

# Computational Storage Does it have a Future?

Scott Shadley

Director, Long-Term Strategy, Solidigm

Executive Board Member, SNIA



# Legal Disclaimers



All product plans, roadmaps, specifications, and product descriptions are subject to change without notice.

Nothing herein is intended to create any express or implied warranty, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, or any warranty arising from course of performance, course of dealing, or usage in trade.

For copies of this document, documents that are referenced within, or other Solidigm literature, please contact your Solidigm representative.

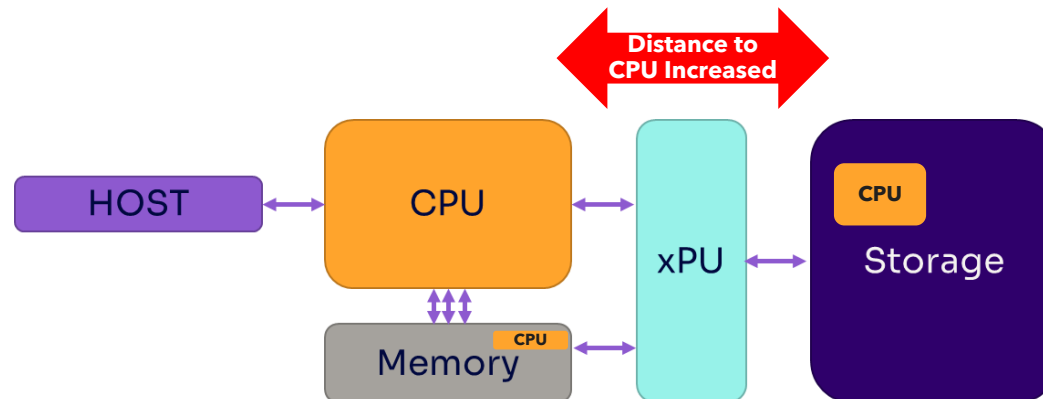
All products, computer systems, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.

© Solidigm. "Solidigm" is a trademark of SK hynix NAND Product Solutions Corp (d/b/a Solidigm). "Intel" is a registered trademark of Intel Corporation. Other names and brands may be claimed as the property of others.

# John von Neumann - The Princeton Architecture



- The Ecosystem today, we have our friend J. v N.
  - CPU/Memory/Storage
- The world is evolving and there is a need for compute in more available locations.
  - Enter the world of “Accelerators” - SmartNIC, xPU, DPU, GPU, IPU
- These are great, but there is room for more!
  - History tells us this much...




# What has Changed? – DATA and AI



## Why Now?


- Storage is no longer 'SLOW'
- Memory is no longer 'Gated'
- Data Gravity, Data Size, Data Locality
- Edge Data Explosion, Transport issue!
- SNIA, NVMe, CXL, OCP, Others are providing new guidance in new areas of implementation



## Key Benefits?

- Faster, Fewer, Easier I/O transfers
- Reducing DRAM/Network tax with new transports, solutions, products
- Redeploy Primary CPU to High Value Work, offer up new services to help
- Improved AI performance due to parallelism for certain workloads
- Better scheduling of data Machine Learning modeling and AI inferencing

New Standardized API & Programming Model




# Amdahl's Law

## Amdahl's Law

Measures the theoretical speedup of a program's execution latency as a function of the number of processors executing it...

**The limiting factor is 'Serial' so move it!!**

- In Memory? Sure...
  - New areas of "Computational Memory"
- In Storage? Why Not...
  - Computational Storage Drives, Arrays
- In the Middle? Of Course...
  - Computational Storage Processors
  - SmartNIC, xPU...

