SEAGATE

Mass Capacity Storage in an AI Era

Library of Congress Conference on Data Storage

March 9, 2026

Jon Trantham
Principal Technologist
Seagate Research

Hard drive storage through the years

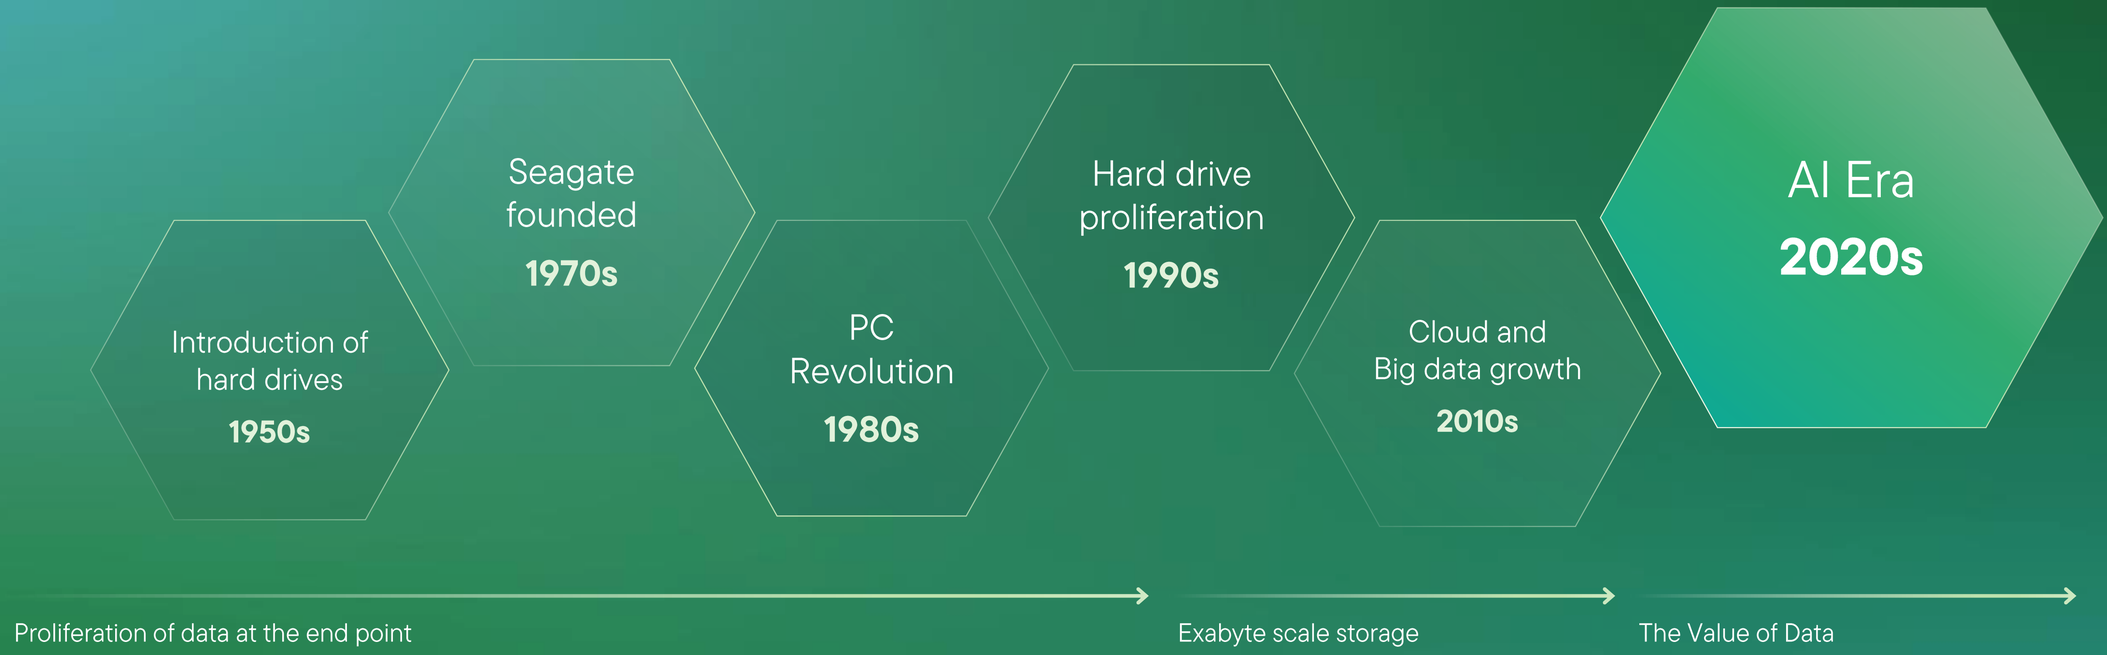

The AI world we live in today...

