SDC | StorageAI™

A SNIA Event

April 29, 2026 • Denver, Colorado

A Rising Tide Floats All Boats

Jason Duquette (Dell Technologies)
StorageAI Chair

Rising Tides: How the StorageAI Community and Open Standards are Accelerating the AI Era

- AI's insatiable appetite for data has created a massive bottleneck. We are building expensive GPU clusters, but GPUs are frequently sitting idle waiting for data. The storage industry has a once-in-a-generation opportunity—and responsibility—to solve this.
- In the tech world, that rising tide is **open standardization**. When competitors collaborate to solve foundational infrastructure problems, everyone scales faster—vendors, customers, and entirely adjacent industries.

Looking Back: How Storage Standards Advanced Connected Industries

A History of Storage Standards & the Connected Industries They Enabled

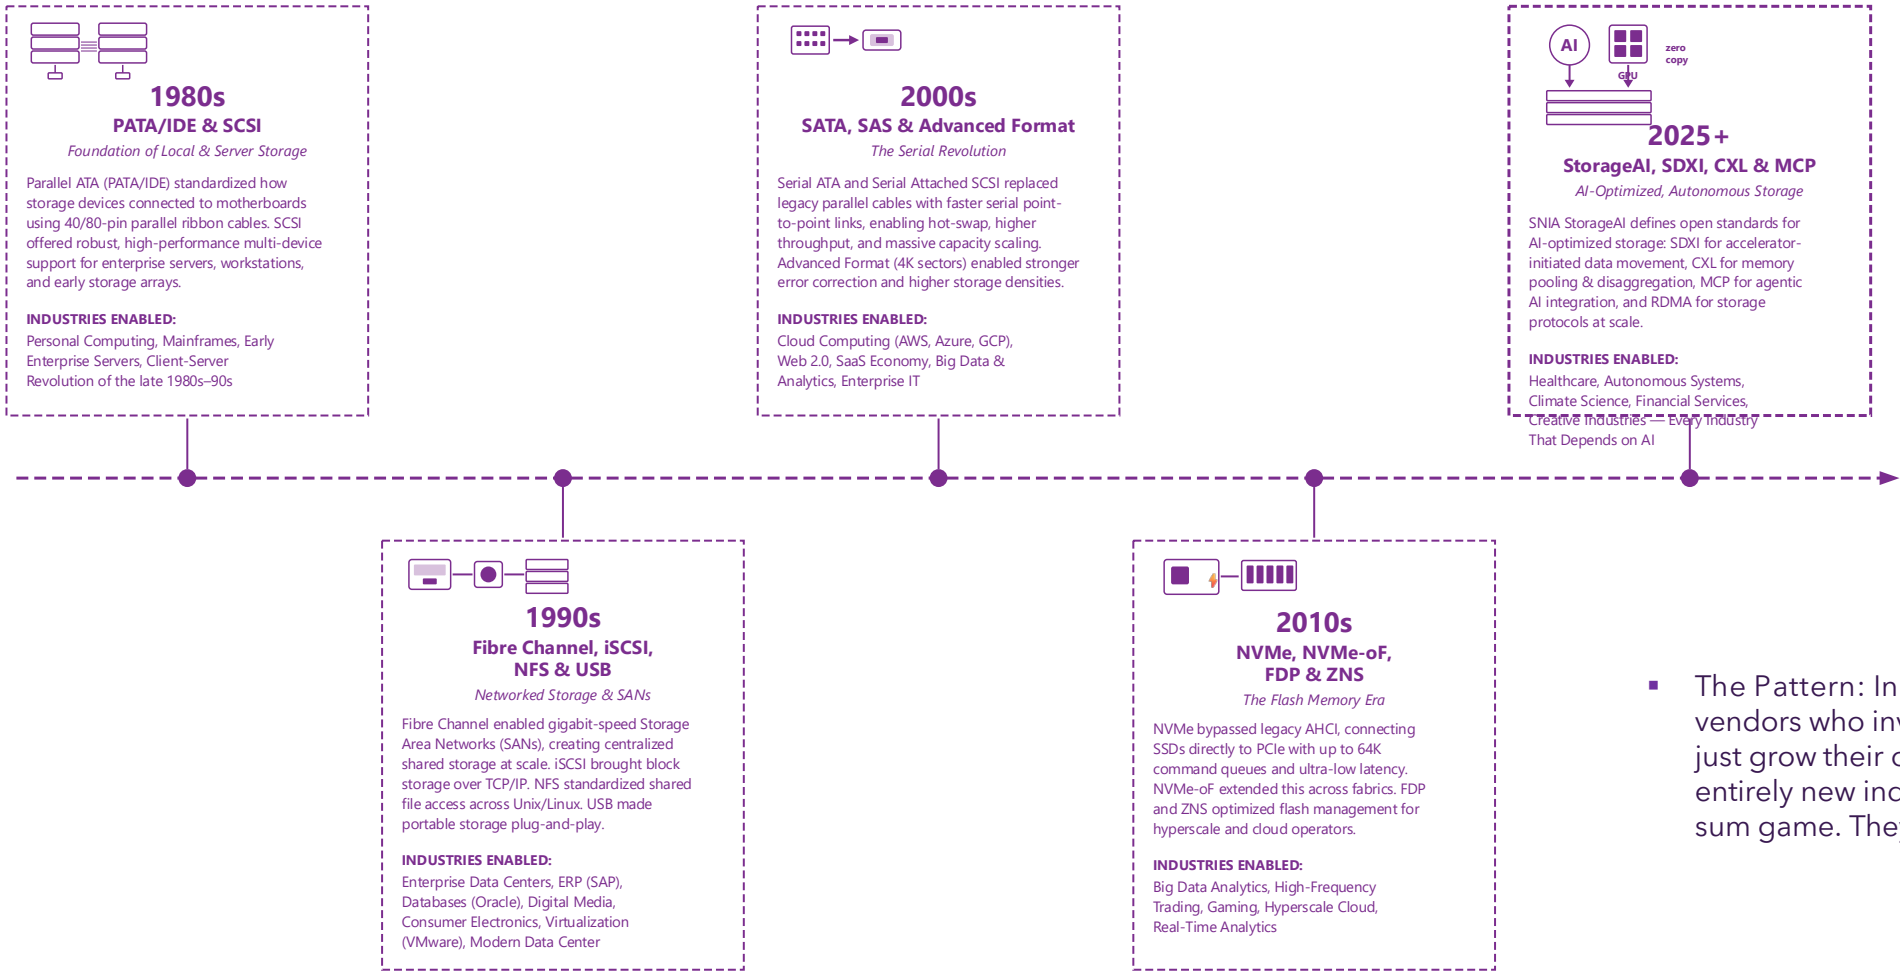

"A rising tide floats all boats — but only if we build the harbor together."

- The Pattern: In every generation, storage vendors who invested in open standards didn't just grow their own market–they unlocked entirely new industries. Standards are not a zero-sum game. They are a multiplier.

SDC | StorageAI™

The Problem: Why AI Demands a New Approach to Storage

- The Current Fragmented Landscape:
  - AI data typically resides on separate storage networks, forcing AI workloads across multiple inefficient network boundaries.
  - Proprietary silos slow down innovation and create costly vendor lock-in.
  - Every AI pipeline stage—ingestion, preprocessing, training, checkpointing, inference—has unique and demanding storage requirements that legacy architectures were never designed to handle.
- The Data Gravity Problem:
  - Models are getting larger, datasets are growing exponentially, and the cost of moving data is becoming the dominant constraint.
  - Without standardization, every organization is reinventing the wheel—building bespoke plumbing instead of advancing AI itself.
- The Opportunity Cost:
  - Every hour a GPU sits idle waiting for data is wasted capital.
  - Every proprietary integration that breaks when a vendor updates firmware is wasted engineering time.
  - The industry cannot afford this fragmentation at the scale AI demands.