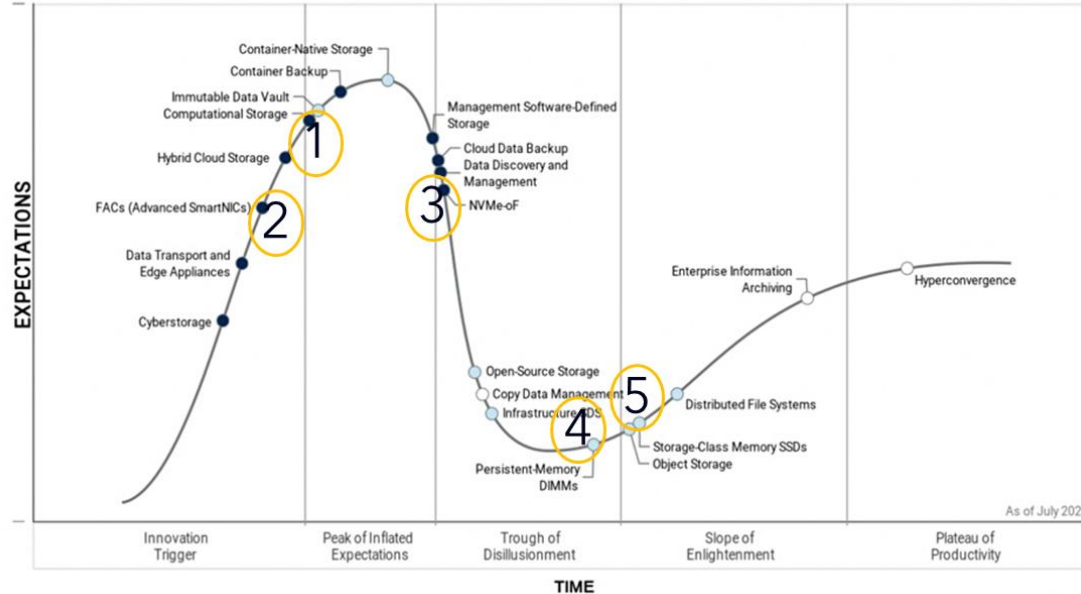


# The “Official” Storage Hype Cycle – One Year later

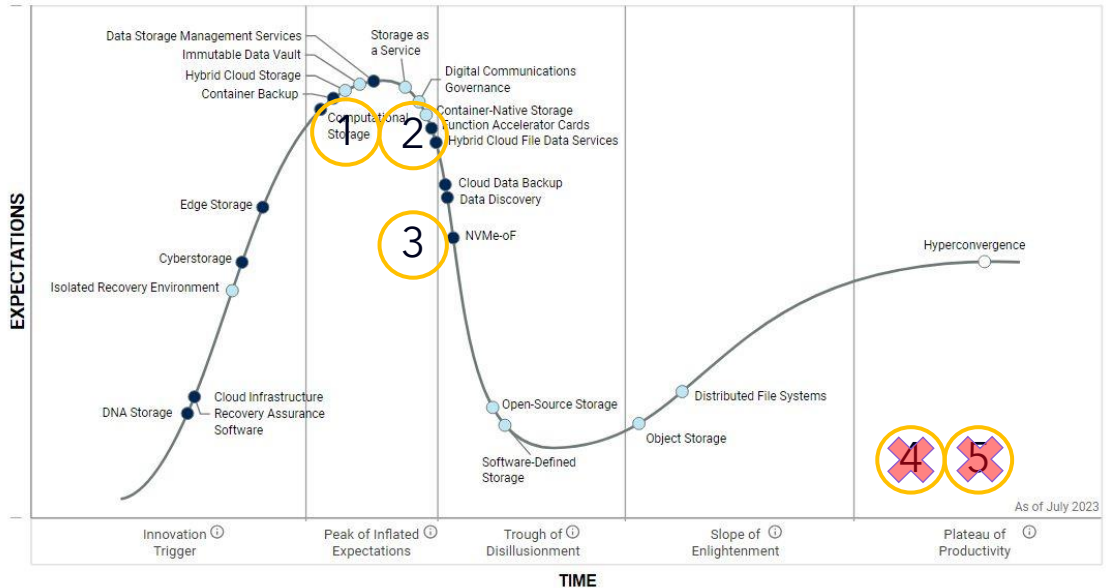


Plateau will be reached: ○ <2 yrs. ● 2–5 yrs. ● 5–10 yrs. ▲ >10 yrs. ✗ Obsolete before plateau

Hype Cycle for Storage and Data Protection Technologies, 2022



Hype Cycle for Storage and Data Protection Technologies, 2023




1. Computational Storage
2. FACs
3. NVMe-oF
4. PM DIMMs
5. Storage-Class Memory SSDs

# The Forefront of Change - DATA, DATA, DATA




Volume of data/information created, captured, copied, and consumed worldwide from 2010 to 2020, with forecasts from 2021 to 2025



Source: [Total data volume worldwide 2010-2025](#) | Statista

Number of data centers worldwide in 2022, by country



Source: [Data centers worldwide by country 2022](#) | Statista

# Workloads Abound for AI with a CSD



- Did you see the Demo in the booth?
  - Data Integrity Solution on a Computational Storage Drive
- Many other workloads exist and support AI Growth
  - Data locality often overlooked when deploying AI solutions
    - Federated Machine Learning
    - Inference
    - Sorting, Matching, Searching
- Stay Tuned for more from NVMe, SNIA, OCP, and Solidigm
  - Standards, Command Structures, Guidelines, Products