11M

Videos created per day with Sora2

600M

Videos created with Kling since launch

250M

Videos created on Veo3 since launch

26M+

AI generated music tracks on TME since launch

10M+

Monthly mobile downloads of Jimeng

5B

Images created with Nano Banana since launch

370M

Videos created on Hailu since launch

AI-generated content has surpassed human content

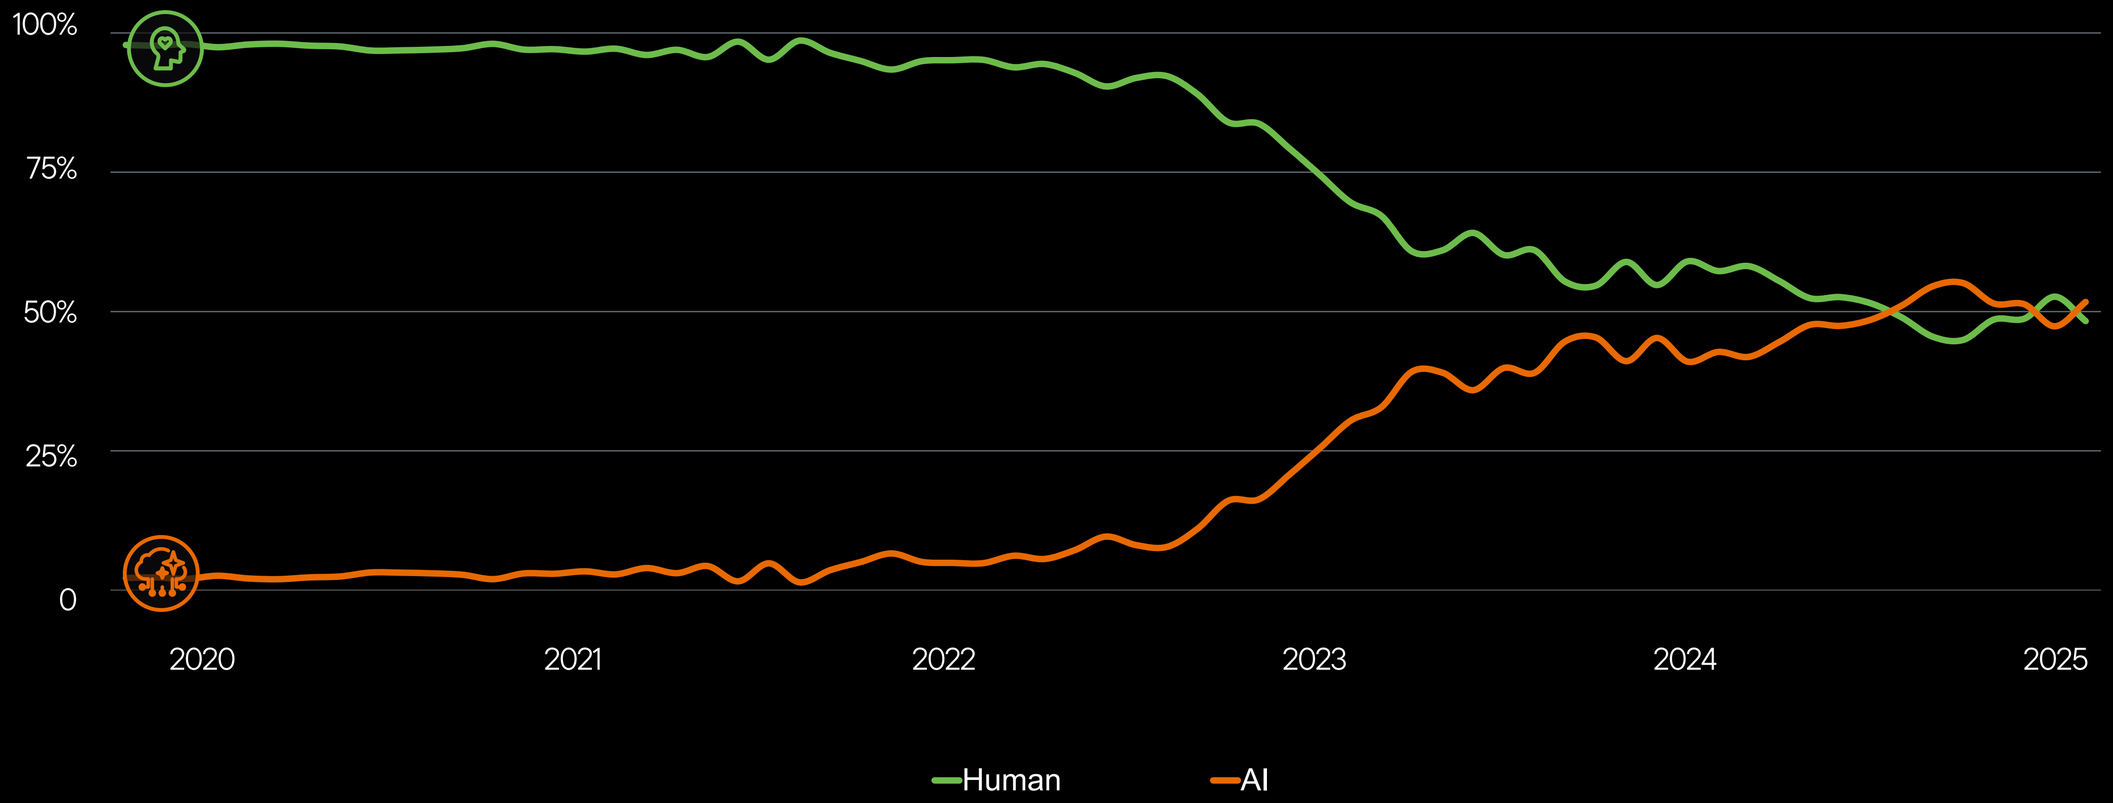

Source: SEO firm Graphite, October 2025

Whether consumer or enterprise, video will make up the bulk of what we store

Video—fueled by streaming, gaming, social media, video analytics, and user-uploaded Generative AI—is over 80% of internet traffic and well over half of all data generated and stored.

- **Nearly 100% of China internet users** consume short videos regularly
- **An estimated 90% of online learners** use YouTube to consume educational long-form videos.
- **93% of sales professionals** reporting that sending personalized video messages (instead of only emails) increased their customer response rates.

