SNIA DEVELOPER CONFERENCE

SDC 25

By Developers FOR Developers

Hyatt Regency Santa Clara, CA
September 15-17, 2025

The Future of Data

Fueling the AI Revolution with SNIA Storage.AI

J Metz, Ph.D
Chair, SNIA Board of Directors

www.sniadeveloper.org

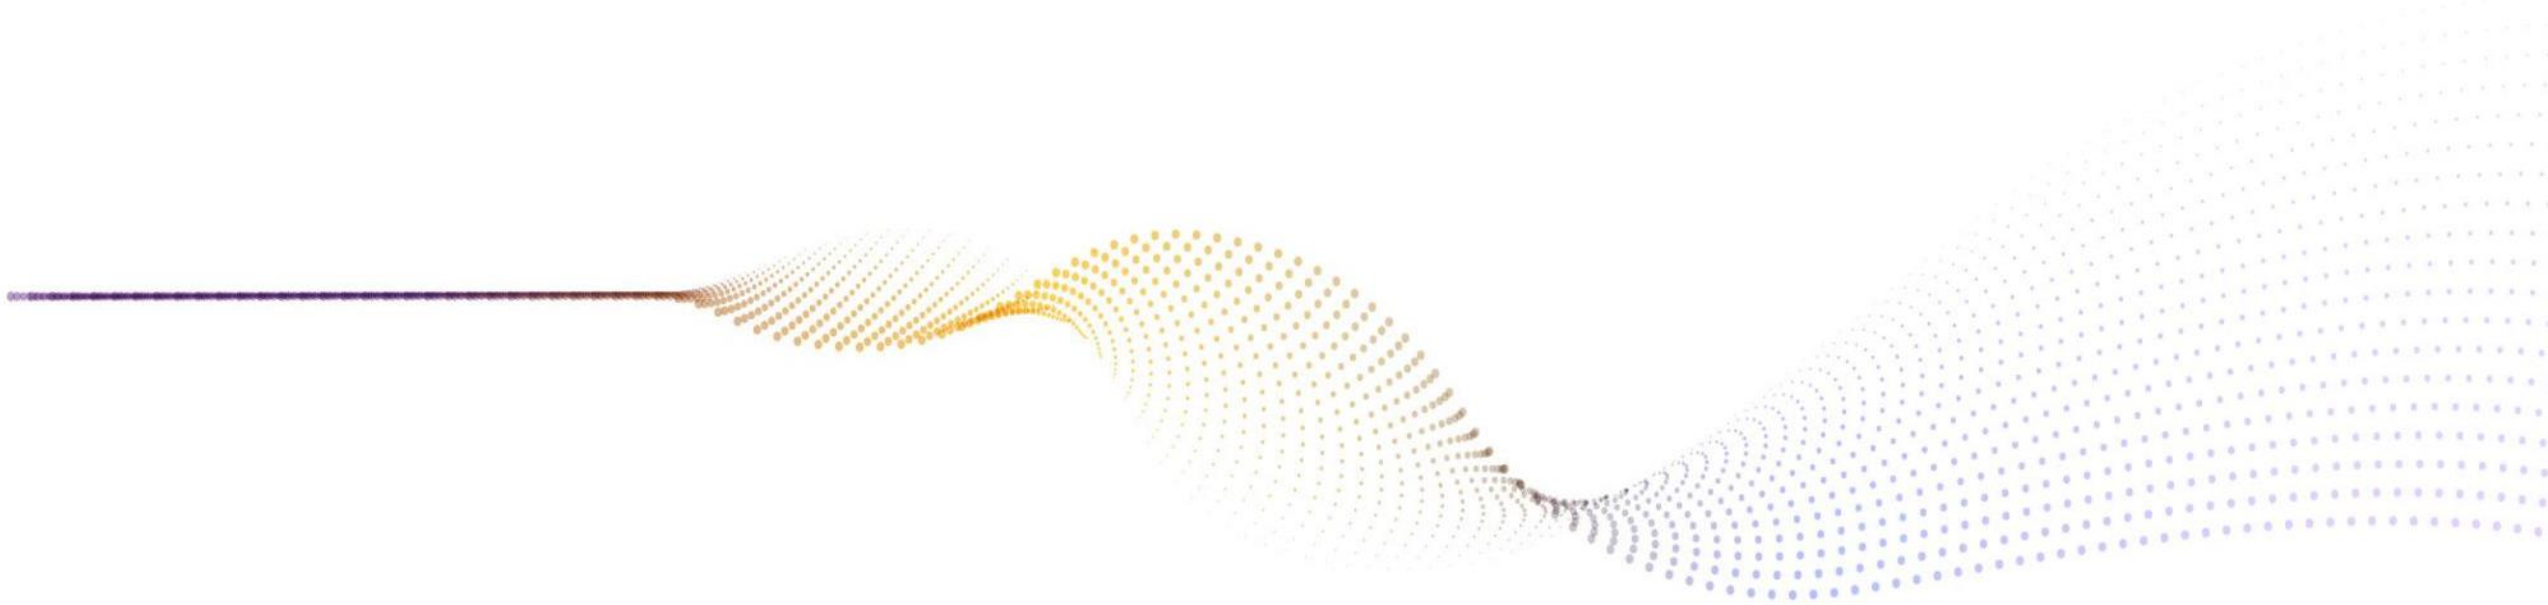

The Problem…

The AI Paradox

- The more data you need…
- …the more data you need in order to process the data!

How People Think AI "Works"

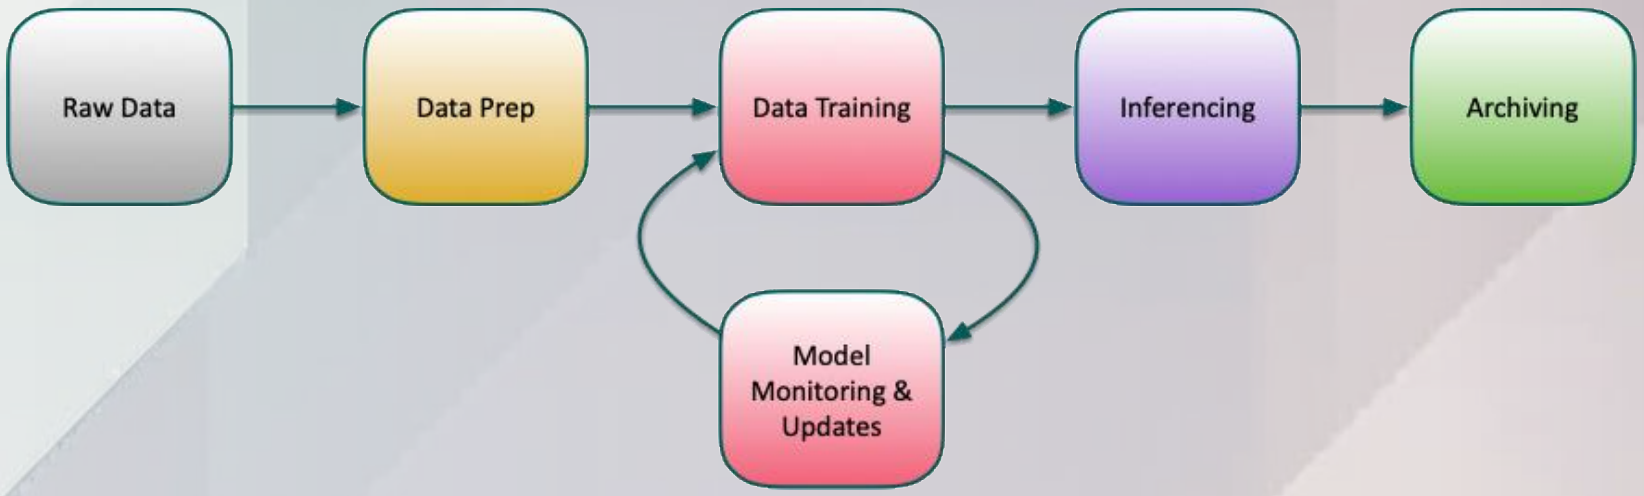

How People Think AI "Works"