Storage Touches Every Stage: Mapping the AI Pipeline's Six Critical Touchpoints

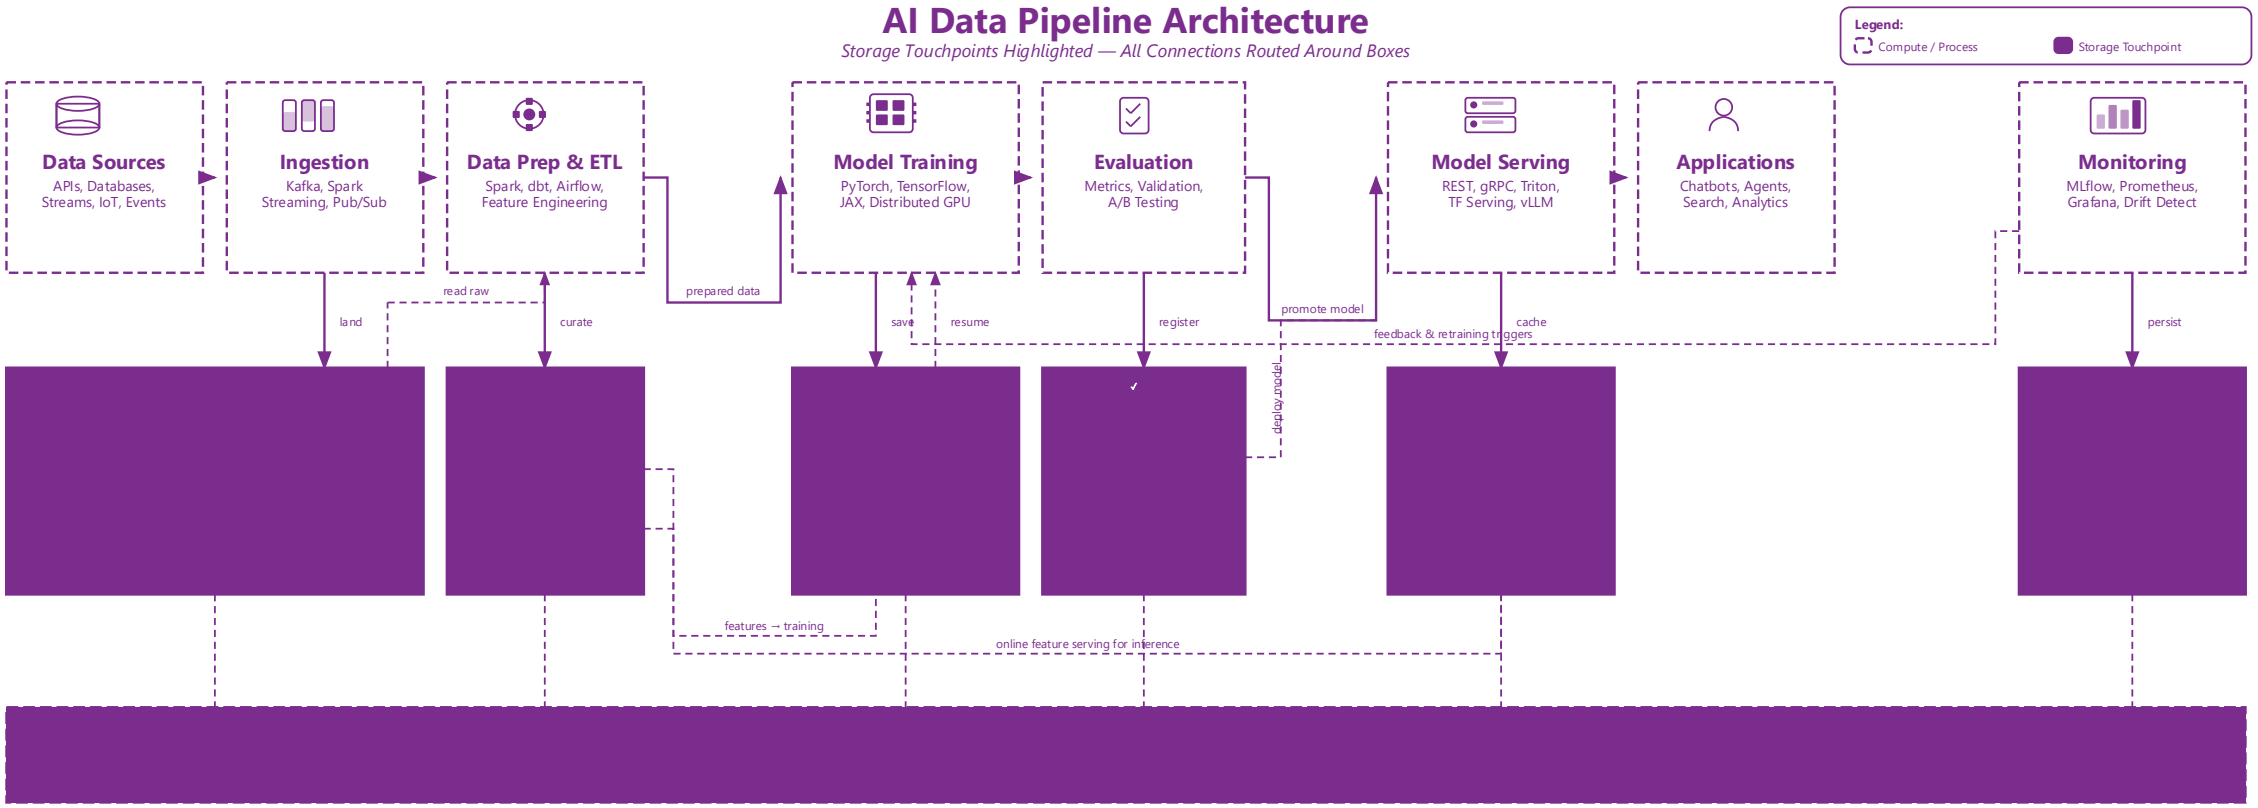

SDC | StorageAI™

What StorageAI is Doing About It

- The Mission: Define open, vendor-neutral standards and educational frameworks that eliminate the storage bottleneck for AI workloads.
- Key Technical Focus Areas:
  - Smart Data Accelerator Interface (SDXI): Enabling standardized, processor-agnostic memory-to-memory data movement that bypasses multiple software abstraction layers.
  - Accelerator-Initiated I/O & GPU Direct Access: Allowing GPUs to pull data directly from storage without relying on the CPU, eliminating a massive bottleneck.
  - RDMA & High-Performance Fabrics: Enabling storage protocols (File, Object, Block, KV) to operate seamlessly over RDMA and emerging fabrics like Ultra Ethernet.
  - StorageAI Infrastructure & Data Discovery: Extending specifications like Swordfish, Redfish, and CDMI so AI agents can autonomously discover, provision, and manage storage resources.
  - Agentic AI Optimization: Driving work on the Model Context Protocol (MCP) and Agent-to-Agent (A2A) communication to transform storage from a passive repository into an active, intelligent participant in AI workflows.
  - Data Resilience, Lifecycle & Sanitization: Addressing new paradigms like ephemeral data, rack-scale failure domains, machine unlearning, and "the right to be forgotten."
  - Power & Efficiency: Measuring and optimizing the energy footprint of storage under AI-specific workloads.
  - Computational Storage & Memory Technologies: Pushing compute closer to data via computational storage, CXL-attached memory, and persistent memory architectures.
- The Mutual Benefit: By creating standardized hardware and software interface definitions, vendors can build interoperable components, buyers aren't locked into single-vendor ecosystems, and startups can innovate on top of a stable foundation rather than fighting infrastructure battles.