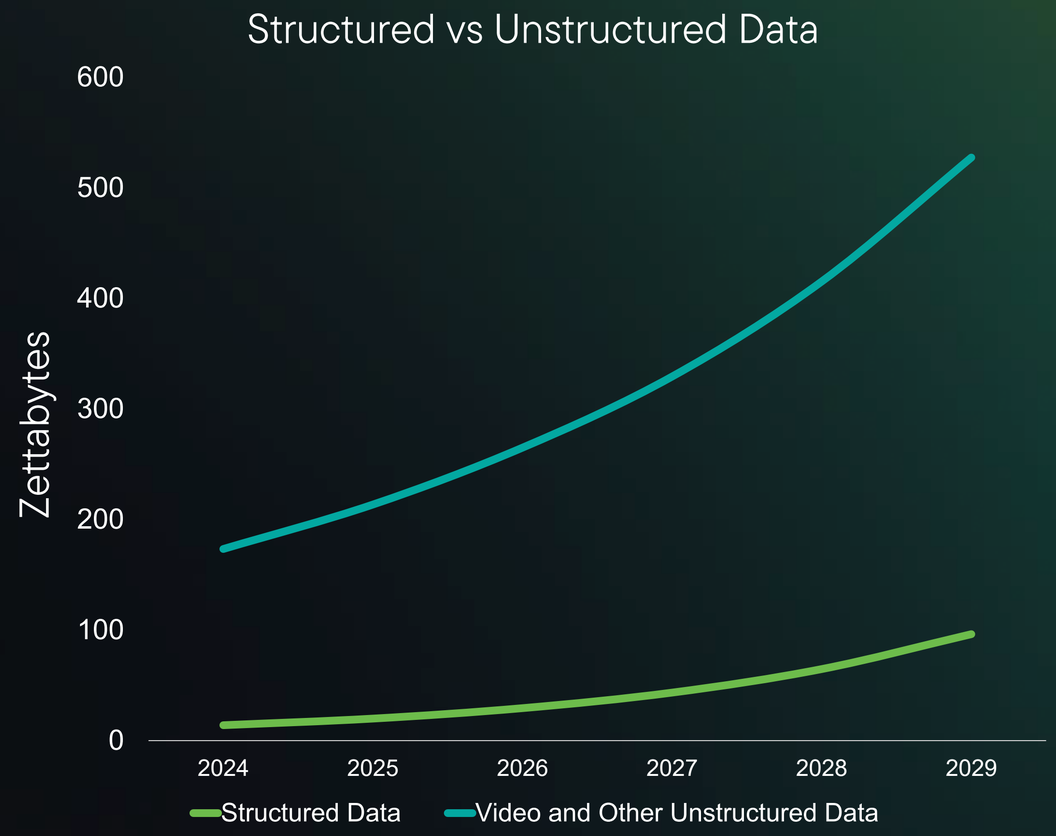

Technology innovation is generating unprecedented volumes of data

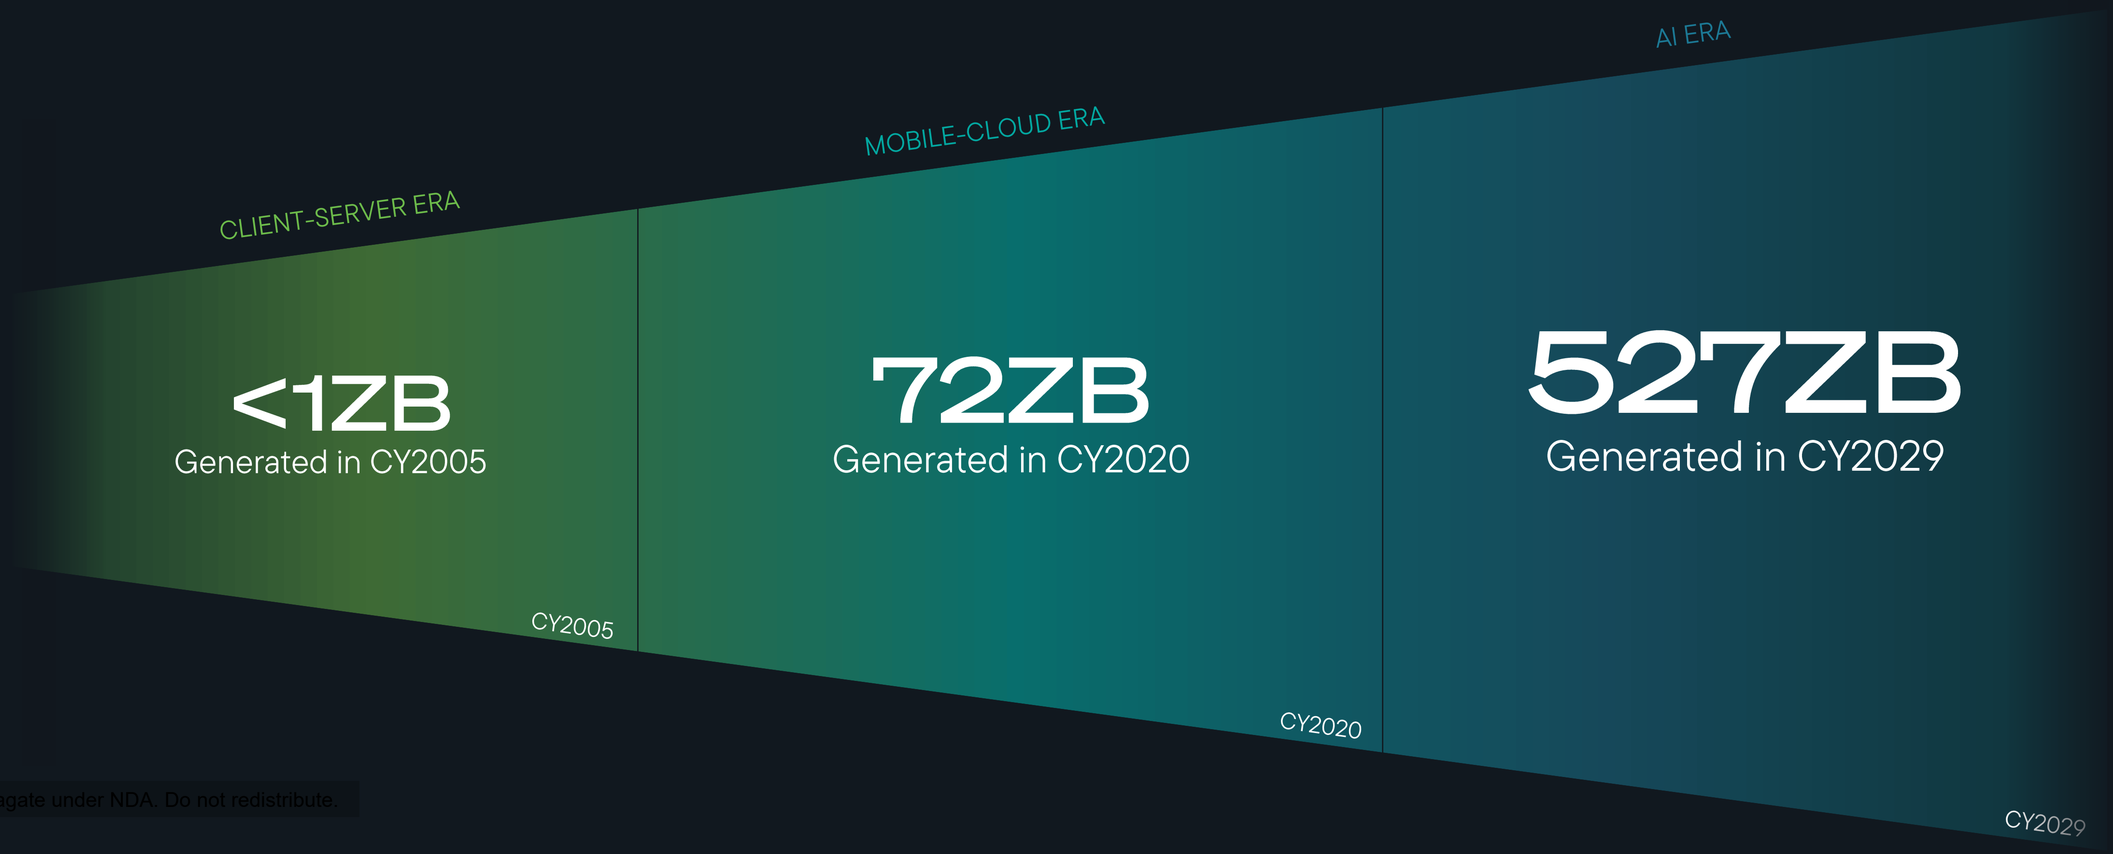

Cloud Storage growth in the AI era

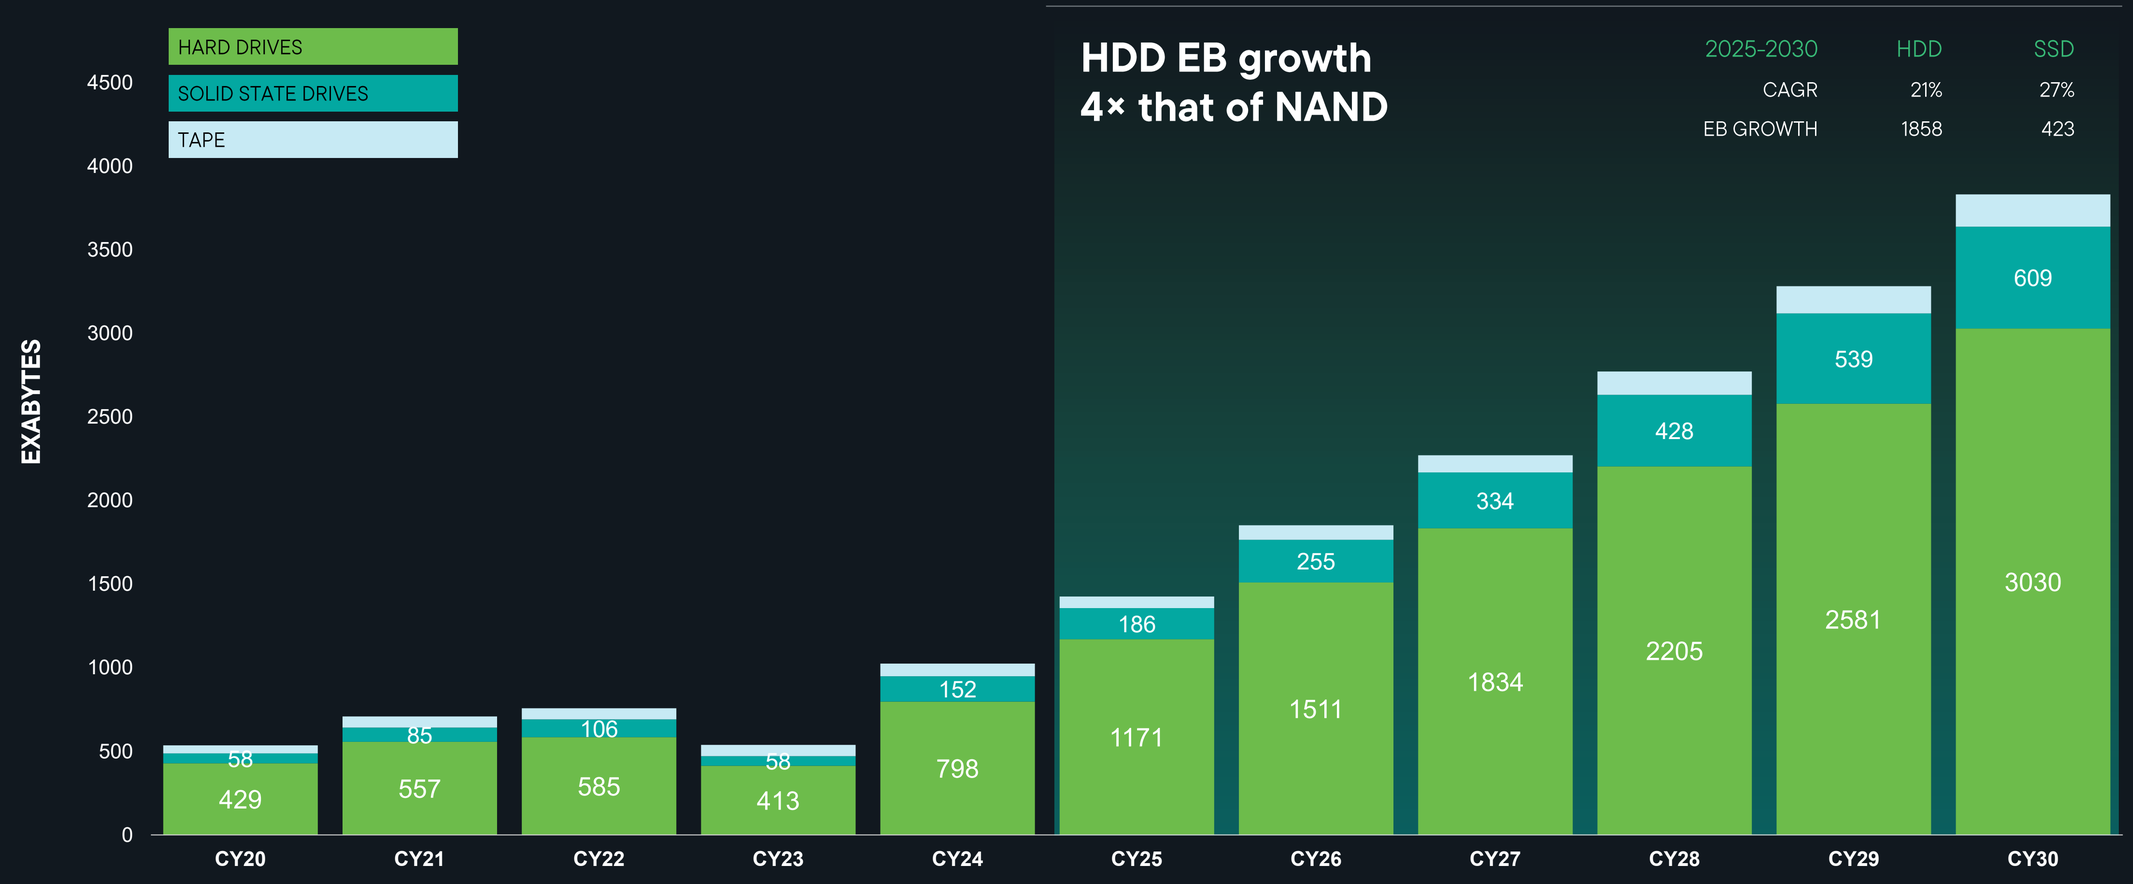