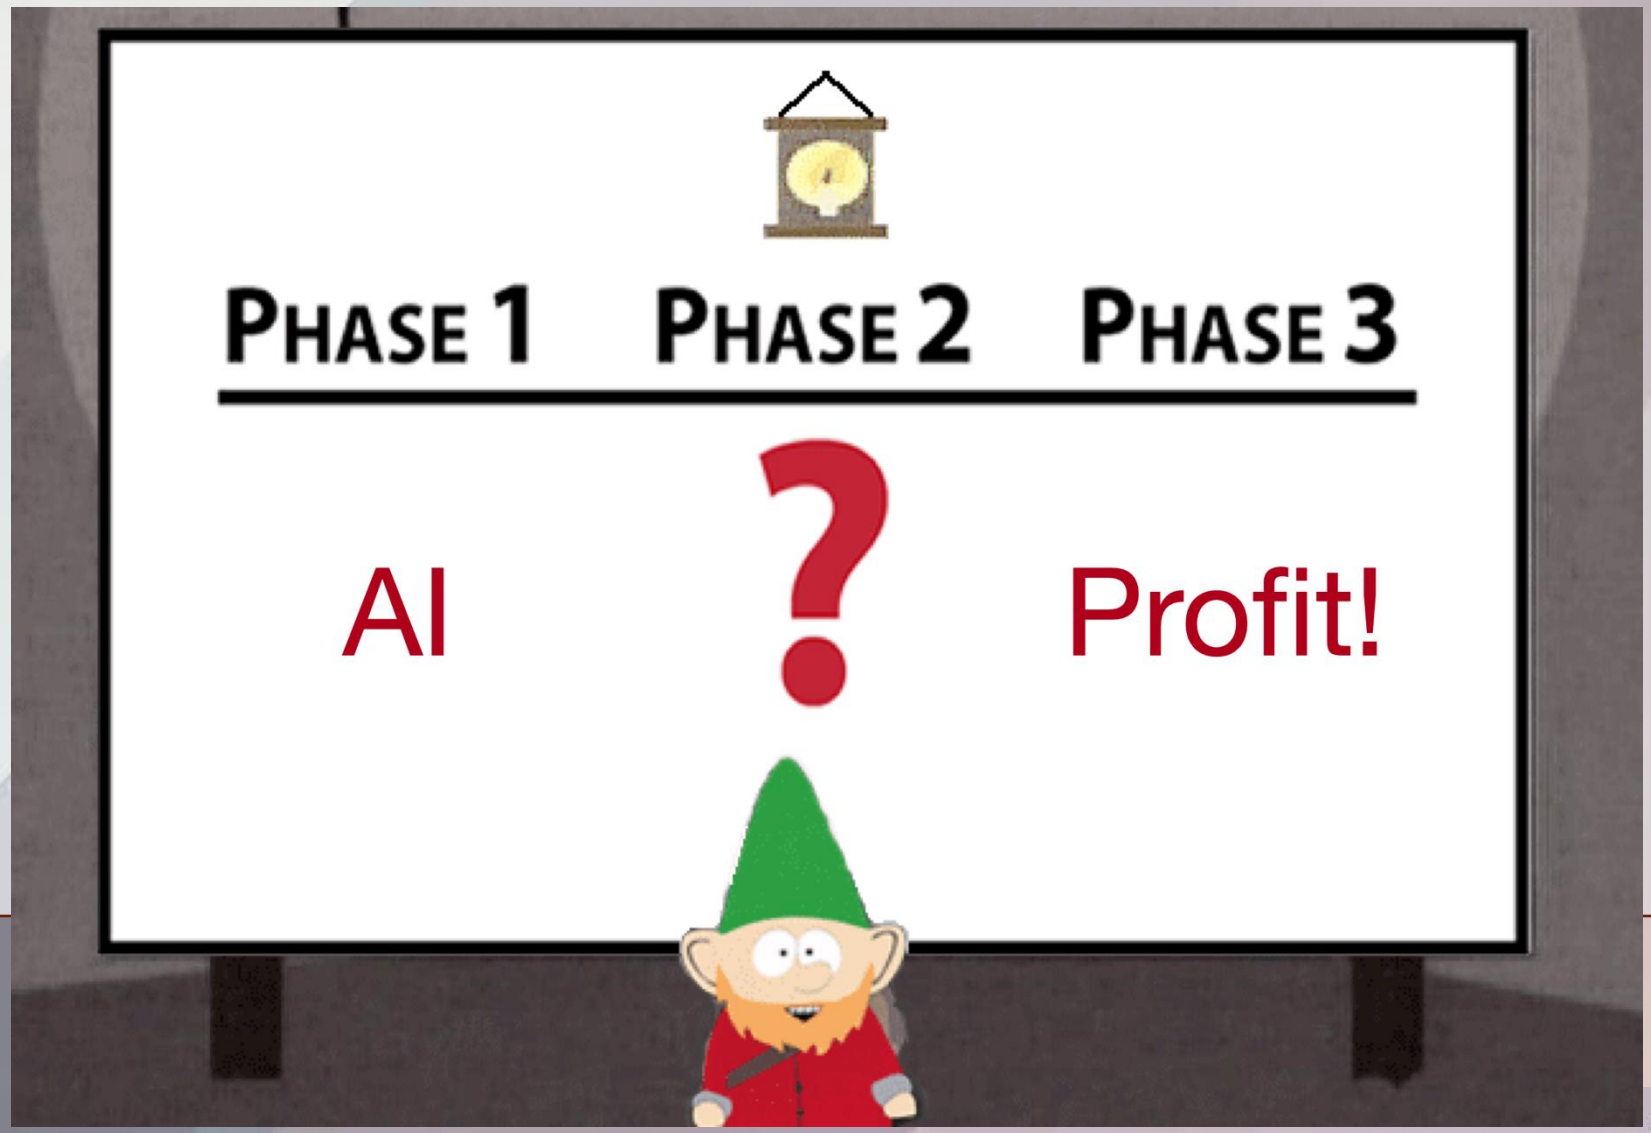

How It Actually Works*

- The “AI Workload” is actually multiple steps of multiple workloads
- Different processors, data types, and network/transports have specific roles
  - Logarithmic expansion of the relationships stresses networks
  - Processing, Memory, and Bandwidth constraints

