Storage for AI

Presented By

Shashidhar Joshi, Principal, Microsoft
Swapna Yasarapu, Principal, Microsoft

AI Model Training

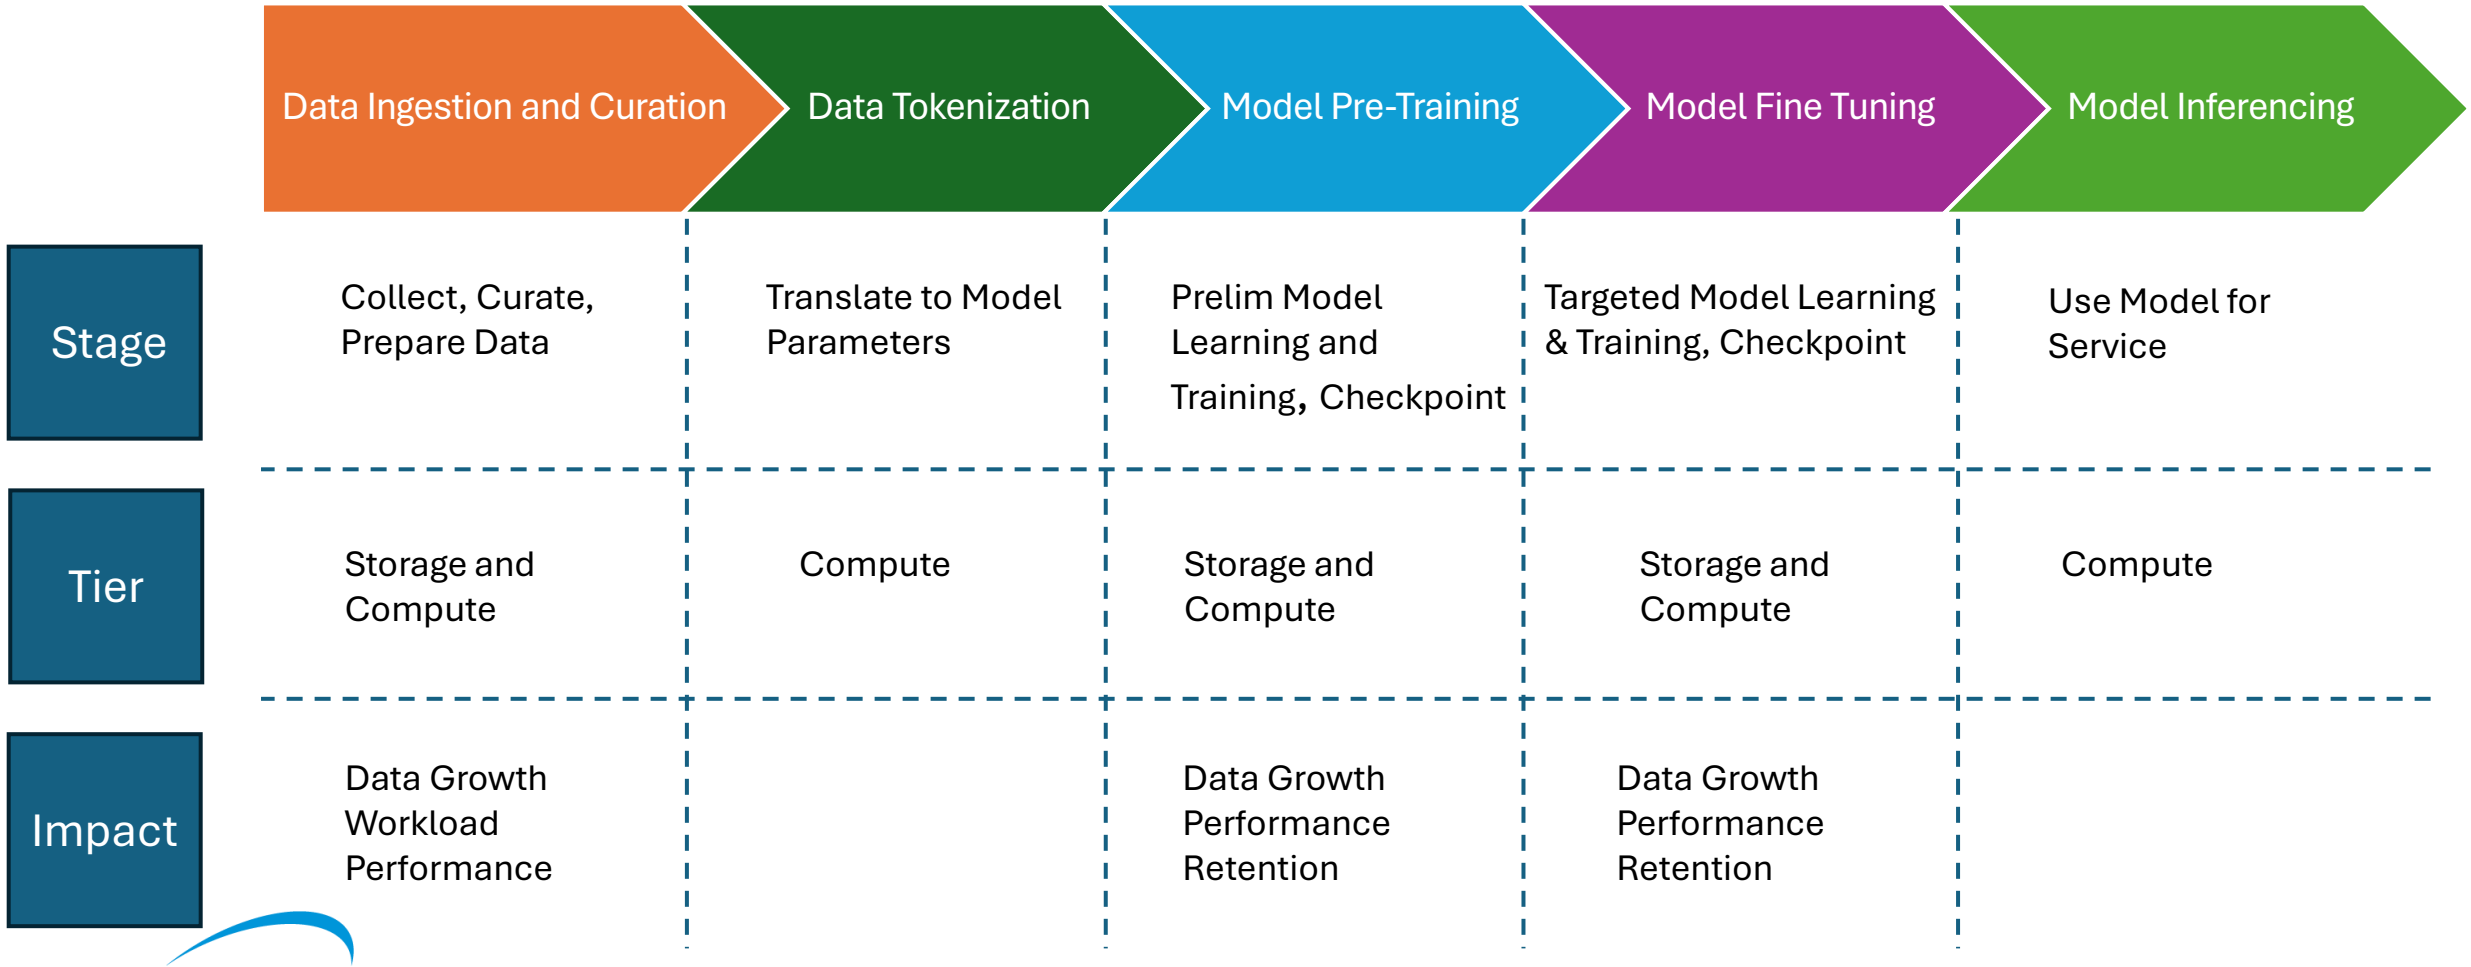

AI Data's Journey through Storage Tiers

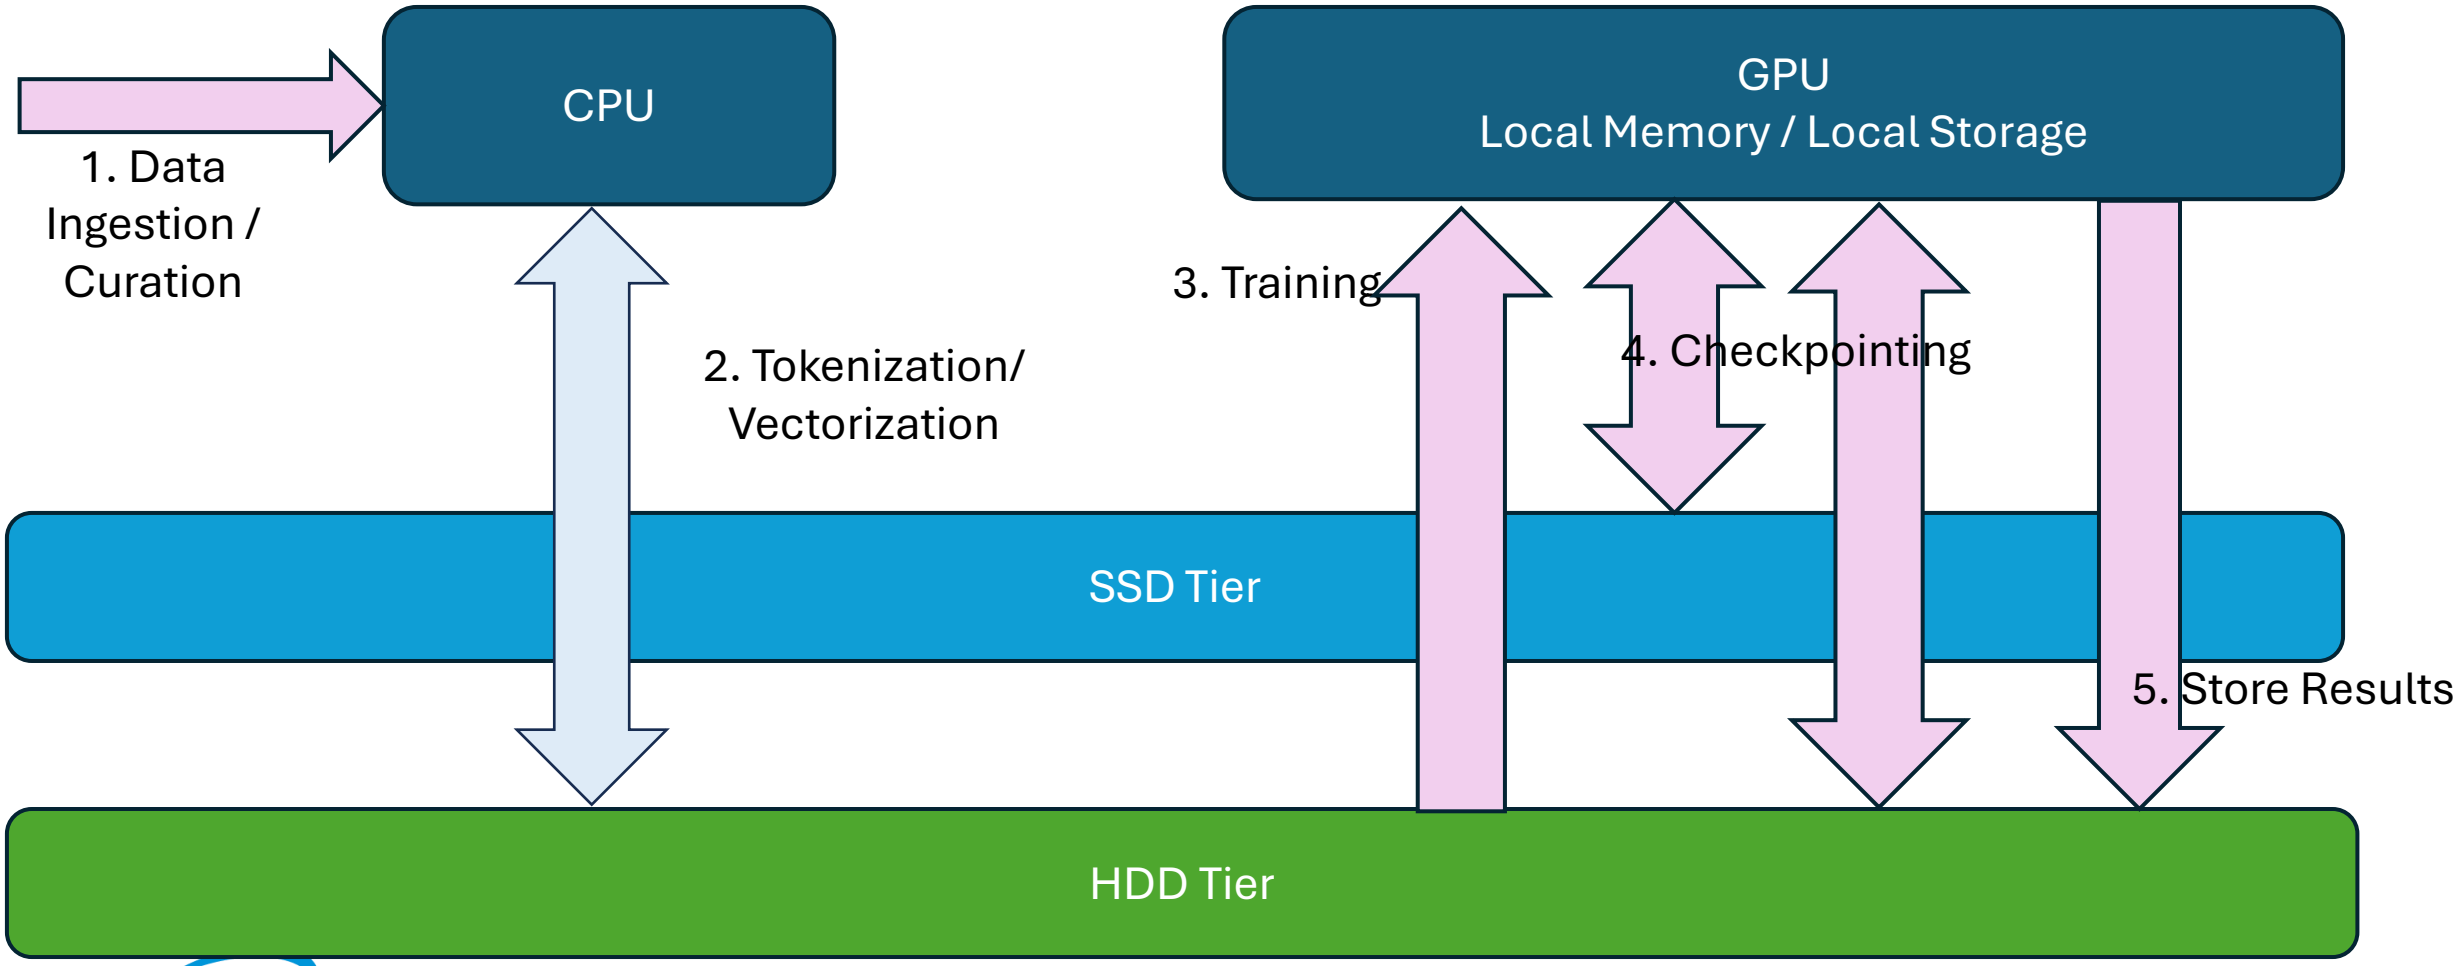

- Ingested and Curated Data is stored in our Blob Storage Service, mostly supported by HDDs

- Most Recent Checkpoint lands on SSD
- Previous Checkpoints land on HDDs for Longer Retention
- Scales with GPUs and Stop/Start Cycles

AI Storage Demand and Performance

	Data Ingestion	Data Curation	Training and Checkpoint
Data Growth	• Data Generation • Data Acquisition	• Model Parameters representative of the Curated Data	• Model Parameters and GPU Count • Frequency of Checkpoint • Data Retention
Performance	• System Level Tbps/PB (Read/Write) • Device Level Mbps (Read/Write) • IOPS (Read/Write) • Latency	• System Level Tbps/PB (Read) • System Level Latency ms (Read) • Device Level Mbps (Reads)	• System Level Tbps/PB (Read/Write) • System Level Latency ms • Device Level Mbps (Read/Write) • Latency
Workloads	• Raw Data • Structured and Unstructured Data • Text, Videos, Images, etc.	• Vectorized/Structured • Reduced footprint • Large block data sets	• High Read and write throughput workloads scaling with GPUs count