Source: HDD = STX Market Research DecEBO/LRP100 / STX Extrapolation, SSD = IDC 2025 StorageSphere_CMI_Q425 SSD / STX Extrapolation, TAPE = Trend Focus Q425 / STX Extrapolation

Growth in Storage is inevitable as the value of data increases

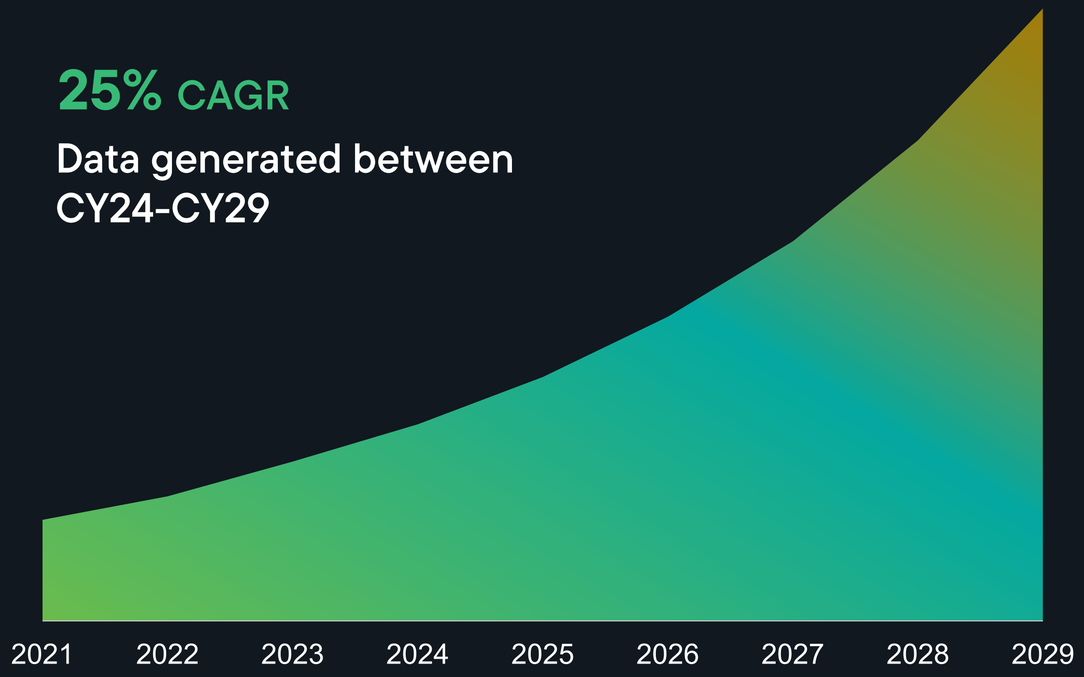

New kinds of content, new data demands

78% of organizations are now generating content they never produced before — everything from 3D product visuals to AI-enhanced training simulations