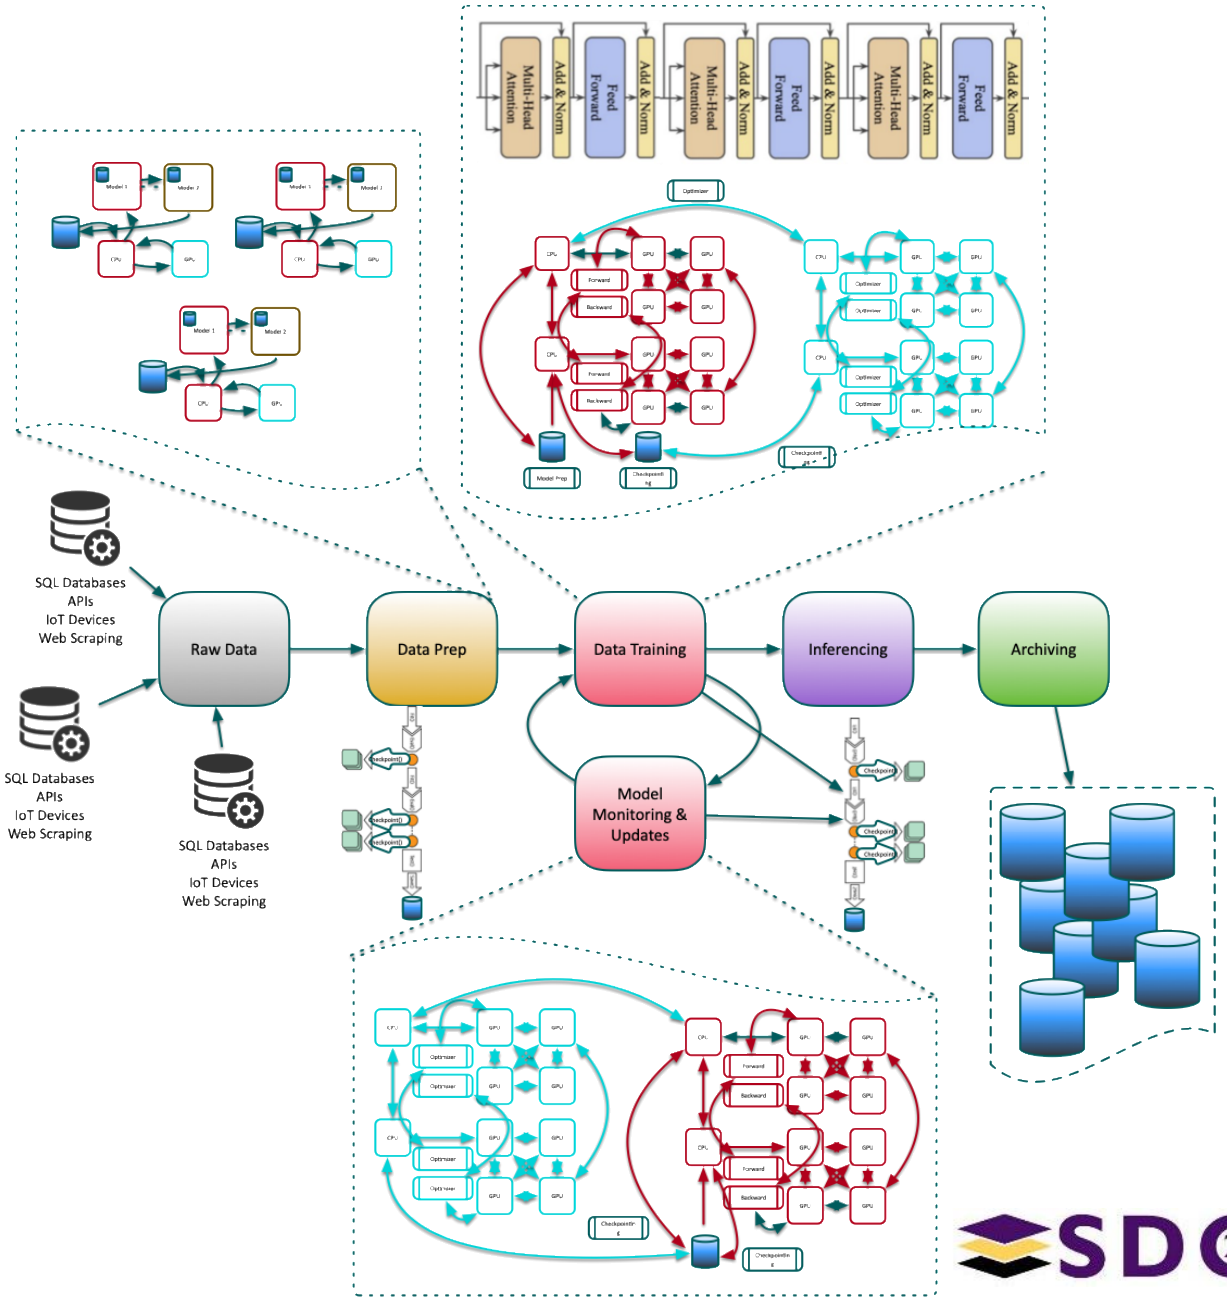

*not to scale

In Technical Terms…

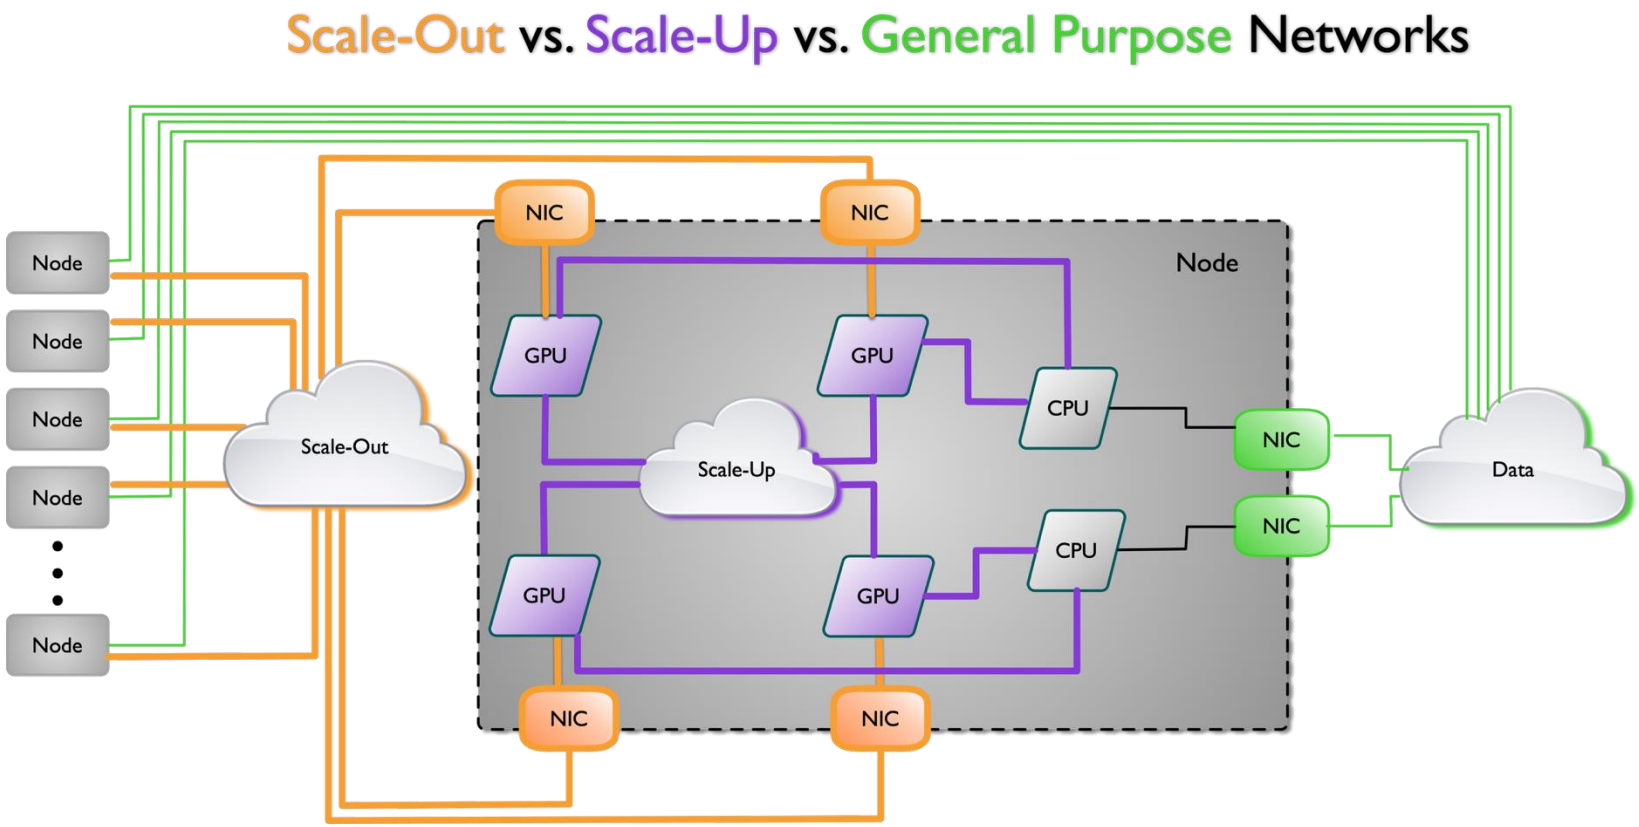

- The data is not where you need it, when you need it
- Data networks are separate from accelerators
- Getting data *into* the networks is non-trivial!

It's Not a "Straight Shot"

- Accelerators do not directly initiate or access data
- Large systems require additional large systems to transfer data *in the correct form* to accelerators

