

STORAGE DEVELOPER CONFERENCE



Fremont, CA
September 12-15, 2022

BY Developers FOR Developers

A **SNIA** Event



DNA DATA STORAGE ALLIANCE

A SNIA Technology Affiliate

The Looming Need for Molecular Storage

How to store Yottabytes on a budget

Aaron Ogus

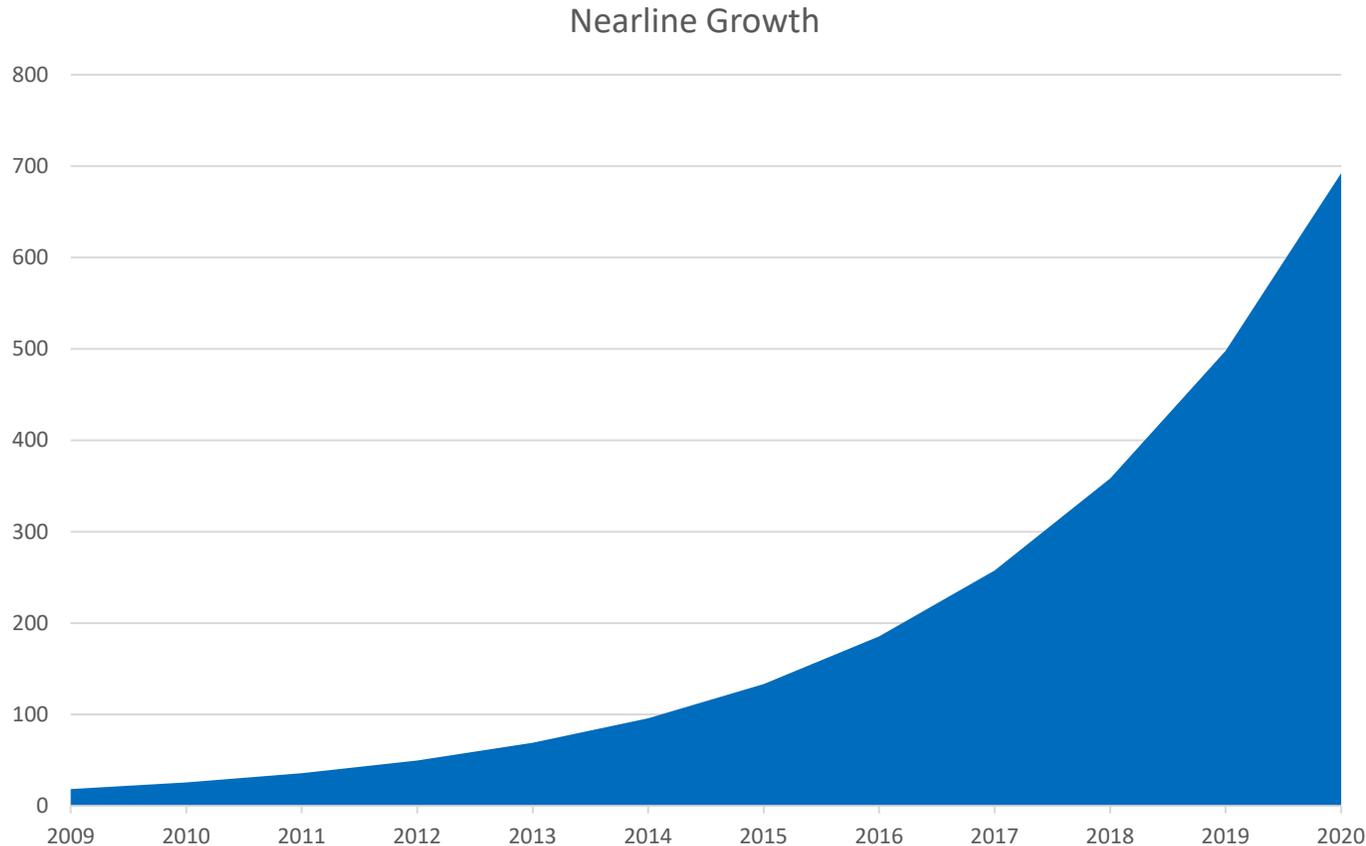
Distinguished Engineer – Microsoft Azure Storage

How to we store Exabytes of Data?



- Many kilometers of Storage Racks, in hundreds of Datacenters.
- Hundreds of MW of power.

