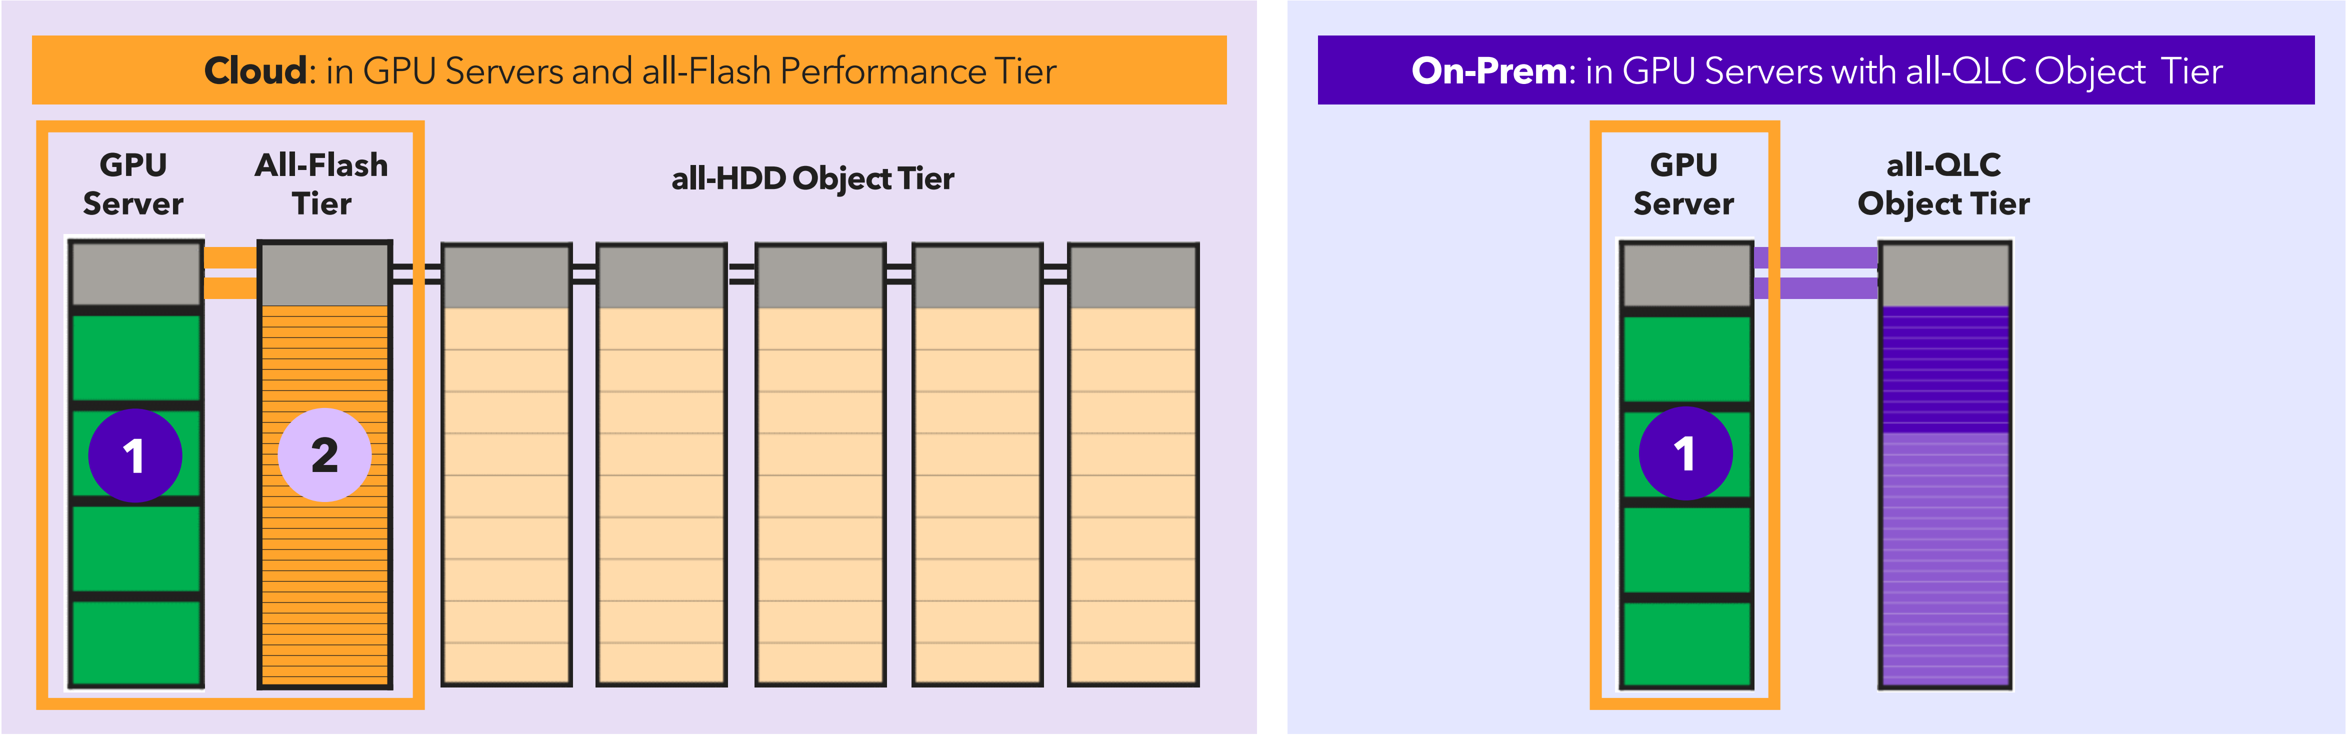

1. As **NVMe Data Cache Drive** in GPU servers
2. In **All-Flash High-Performance Tier** supporting lower-performing HDDs