Checkpoints: Driving Capacity and Thruput on Storage

A simple monolithic AI Training HW System

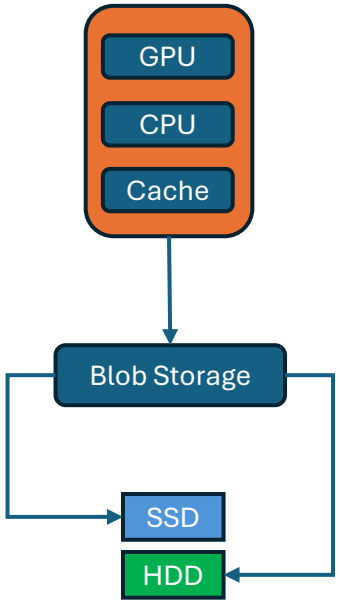

Scaling of GPUs as a function of Training Demand and Fetching of Checkpoints

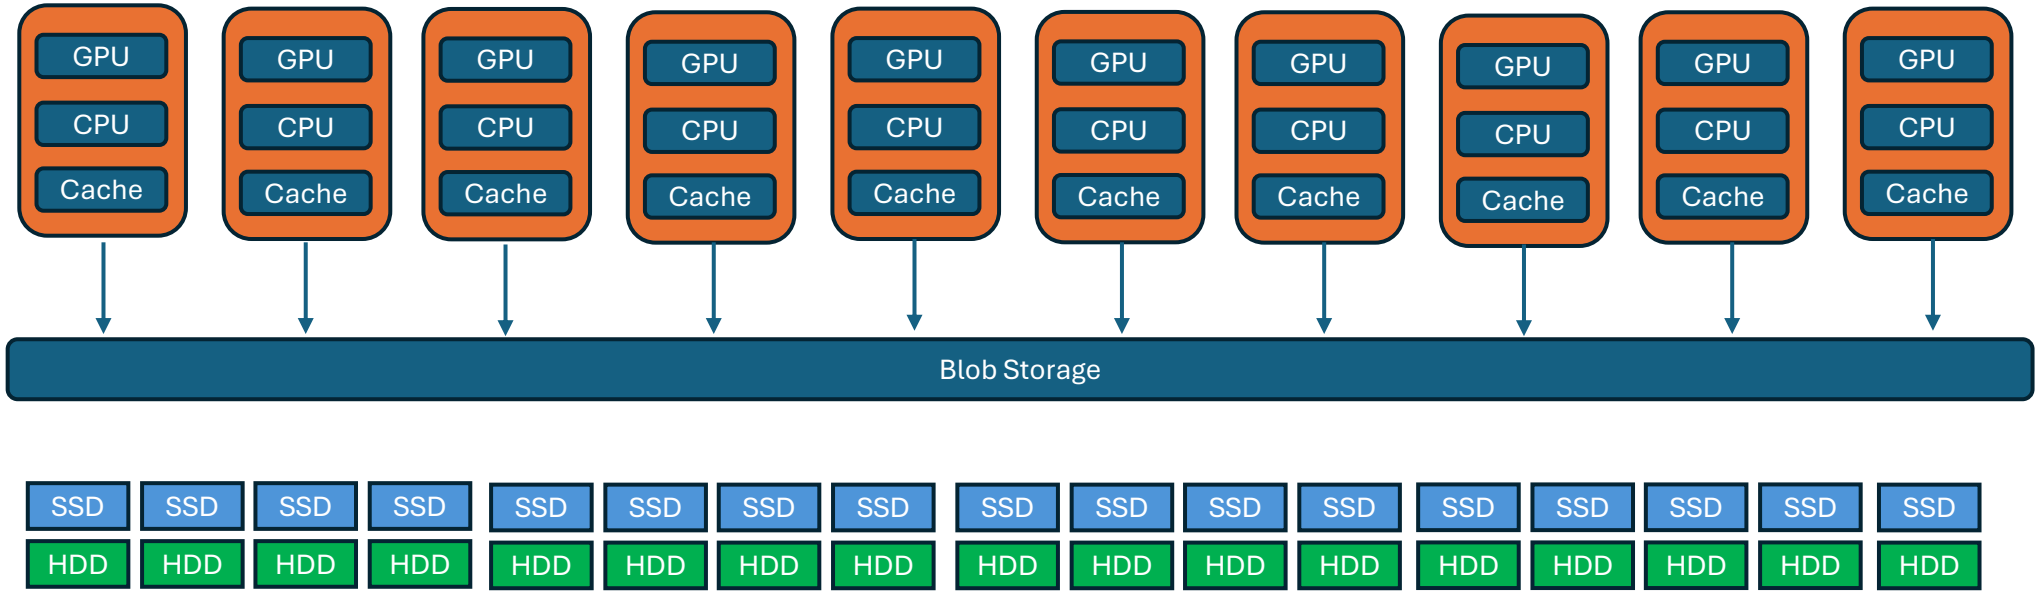

- Checkpoint Frequency: Model and Customer Unique. More Frequent Checkpoints lead to more Storage
- Most Recent Checkpoint Data: Most Recent copy available in SSD tier for low latency access
- Older Checkpoint Data: On HDD Blob Storage Tier, available but slow to access when needed
- GPU Scaling: All at once access from Blob Storage Tier drives high thruput