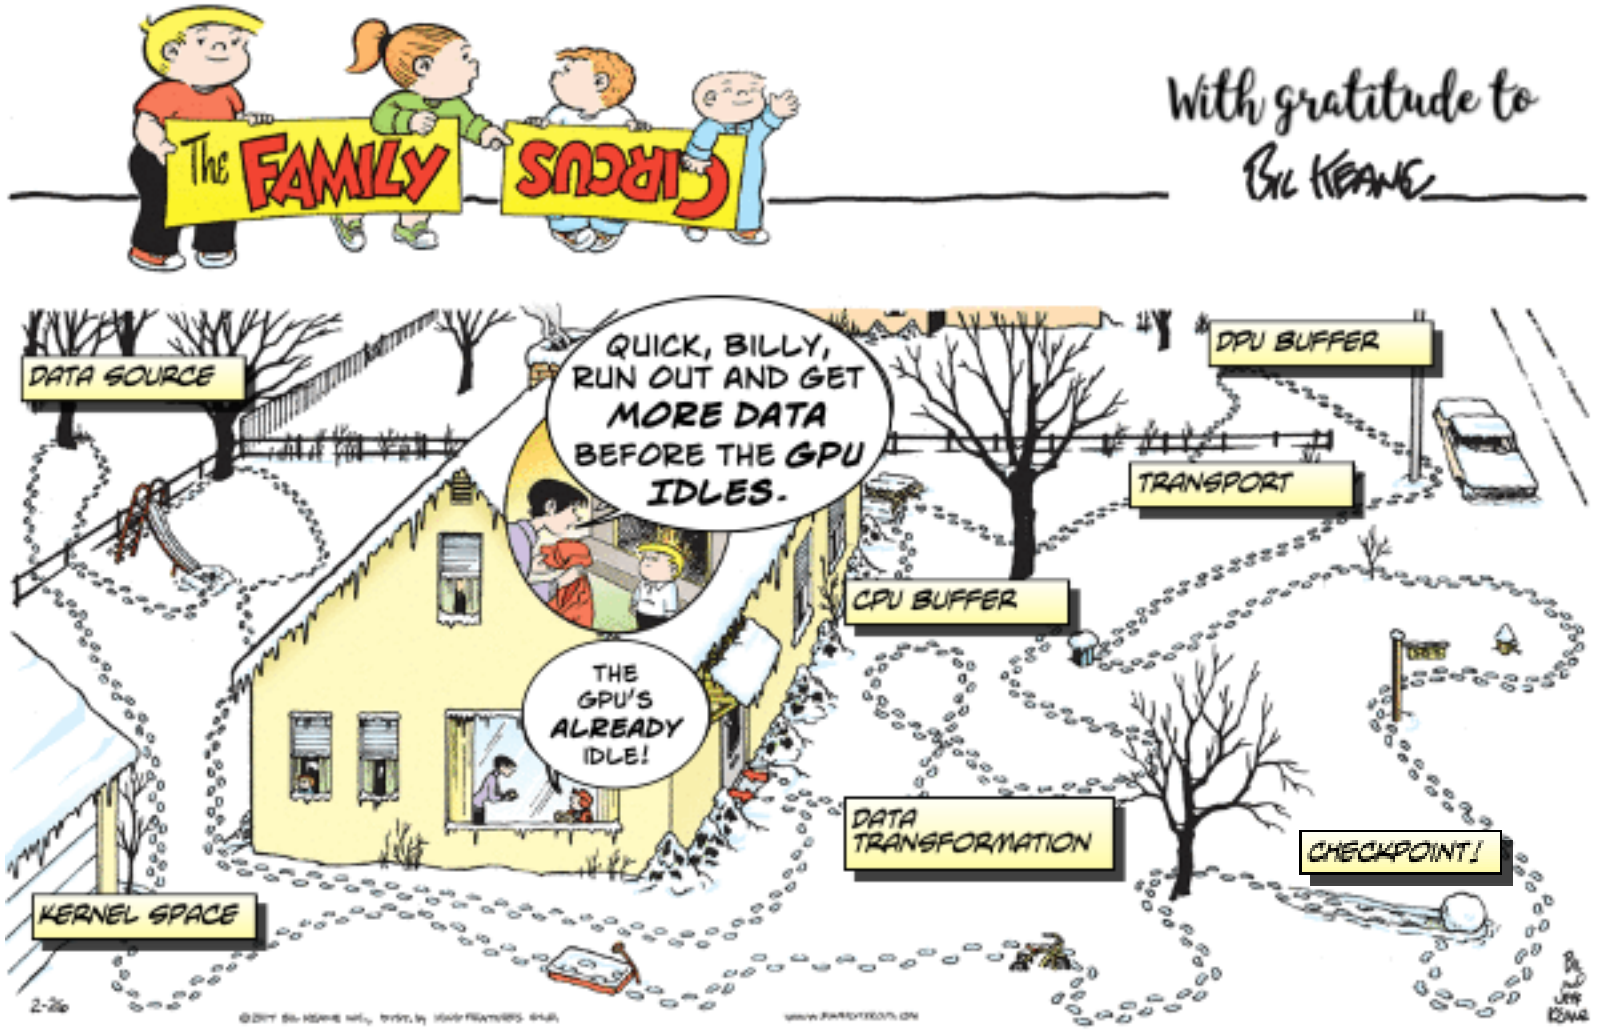

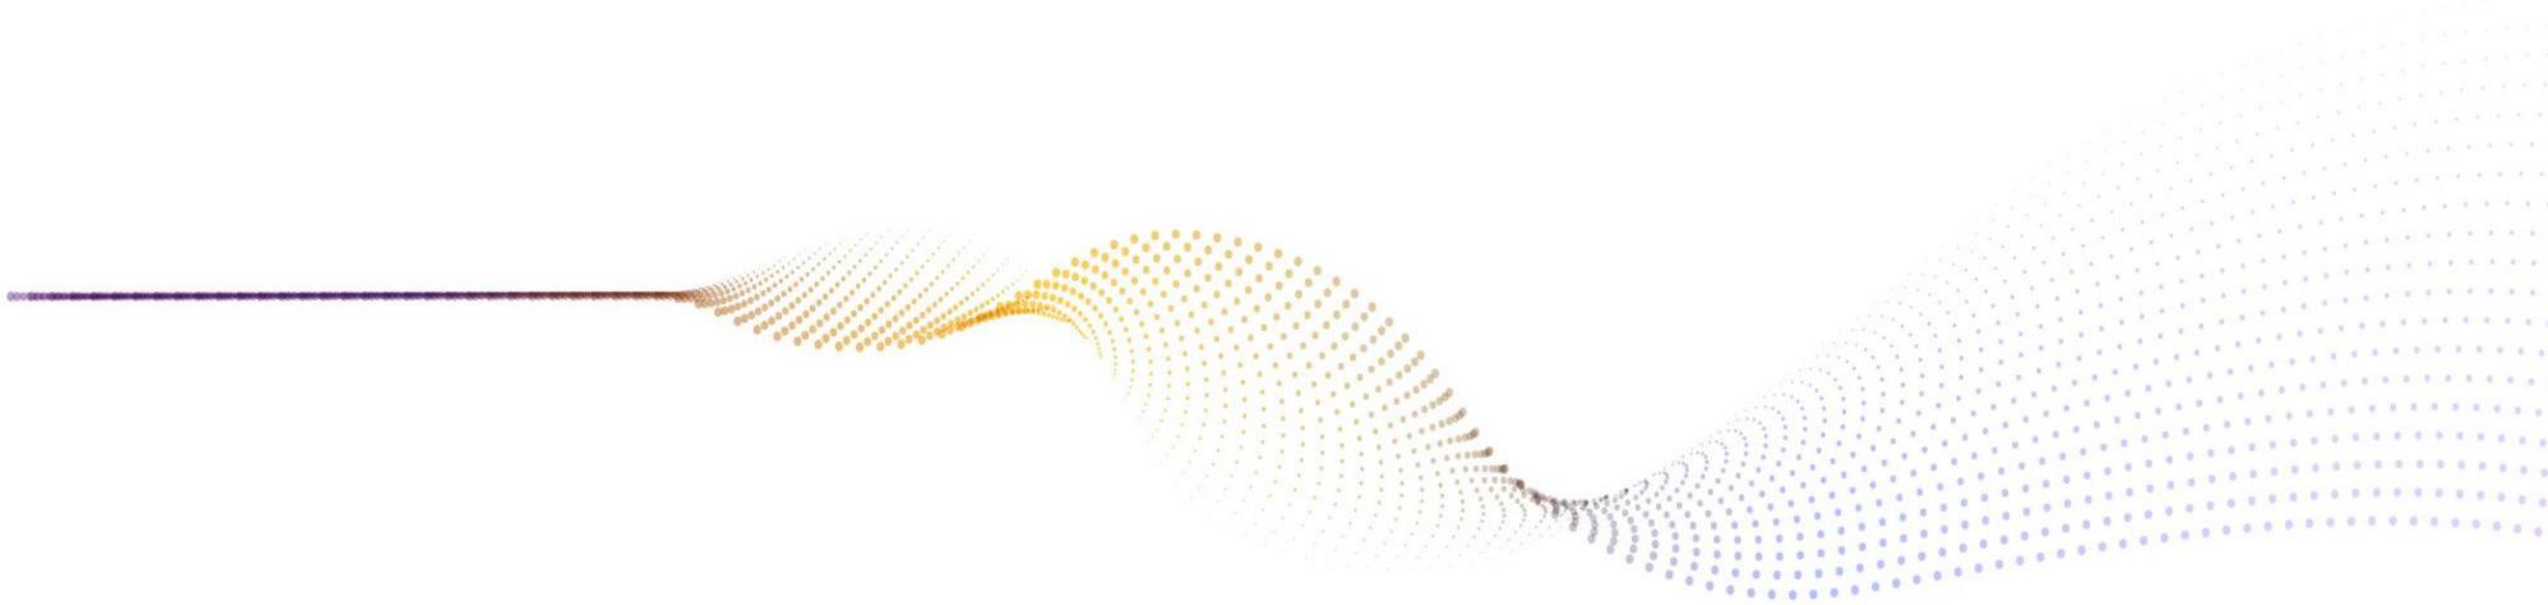

SNIA Storage.AI

snia.ai

What is Storage.AI?