The vast majority of data is moving to the Cloud



- In 2010's HDD shifted from Consumer to DC. By 2030 almost all HDD will be nearline in the cloud.
- Underlying data growth signal is unwavering, and projects to over **7 ZB** per year by 2030 <not shown>
- Data CAGR about 40%
- HDD Capacity CAGR < 20%

Storage Efficiency Improvements 2008 to 2022



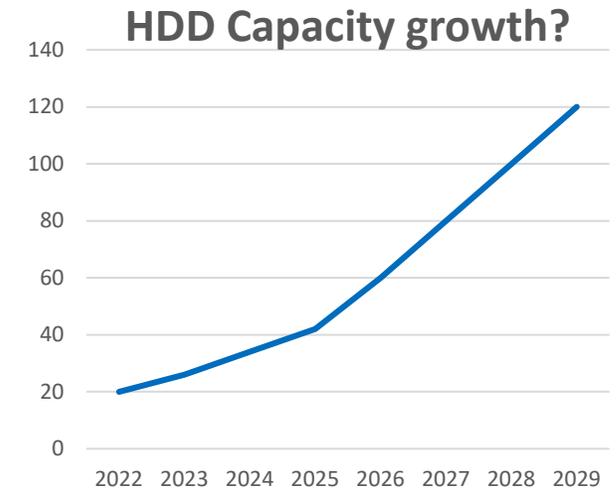
- 2006: 500GB HDD, 3 Replica, 4 HDD/Server
- 2008: 1 TB HDD, 3 Replica, 12 HDD/Server
- 2021: 20 TB HDD, <1.3 Replica, 80+ HDD/Server



- **> 99.5% reduction in online storage cost over 15 years**

Efficiency Improvements for HDD continue... but.

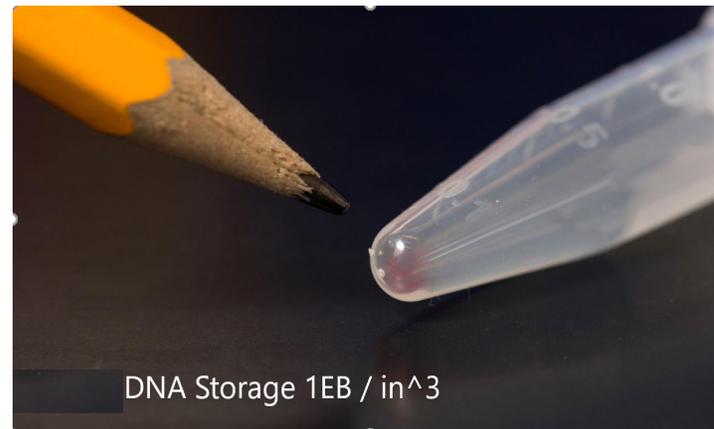
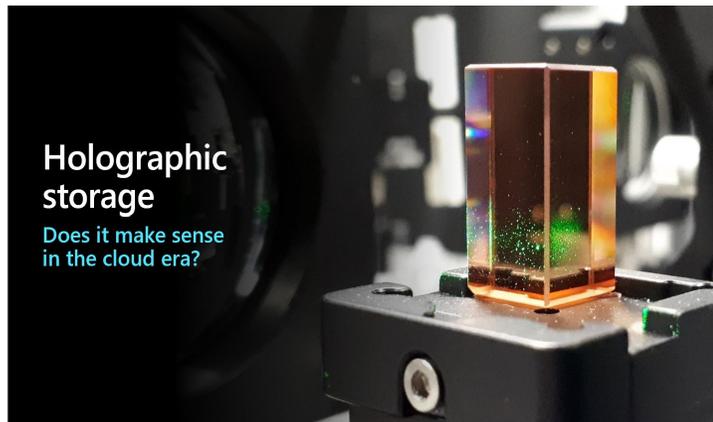
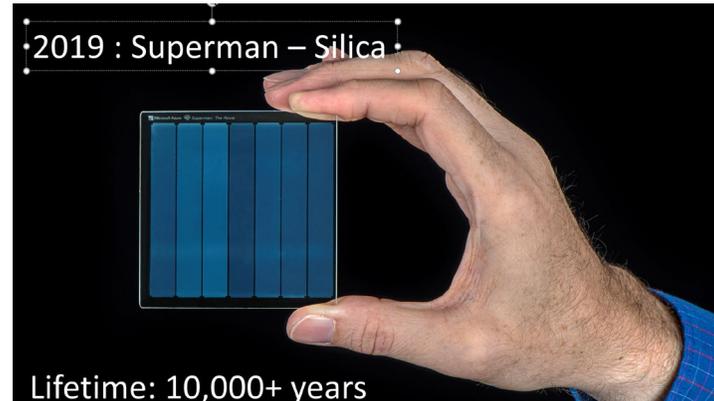
- HDD Suppliers are shifting to MAMR and HAMR
- Only HAMR has legs to 60 TB+
- HAMR might run out of steam around 100TB
- Optimistically if we assume HDD growth to 230TB+, we will require 5% of current US Power capacity to be dedicated to spinning HDDs by 2042, and 60% by 2050