AI Workload Implications: MB/s/TB

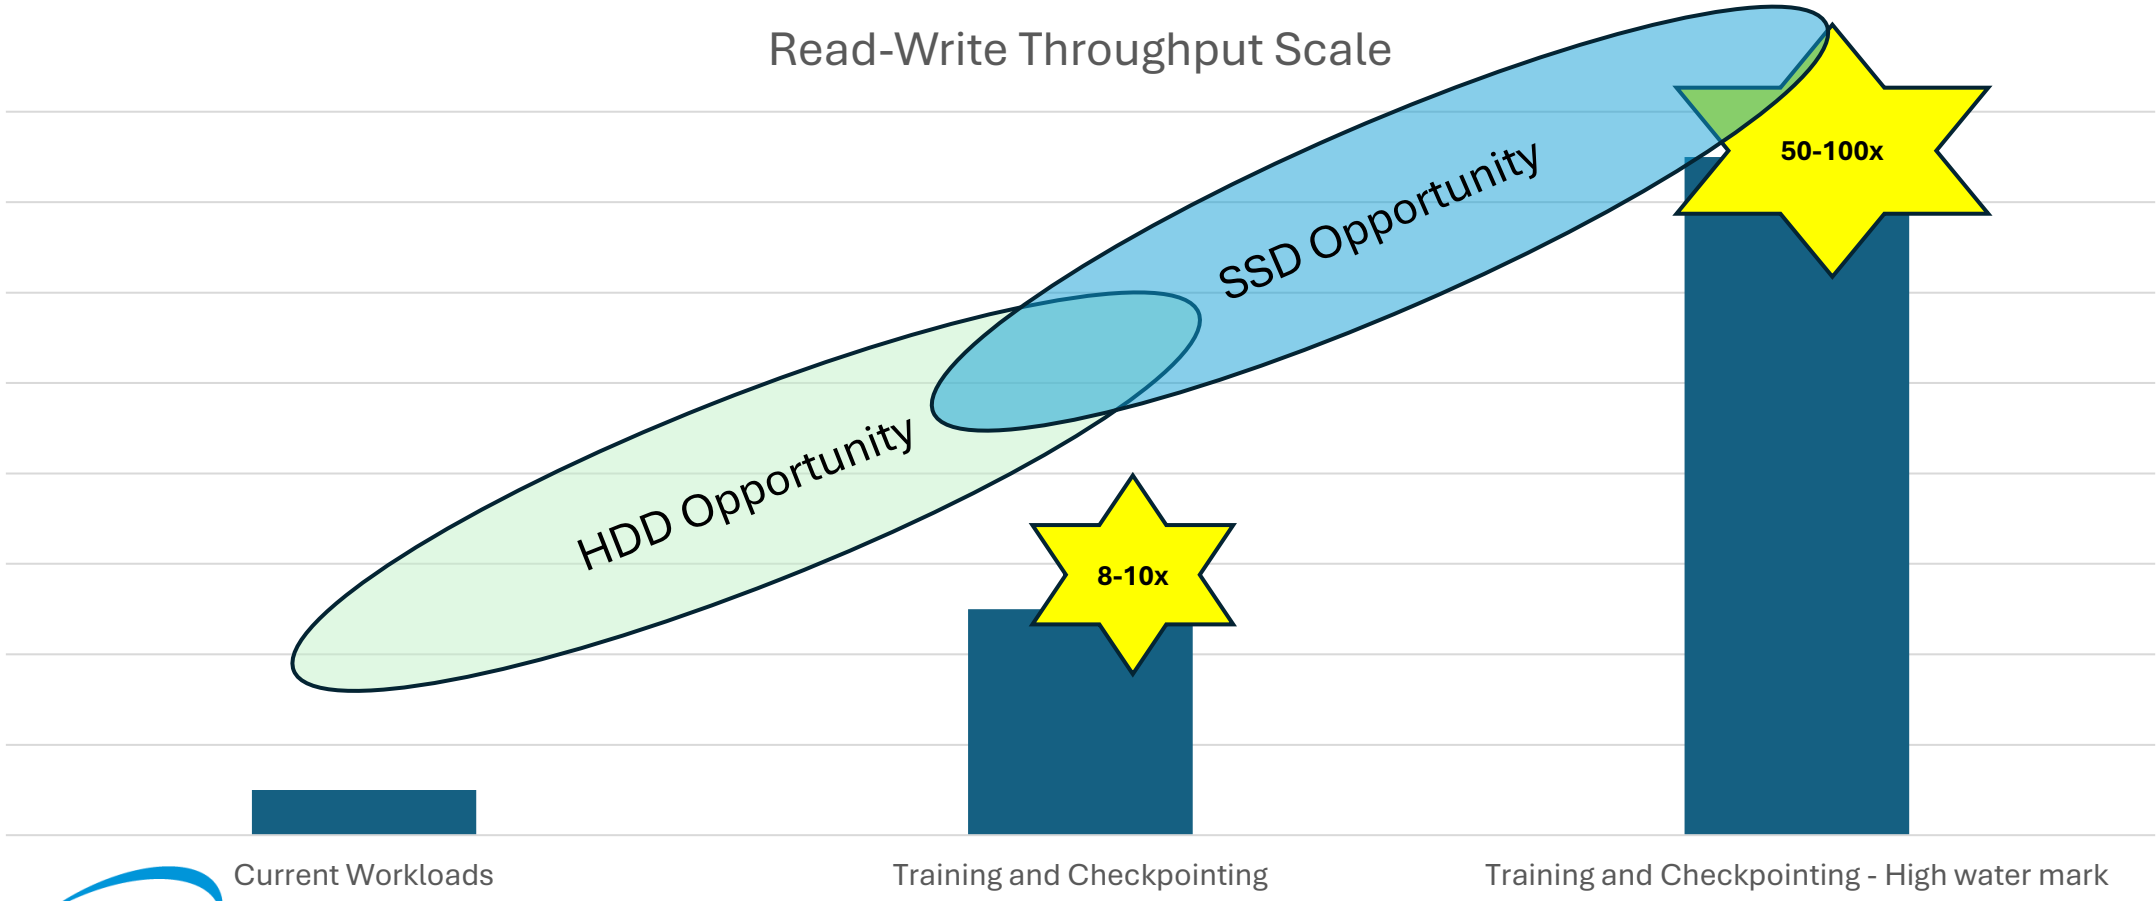

Media Technologies - Call to Action

HDDs

- Power Efficiency Management
- Scale Throughput Density with HDD Capacity

SSDs/Flash

- Density Scaling
- Drive TCO Improvement while delivering on the throughput-Density-Power
- Step Function Improvement in Bit Cost

Questions?