©2024 Conference Concepts, Inc.
All Rights Reserved

QLC can efficiently **displace HDDs in AI data pipelines**

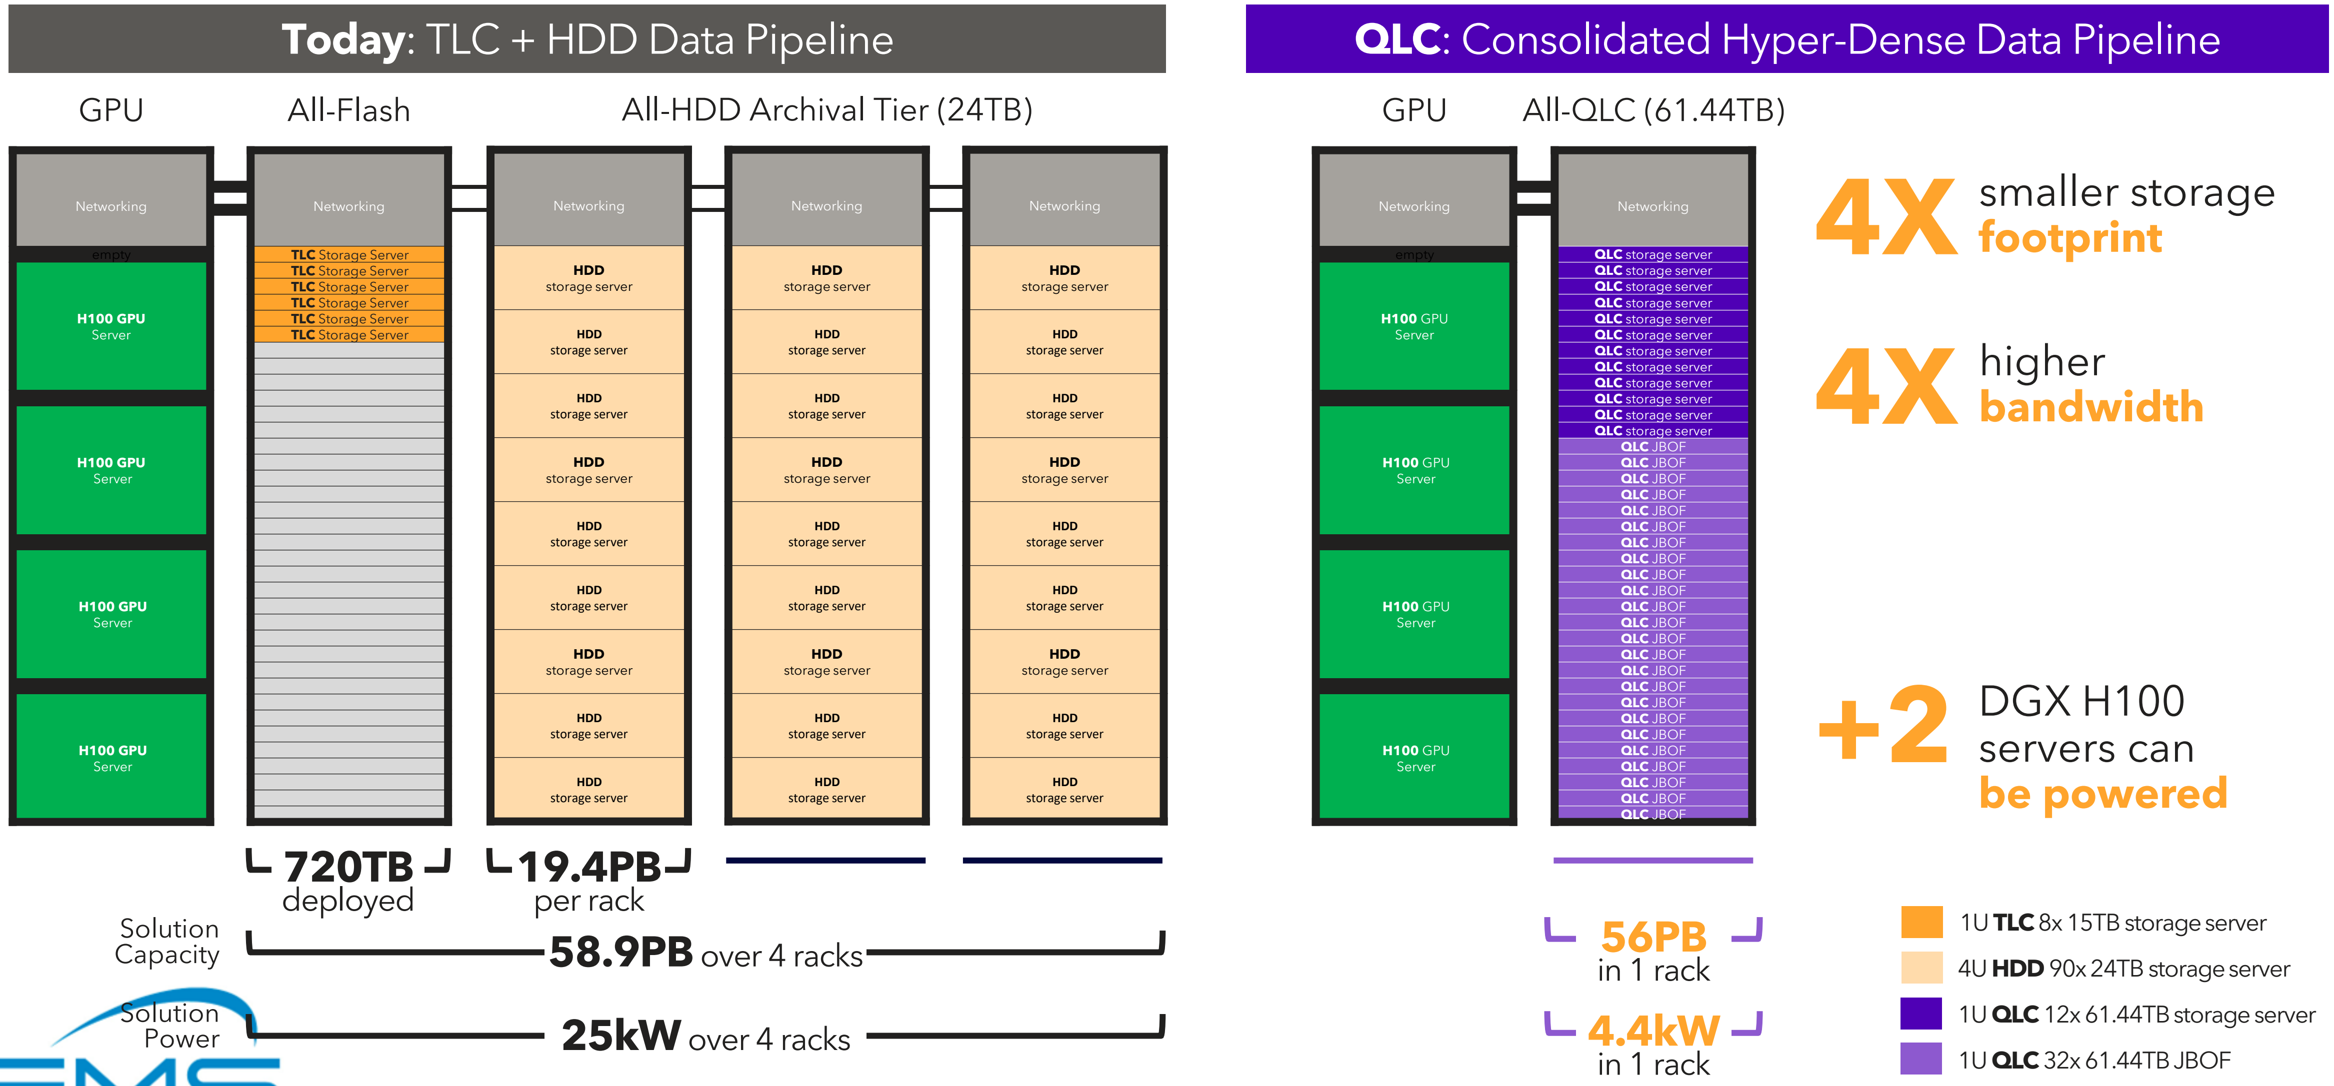

4X smaller storage **footprint**

4X higher **bandwidth**

+2 DGX H100 servers can **be powered**

- 1U **TLC** 8x 15TB storage server
- 4U **HDD** 90x 24TB storage server
- 1U **QLC** 12x 61.44TB storage server
- 1U **QLC** 32x 61.44TB JBOF

©2024 Conference Concepts, Inc.
All Rights Reserved

QLC SSDs Overcome HDD Limitations in AI

	Capacity	Performance		Reliability	Power Efficiency
	Data Set / Model Size	Training Performance	Ingest / Checkpointing	Continuity & Cost Efficiency	Scalability
	up to **24TB** widely available	up to **212 IOPS** random read	up to **298 MB/s** sequential write	**1.7%** annual failure rate[3]	**3.95 TB/W** efficiency[5]