Something has to change:

- Data growth has to slow.
- Power Generating capacity needs to vastly increase.
- Data Storage Technology needs to change.

Azure Storage + MSR researching new media types



- Data storage in DNA
- Data storage in Silica Glass
- Holographic Data storage in Sodium Niobate crystals

Molecular Storage the Panacea for density

- Highest Storage Density of Any Technology
- DNA Storage Alliance is building standards around DNA storage.
 - First commercial products based on data storage in molecules.
- Problem: Read and Write Rates
- Electronic Molecular interface... based on medical research
- Can we build fast Molecular Storage?

Where is most of Humanity's Data Stored?

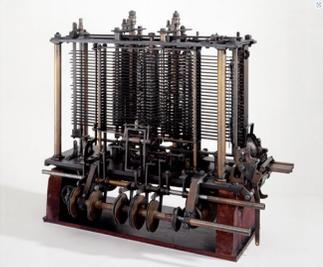
- HDD industry shipped 1ZB for the first time in 2021
- 1,000,000,000,000,000,000,000 bytes 10^{21}
- Power if on 20TB HDDs @10W each would be:
50Million HDDs and 500MW.

500 MW/ZB

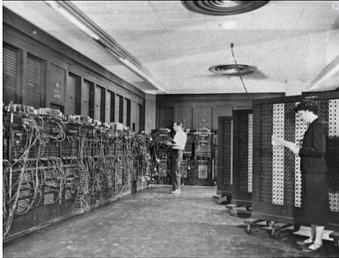
- What about Human Brains?
- 8 Billion Humans
- 2.5PB per brain (estimated)
- Power @ 20W per brain: 160GW
- 2 Yottabytes in Brains (About 1000x data stored on HDD)

8 MW/ZB

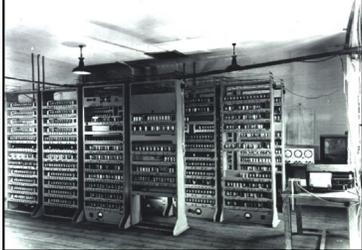
Compute Evolution



Analytic Engine
Designed 1820
Metal Gears
Mechanical



ENIAC
1945
Vacuum Tubes
Crystal Diodes

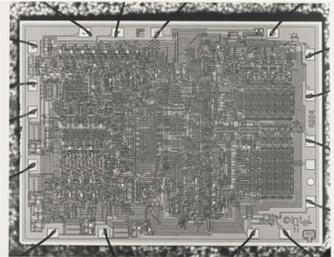


EDSAC
1949
Thermionic Valves

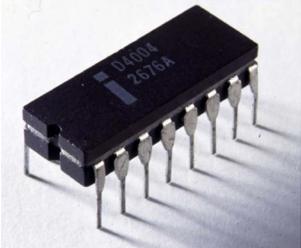
1955 Burroughs
Atlas
Transistors



1957
Univac
Transistors



Intel 4004: First computer on an IC
1971



Intel Xeon SPR
2021, 100,000,000x
improvement over 4004

Storage needs a Platform shift for YottaScale (2042)

- Today's Storage Platforms:
 - Magnetic Charge (Tape/HDD/Core Memory)
 - Electric Charge (Flash, RAM)
 - Optical
 - Holographic (Lithium Niobate and Lasers)
 - Molecular (DNA)
- Future do we need fast Molecular Storage?
 - DNA based?
 - Can we build the toolchain for molecular machines?

What's next?

- Research in Molecular Simulation is Ramping:
 - Alpha Fold - Google
 - AI4Science (announcement) - Microsoft
 - Medical Research
- Resolving the Molecular / Electronic interface is critical
- Possibility to leverage patterns from Biology? Molecular Machines?
- YottaScale 2042(ish), XanoScale 2063(ish)



Please take a moment to rate this session.

Your feedback is important to us.