File size growth with compounding effects

More than 70% say files are already larger and richer, with video and image formats driving massive increases in storage requirements

Archiving with purpose

42% are already adopting tiered and archived storage strategies to manage long-term retention efficiently.

https://www.seagate.com/resources/content-creation-in-the-age-of-generative-ai-implications-for-value-and-scale/

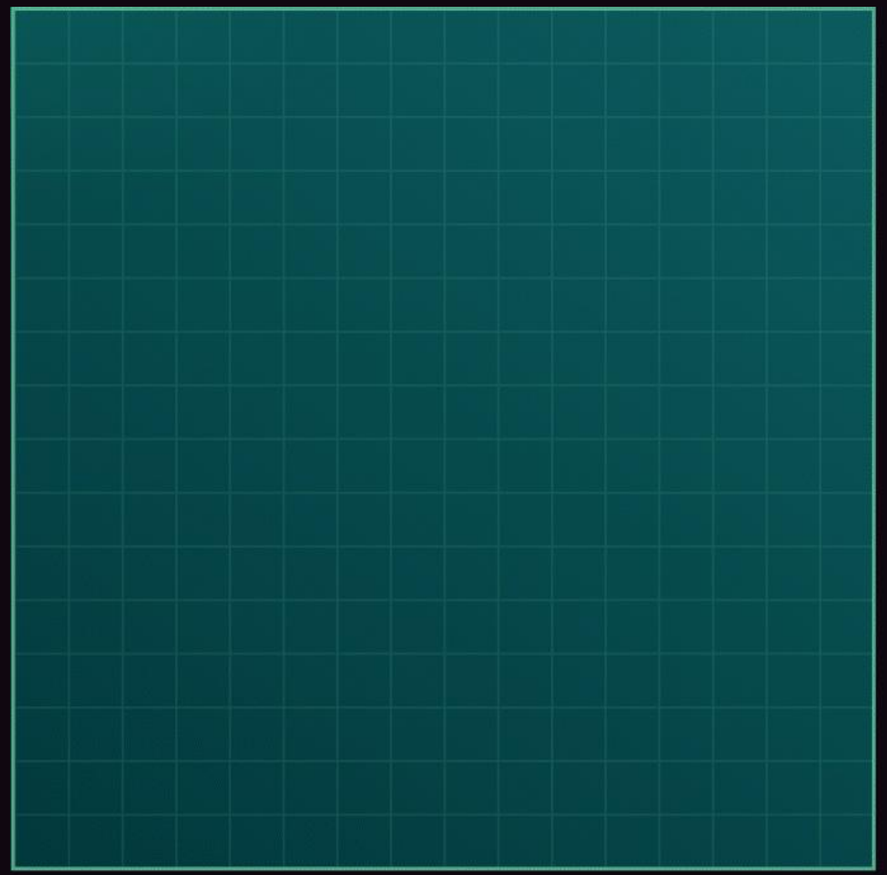

The value of data is timeless.

The system can reliably reach that data again, there is no need to keep everything cached forever.
This reduces cost and avoids large investments that do not scale—especially in cache and compute—while still supporting strong customer experiences.

ENGINEERED FOR VALUE DENSITY

Every Square Foot Is a Strategic Asset

- Leading hyperscalers are unconstrained by per-device IO/TB.
- They leverage software-driven, HDD-tuned operations at fleet scale.
- As a result, they're realizing the Mozaic advantage, increasing capacity-per-square-foot and accelerates their path to customer growth.