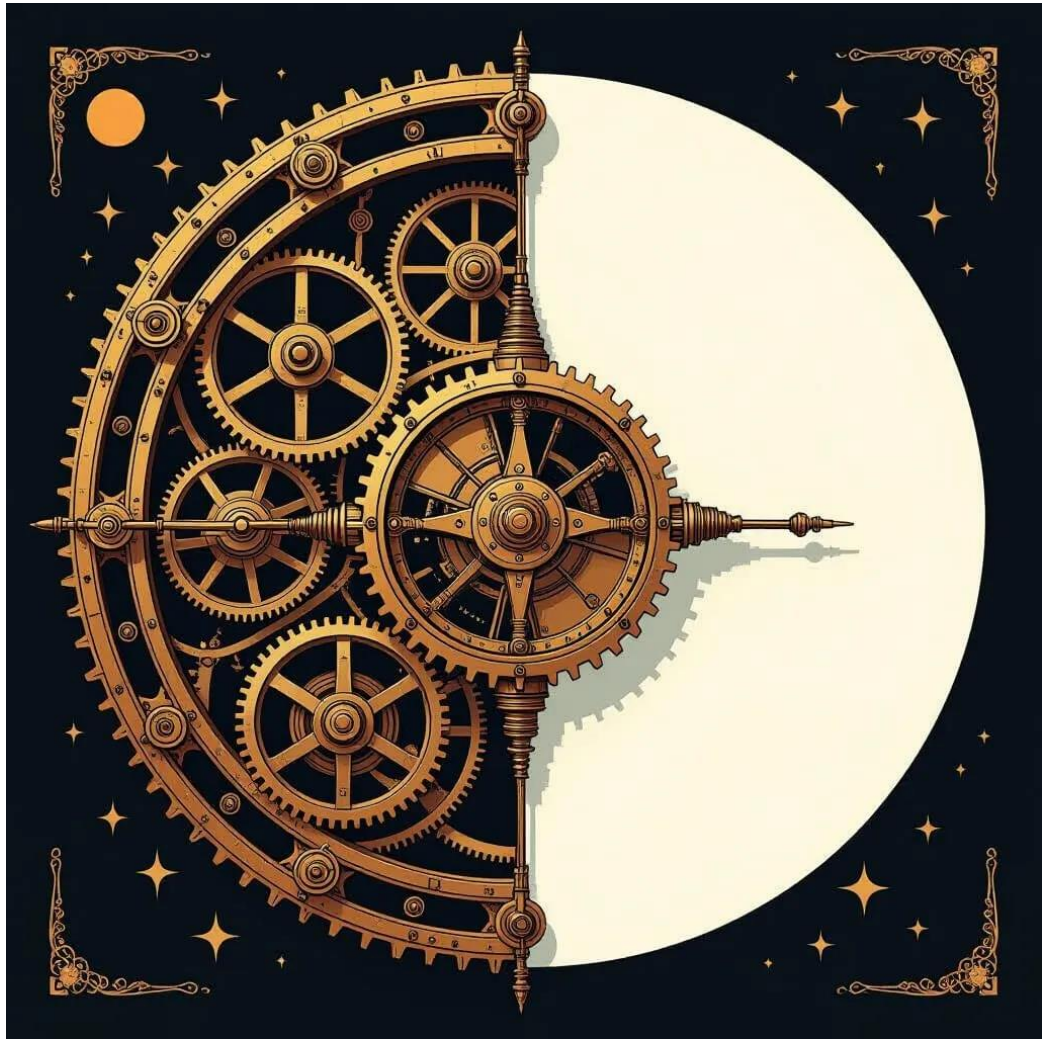

- A vendor-neutral, open-standards effort tackling the most pressing AI data service challenges—such as memory tiering, latency, data movement, compute-near-storage, and storage efficiency
- Or…
- Making data work better for AI workloads

The Goals of Storage.AI

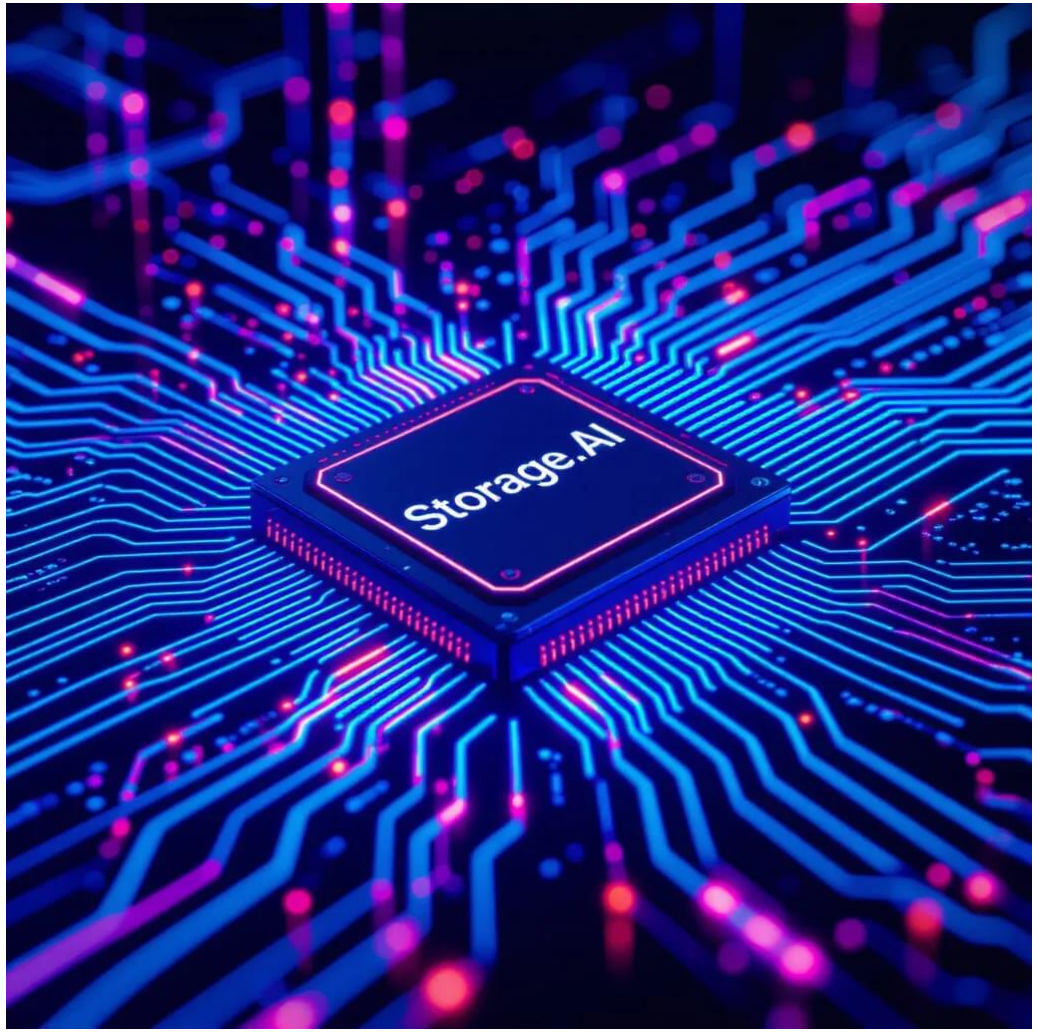

- Reduce I/O data amplification, optimize I/O consumption for accelerators
- Efficient, secure, and reliable data movement through the AI workload lifecycle
- New initiation and consumption model for accelerators
- Standardized hardware/software interface definitions
- Open programming models
- Secure data movement and capacity

The Approach, The Vision, The Principles

- Champion **open** standards for efficient data services
- Flexible integration of data services standards
  - Each technology improves performance, efficiency, and effectiveness; works even better together
- Integration with industry standards organizations

