

The logo for SDC | StorageAI. It features a stylized icon of three stacked horizontal bars on the left, followed by the text "SDC | StorageAI" in a bold, white, sans-serif font. The background of the entire slide is a dark blue field with a complex network of glowing lines and dots in shades of blue, green, and orange, suggesting a data network or AI architecture.

SDC | StorageAI™

A SNIA  Event

April 29, 2026 • Denver, Colorado

From Components to Systems: Chair's Welcome

J Metz, Ph.D

AI is Data Movement Bound

- No longer compute-bound

AI is a Single Pipeline

- A pipeline that has seams:
 - Between memory and storage
 - Between storage and network
 - Between network and compute

Workloads Don't Respect These Boundaries

- Burst-Driven
- Latency-sensitive
- *Data-orchestrated!*



The AI Architecture Paradox

- Optimize every layer;
- System still underperforms



One Path: Three Responsibilities

- Placement
 - Where data lives
- Movement
 - How data gets there
- Access
 - How compute consumes data

Every workload depends on **all three**.

This is not a component problem:
It's a coordination problem

Mapping The Framework

- SNIA
 - Placement
- NVM Express
 - Access
- Ultra Ethernet
 - Movement

What Collaboration Means

Not messaging - *mechanics*

1. Semantic alignment
2. Control Loop Integration
3. Data Locality Awareness

Why SDC Matters

- A convergence point:
 - StorageAI
 - NVMe Technical Work
 - Ultra Ethernet development

The Future

- Not faster parts, but a coordinated system
- The future won't be defined:
 - By the best GPU
 - The fastest SSD
 - The most advanced network

But how they **work together**

Morning Agenda

7:30 AM - 9:00 AM	Registration Open - Coffee & Continental Breakfast
9:00 AM - 9:10 AM	Welcome and Opening Remarks – J Metz (SNIA - Chair, SNIA Board of Directors; AMD - Technical Director for Systems Design)
9:10 AM - 9:50 AM	Storage for AI 103, An Introduction to Storage for Inference – Curtis Ballard (SNIA - Co-Chair, SNIA Technical Council; HPE - Strategist, Technology Enablement)
9:50 AM - 10:30 AM	AI Server Clusters - Networking and Storage – David Black (Dell Technologies - Sr. Distinguished Engineer, Infrastructure Systems Group)
10:30 AM - 10:50 AM	Morning Break
10:50 AM - 11:20 AM	StorageAI: A Rising Tide Floats all Boats – Jason Duquette (SNIA - Co-Chair, SNIA StorageAI Community; Dell - Distinguished Engineer, Chief Platform Architect)
11:20 AM - 12:00 PM	Demystifying Data Flows through Typical LLM Training – Bill Lynn (AMD - Fellow, Systems Design Engineering)
12:00 PM - 1:00 PM	Lunch

Thank you to our sponsor:



Afternoon Agenda

1:00 PM - 1:30 PM	An Update on Accelerated Object Storage for AI/ML – Jason Goldschmidt (SNIA - Chair, Accelerated Object I/O TWG; Dell Technologies - Distinguished Engineer)	Addressing QoS issues through IO Prioritization in NVMe – Anthony Constantine (NVMe - Board Member; Micron Technologies - Distinguished Member of Technical Staff)
1:35 PM - 2:05 PM	From GPUs to Flash: An AI-Native Storage Platform for Scalable LLM Inference – Eshcar Hillel (Asteralabs - Principal Research Scientist)	NVMe LBA Access Control for GPU-Direct Storage in AI/HPC Workloads – Chaitanya Kulkarni (NVIDIA - Senior Director)
2:10 PM - 2:40 PM	Scaling Inference with KV Cache Storage Offload and RDMA Accelerated Architecture – Ugur Kaynar (Dell - Distinguished Engineer, Storage Technologist - Storage CTO)	AiSIO: Orchestrating Storage I/O Across CPUs and Accelerators – Simon Lund (Samsung - Principal Engineer)
2:55 PM - 3:25 PM	Scaling Beyond Memory: Fine-grained GPU Access to Unbounded Data for Vector Databases and Graph Neural Networks – CJ Newburn (NVIDIA - Distinguished Engineer)	
3:30 PM - 4:00 PM	Enabling AI Storage Benchmark Evolution at the Pace of AI – Wes Vaske (Micron Technology - Senior Member of Technical Staff)	PNFS Past, Present, and Future: What's All the Excitement About? – Gary Grider (LANL - Leader, High Performance Computing (HPC) Division)
4:05 PM - 4:35 PM	Data Storage Innovations for Scalable AI Infrastructure – Erich Haratsch (Marvell - Senior Director)	DiskANN - VectorDB Indexing Offload – Alessandro Goncalves (Solidigm - Solutions Architect)
4:40 PM - 5:10 PM	AI Storage Workloads: What is Different and Impacts on SSD Design – Rory Bolt (KIOXIA America - Senior Fellow)	Take a Break: A Deep Dive into Checkpointing – John Mazzie (Micron Technology - Member of Technical Staff, Systems Performance Engineer)
5:15 PM - 6:15 PM	Birds of a Feather: Storage and Networking Considerations for AI Deployments	Birds of a Feather: Rearchitecting the Stack for AI: Data Access, Benchmarks, Scale, and SSD Innovation
6:15 PM - 7:15 PM	Birds of a Feather: Building a Standards-Driven Ecosystem for AI Storage	Birds of a Feather: Community Discussion of Storage for AI Workload Definitions, Characterization, and Benchmarking Needs and Wants



Rate the Sessions!

Please remember to rate the sessions on the tablets at the back of each room after the session is over.

SDC'26

SNIA DEVELOPER CONFERENCE



BY Developers FOR Developers

- Taking place in Santa Clara on September 28-30
- Call for Presentations is currently open - Submit by May 15
- Early Bird registration ends on September 5th
- SDC Regional attendees receive an additional \$200 off by using code SDC26AIDis

The logo for SDC | StorageAI. It features a stylized icon of three stacked horizontal bars on the left, followed by the text "SDC | StorageAI" in a white, sans-serif font. The background of the entire slide is a dark blue gradient with abstract, glowing lines and dots in shades of blue, green, and orange, suggesting a digital or data environment.

SDC | StorageAI™

A SNIA  Event

Thank You