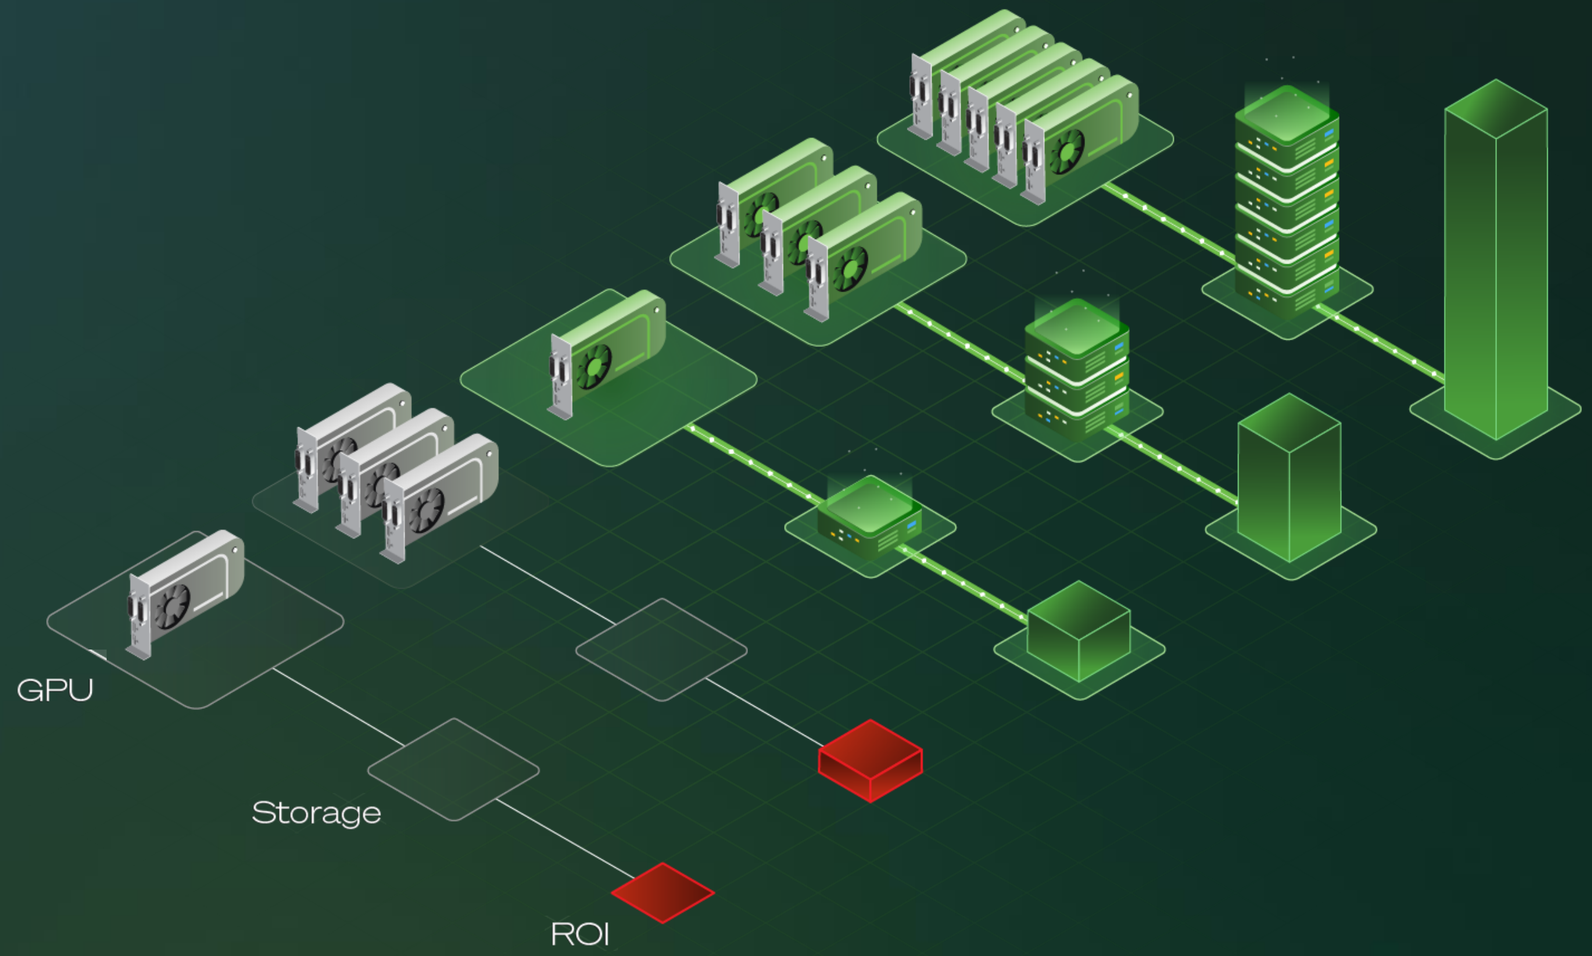

Challenges to achieve high areal density

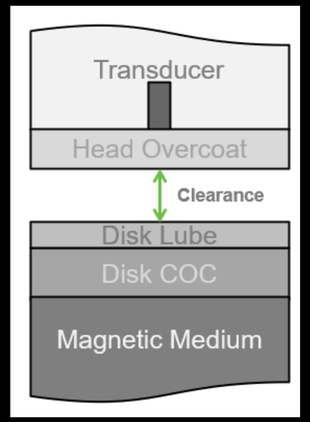

Head-media spacing reduction:

Reducing surface roughness and lube monolayer thickness and new materials integration will be required.

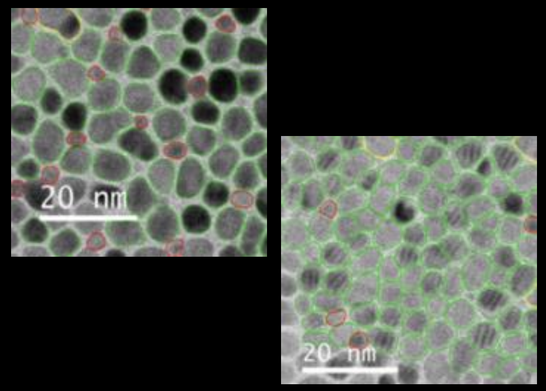

Media grain size reduction:

Need to overcome the loss of moment and anisotropy and widening of Tc distribution caused by finite size effects.

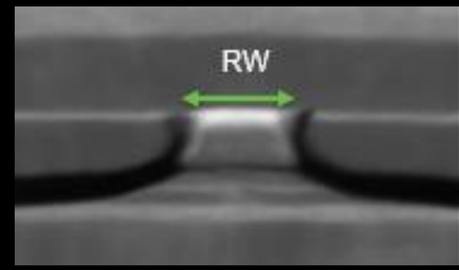

Reader scaling:

Reader geometry limited by lithographical capability & thermal magnetics noise challenges at small sensing layer volumes