Potential Applications of Storage.AI: Existing SNIA Projects

SNIA Standard / Initiative	Enables…
SDXI (Smart Data Acceleration Interface)	Processor-agnostic DMA acceleration and data transformation
Computational Storage API & Architecture	Near-data compute
NVM Programming Model	Unified software interface for accessing memory tiers
Swordfish / Redfish Extensions	Storage management
Object Drive Workgroup	Standard interfaces for object storage
Flexible Data Placement APIs	Optimize data layout and streaming throughput
Security	3 guesses…
Green (Energy)	Environmental regulations, power and thermal solutions

Potential Applications of Storage.AI: New SNIA Projects

- Members have expressed interest in creating the following technical workgroup projects:

SNIA Standard / Initiative	Enables…
File/Object over RDMA/UE	Hybrid file/object storage backends combined with remote memory access
Accelerator Direct CPU Bypass	Bypasses CPU for data movement between accelerators and data
Accelerator-Initiated Storage I/O	Reduction of CPU bottlenecks

- More to come!

Right Tool For The Job

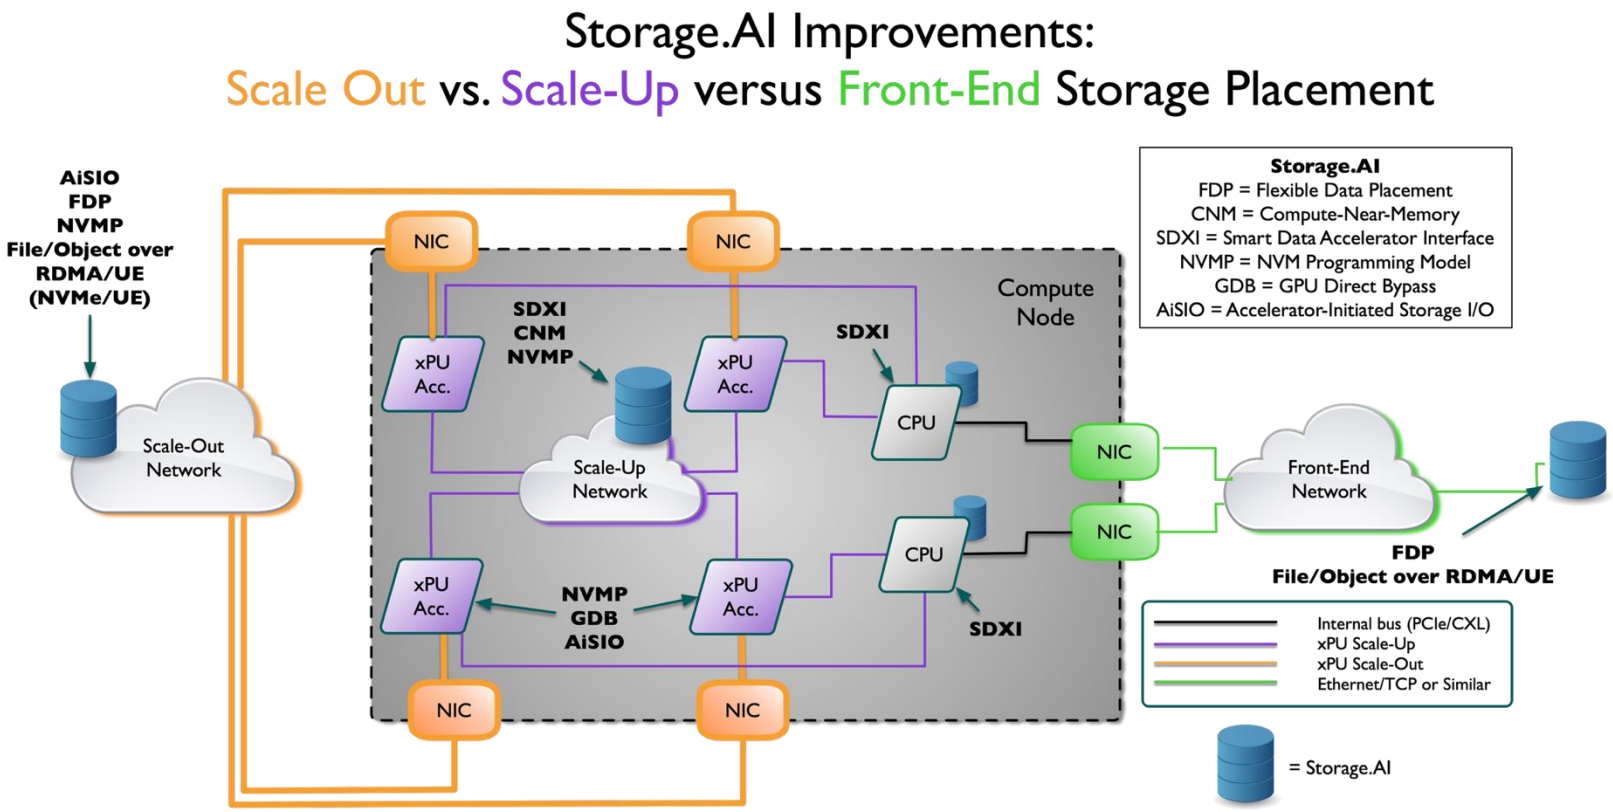

- SNIA technical work addresses pieces of the puzzle
- Storage.AI looks at the big picture

Communities and TWGs: The Relationship

- Communities: What work needs to happen
  - What are the problems that need to be solved?
  - How does it all work together?
- TWGs: Where the work gets done
  - Let's make it work!