	2.5x bigger	**4.5Kx** faster	**10x** faster	**8.5x** better	**2.7x** better
	up to **61.44TB** and counting	up to **1M IOPS** random read	up to **3,000 MB/s** sequential write	**< 0.2%** annual failure rate[4]	**10.78 TB/W** efficiency[5]

Superior storage improves AI development performance and reliability while saving power and rack space

©2024 Conference Concepts, Inc.
All Rights Reserved

Solidigm SSDs for AI Data Storage

Object Storage Tier | Compute + NAND Cache Tiers | Object Storage Tier

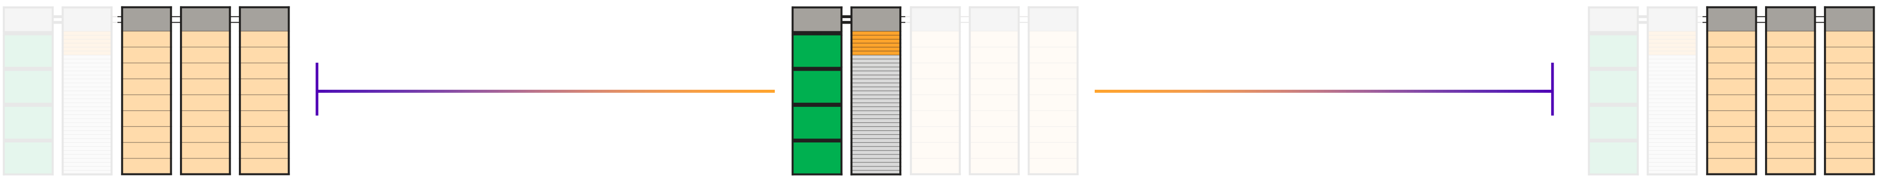

Stage	Data Ingest	Data Prep	Training	Checkpointing	Inference	Archive
Storage Requirements	High capacity and sequential write performance	Sequential read and write performance	Random read performance	Sequential write performance	Random read and write performance	High capacity
	Infrastructure efficiency: power utilization, cooling, rack space, and reliability					
Recommended Solution	 **Solidigm D5-P5336** PCIe 4.0 QLC SSD Capacity Read Write	 **Solidigm D7-PS1010** PCIe 5.0 TLC SSD Capacity Read Write — or —  **Solidigm D5-P5520** PCIe 4.0 TLC SSD Capacity Read Write — or —  **Solidigm D5-P5430** PCIe 4.0 QLC SSD Capacity Read Write				 **Solidigm D5-P5336** PCIe 4.0 QLC SSD Capacity Read Write
	Maximize TB / Watt	**Maximize Performance / Watt**				**Maximize TB / Watt**

FMS

©2024 Conference Concepts, Inc.
All Rights Reserved

Summary

- NAND SSDs offer better density, throughput and low power solutions compared to HDDs
- FIO based workload shows that NVMe SSDs are better suited for AI and ML applications
- Solidigm offers TLC and QLC based NVMe SSDs for AI and ML applications

©2024 Conference Concepts, Inc.
All Rights Reserved