Education: The Catalyst for Ecosystem Growth

- Standards Need Evangelists: Creating a technical specification is only half the battle. If nobody understands how to implement it, it sits on a shelf. Teaching the industry how to use standards is the other critical half.
- The Educational Mandate of Storage.AI:
  - SNIA leadership has emphasized that developing specs must be paired with educating developers, architects, and organizations.
  - The Storage.AI group is actively creating white papers, primers, and webcasts to demystify complex topics.
  - Educational content is being developed across key areas including KV Cache architectures, RDMA networking protocols, data resilience for AI, and data sanitization/machine unlearning.
- Demystifying the AI Pipeline: AI workloads are often viewed as an "opaque box" by storage professionals. The community is working to educate architects on the unique storage touchpoints at every stage:
  - Data Ingestion & Pre-processing – Where does the data live, and how do you get it to the pipeline efficiently?
  - Training (Forward/Backward Passes) – Data sharding, distributed training, and the I/O patterns that make or break GPU utilization.
  - Checkpointing & Archiving – High-stress, bursty writes that can overwhelm traditional storage. A single checkpoint for a large model can be terabytes.
  - Inference & Serving – Low-latency, high-throughput access patterns that differ fundamentally from training.
- Lowering the Barrier to Entry:
  - Education reduces the barrier to entry for startups and smaller enterprises, ensuring standards are actually adopted and implemented in the wild.
  - Conferences like the Storage Developer Conference (SDC) serve as critical venues for hands-on education and community building.
  - The group is also exploring guest speakers, cross-pollination with partner organizations (NVMExpress, UEC, OCP, DMTF), and leveraging industry events like face-to-face meetings in Denver to accelerate knowledge sharing.
- The Multiplier Effect of Education:
  - A well-educated ecosystem adopts standards faster.
  - Faster adoption means larger markets.
  - Larger markets mean more investment, more innovation, and more boats lifted by the rising tide.

The Rising Tide in Action: What This Means for the Future

- For Storage Vendors: Faster time-to-market, reduced engineering costs from proprietary integrations, expanded addressable markets, and the ability to compete on value rather than lock-in.
- For AI Developers & Startups: Lower infrastructure costs, fewer integration headaches, the freedom to choose best-of-breed components, and the ability to focus on model innovation rather than storage plumbing.
- For Data Center Architects: Simplified procurement, interoperable multi-vendor environments, consistent management interfaces, and confidence that today's investments won't become tomorrow's stranded assets.
- The Ripple Effect Across Industries: By solving the AI storage bottleneck, we accelerate the industries that rely on AI:
  - Healthcare: Faster drug discovery, genomics analysis, and medical imaging.
  - Autonomous Systems: Self-driving vehicles, robotics, and smart logistics.
  - Climate Science: Large-scale climate modeling and environmental monitoring.
  - Financial Services: Real-time fraud detection, risk modeling, and algorithmic trading.
  - Creative Industries: Generative content creation, media processing, and design.
  - The pattern is clear: every industry that depends on AI depends on the storage infrastructure beneath it. When we raise the tide for storage, we raise it for all of them.

Conclusion & Call to Action

- The Lesson of History: Every major leap in computing—from mainframes to client-server, from the web to cloud, from mobile to AI—was built on a foundation of open standards. The storage industry has been a key architect of that foundation in every generation. We have the opportunity—and the responsibility—to do it again.
- The Power of Community: Storage.AI represents something rare: an industry coming together not because a regulatory body forced it, but because the participants recognize that the challenge of AI infrastructure is bigger than any single company. That voluntary collaboration is the true rising tide.
- The Ask:
  - Participate: Join the Storage.AI community. Attend meetings. Contribute your expertise.
  - Volunteer: Lead or join sub-groups on the key educational topics—KV Cache, RDMA & Networking, Data Resilience, Data Sanitization, and more.
  - Educate: Help create the white papers, webcasts, and conference content that will bring these standards to life.
  - Evangelize: Carry these ideas back to your organizations. Advocate for open standards in your procurement decisions, your architecture reviews, and your engineering roadmaps.
- The Closing Thought: "A rising tide floats all boats—but only if we build the harbor together." The specifications, the education, the community—that's the harbor. Let's build it.


SDC | StorageAI™
A SNIA Event
Thank You