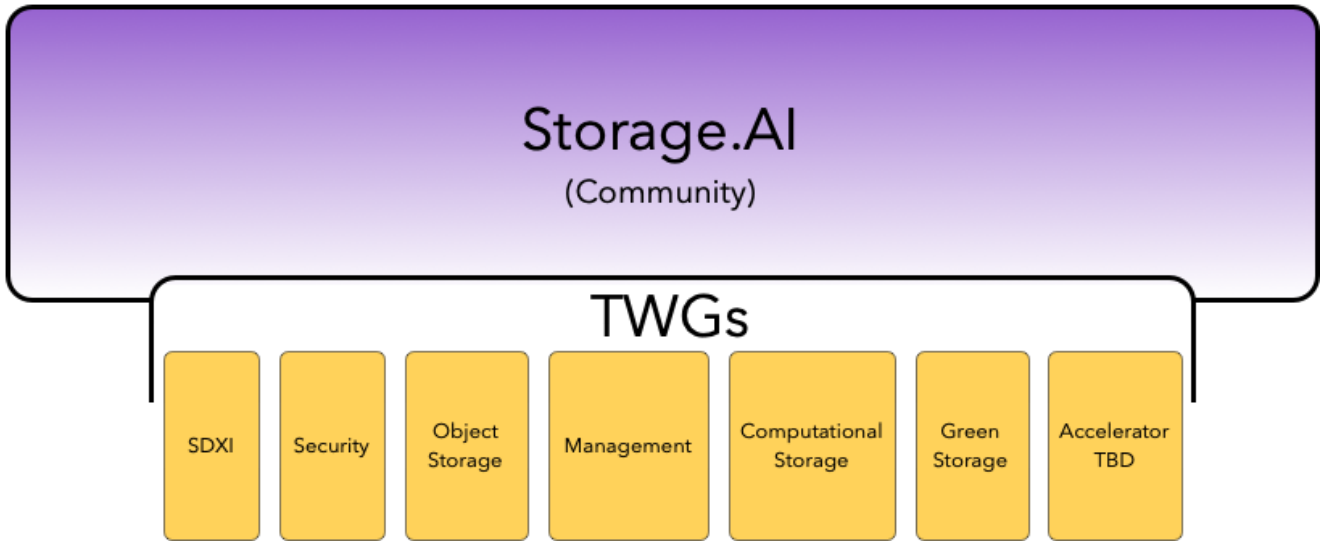

The Process and the (Intended) Deliverables

- Standards
- Software tools (including Management)
- Programming Models
- Performance and Testing Tools
- Architecture and Reference designs
- Education materials (including Developer!)
- Regulation positions
- Joint industry workstreams

The Big Three

- SNIA Members
  - Reduced R&D burden
  - Faster Adoption Cycles
  - Efficiency, Reuse Common Building Blocks
  - Open Collaboration boosts Credibility and Trust
- AI Implementers
  - Streamlined Pipelines
  - Better GPU utilization
  - Vendor Neutrality
  - Future-Proofing
  - Operational Efficiency
- Regulations
  - Transparency & Accountability
  - Sustainability
  - Risk Reduction
  - Policy Confidence
  - Global Harmonization

The Future is Collaborative

- Broad industry ecosystem partnerships
  - UEC
  - NVM Express
  - OCP
  - OFA
  - DMTF
  - PCI-SIG
  - The GreenGrid
  - SPEC
  - UALink (*coming soon!*)
- Join the Storage.AI community
  - --> Use the QR Code!

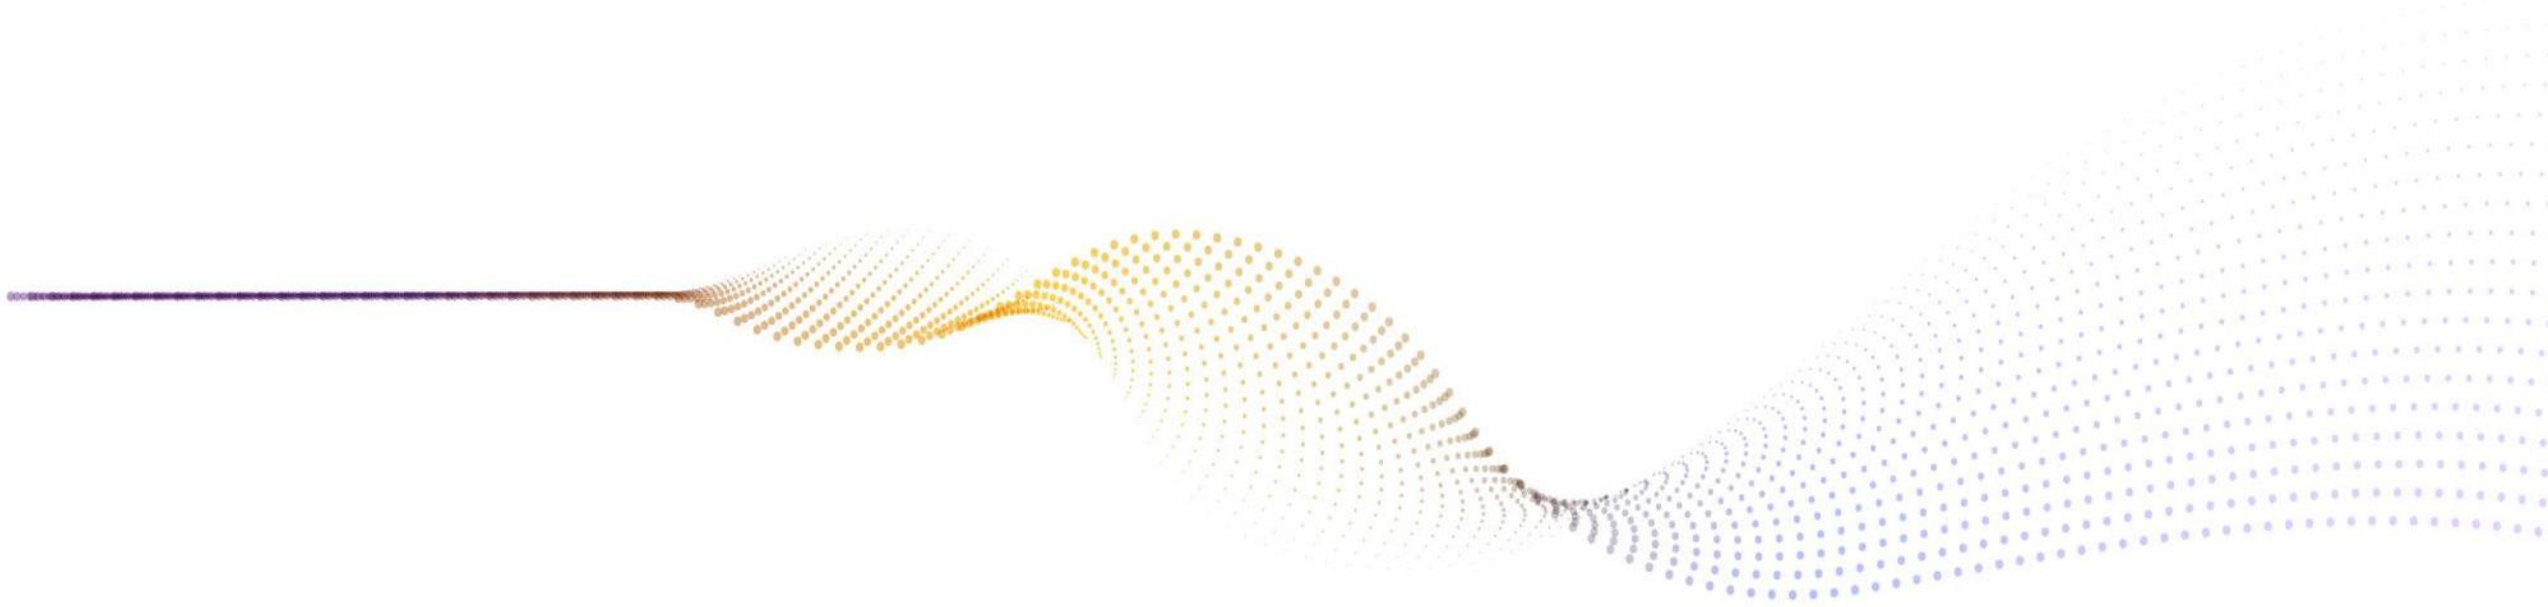

Thank you for attending!

Please remember to rate this session. You get access the presentations at
http://sniadeveloper.org/conference

