



# Storage Scaling ZettaScale to XanoScale

Aaron Ogus  
Distinguished Engineer  
Microsoft Azure Storage

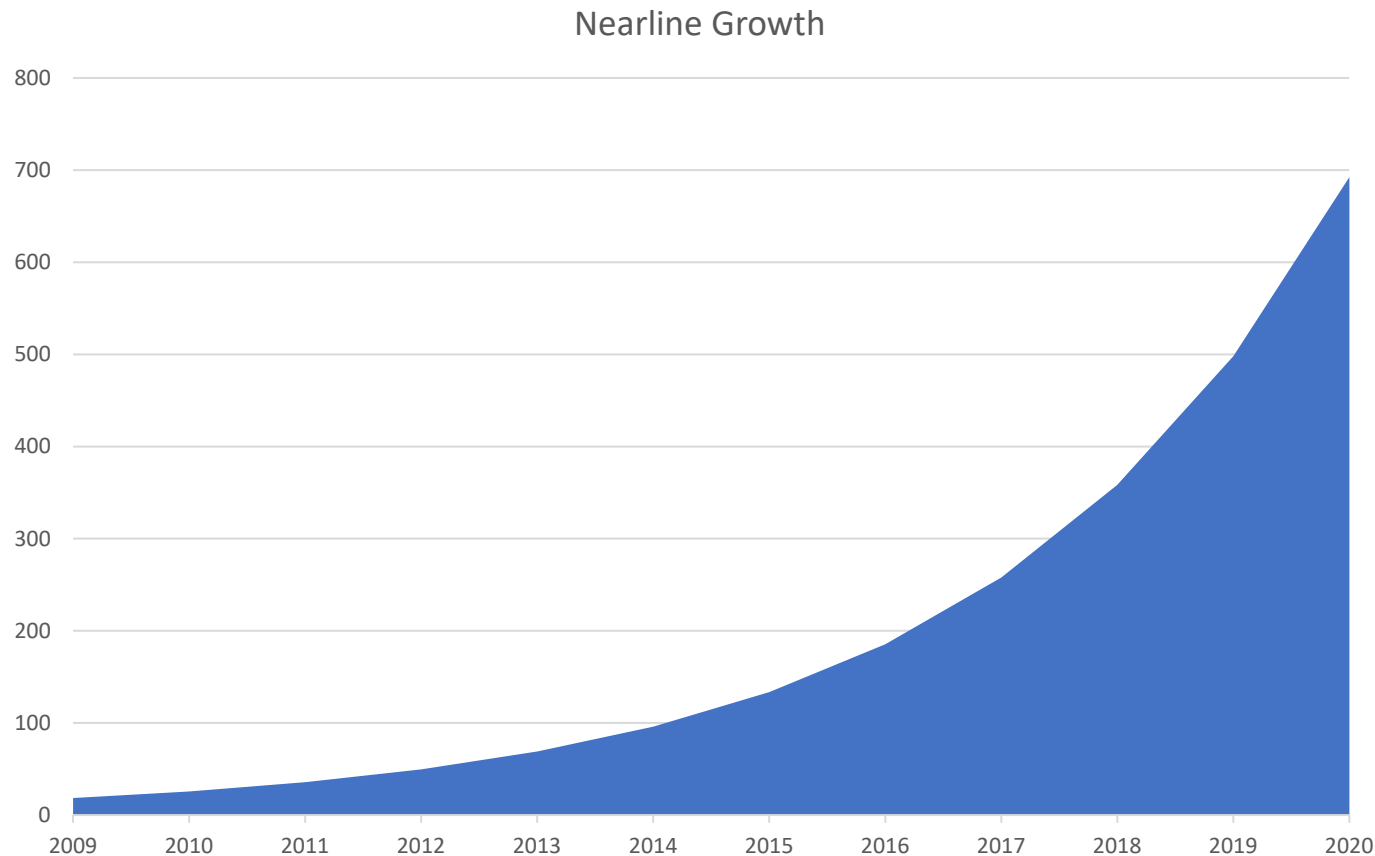
- March 28, 2023



How do we store EBs of Data?

- Literally Miles and Miles of Storage Racks
- Hundreds of MW

# Data Headed to the Cloud



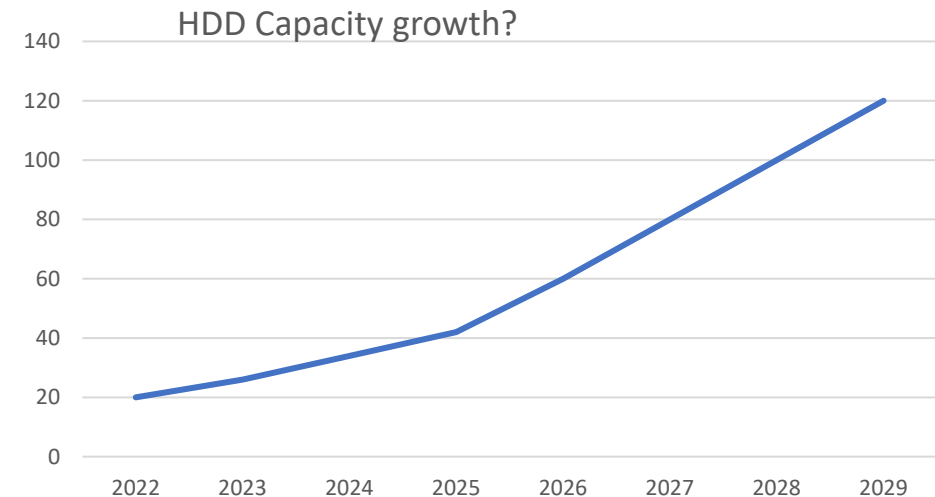
In 2010's HDD shifted from Consumer to DC. By 2030 almost all HDD will be nearline in the cloud.

Underlying data growth signal is unwavering, and projects to over **7 ZB** per year by 2030 <not shown>

Data CAGR is 40%  
HDD Capacity CAGR < 20%

# HDD Roadmap

- 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,  
Generating capacity needs to vastly increase  
Storage technology needs to change

# Storage Research – MSR

Azure Storage and MSR collaborate on research and commercialization of new media types for storage.



# Where is most of Humanities Data Stored?

HDD shipped 1ZB for the first time in 2021

1,000,000,000,000,000,000,000 bytes  $10 \cdot 10^{22}$

500 MW/ZB

Power if on 20TB HDDs @10W each would be: 50Million HDDs and 500MW.

What about Human Brains?

8 Billion Humans

2.5PB per brain (estimated)

8 MW/ZB

Power @ 20W per brain: 160GW

2 Yottabytes in Brains (About 1000x data stored on HDD)

# 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)