HDD Technology Transitions in Scale

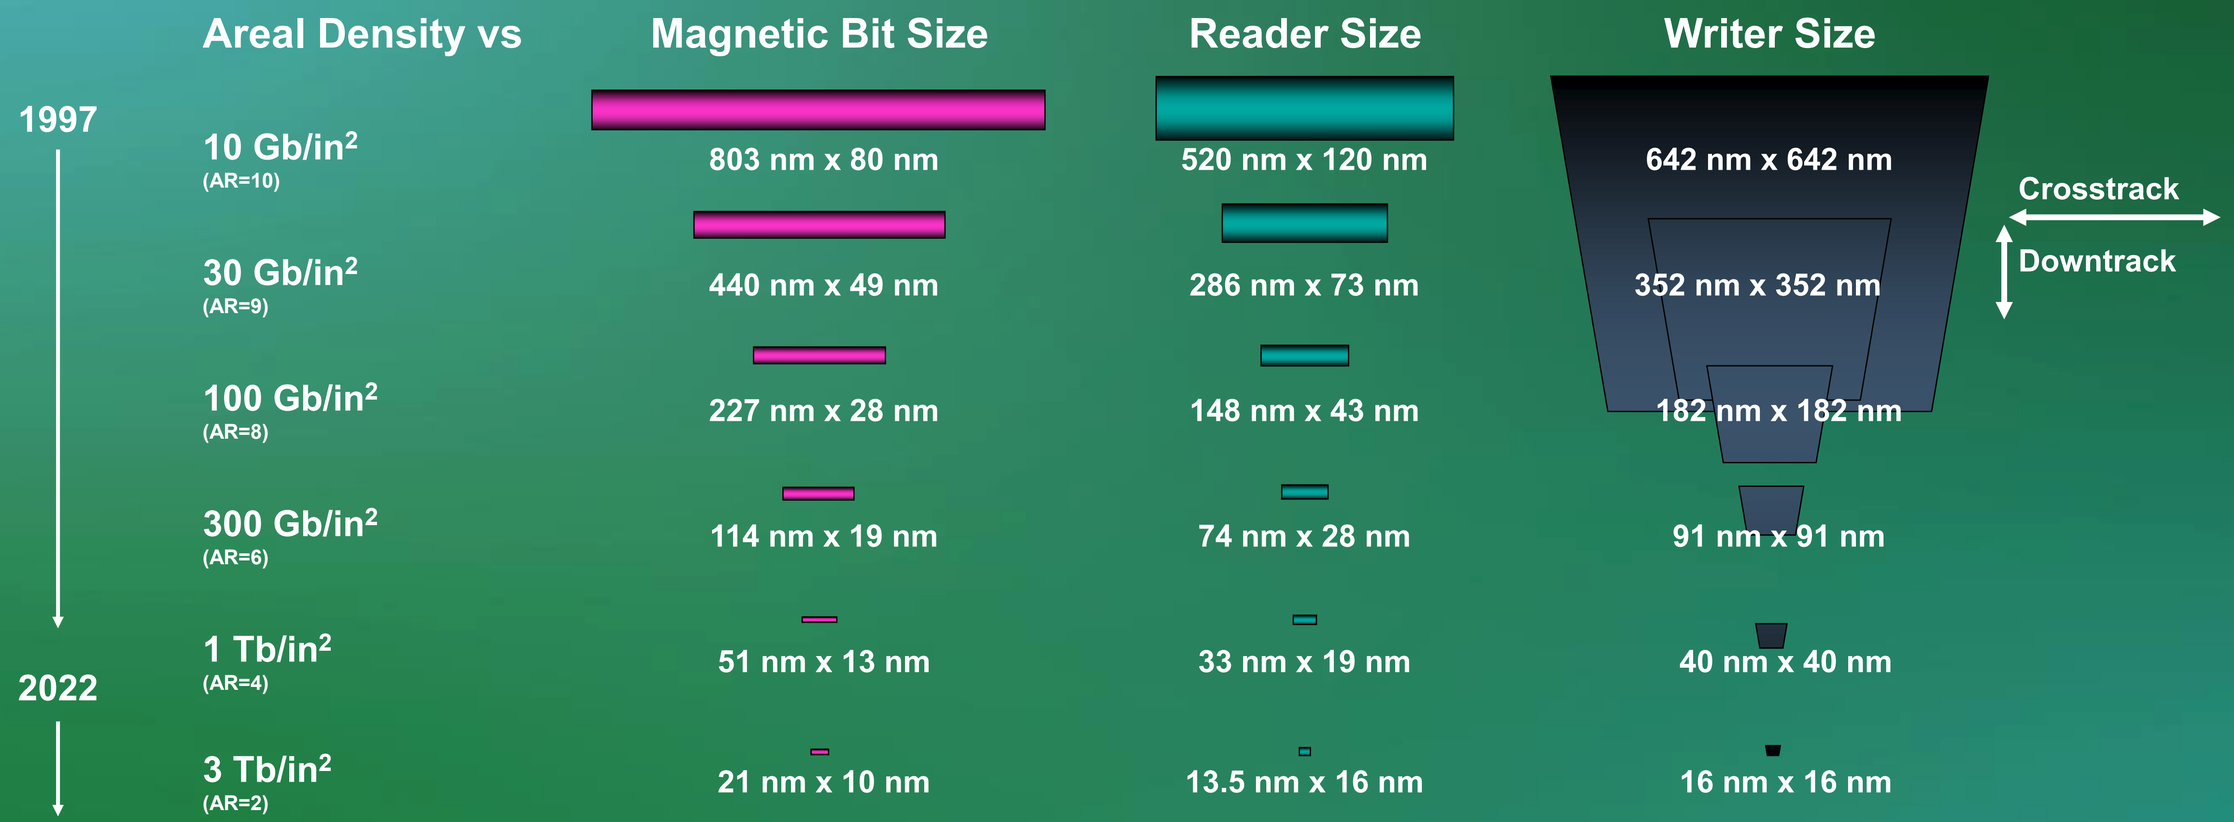

HAMR Areal Density Demonstration

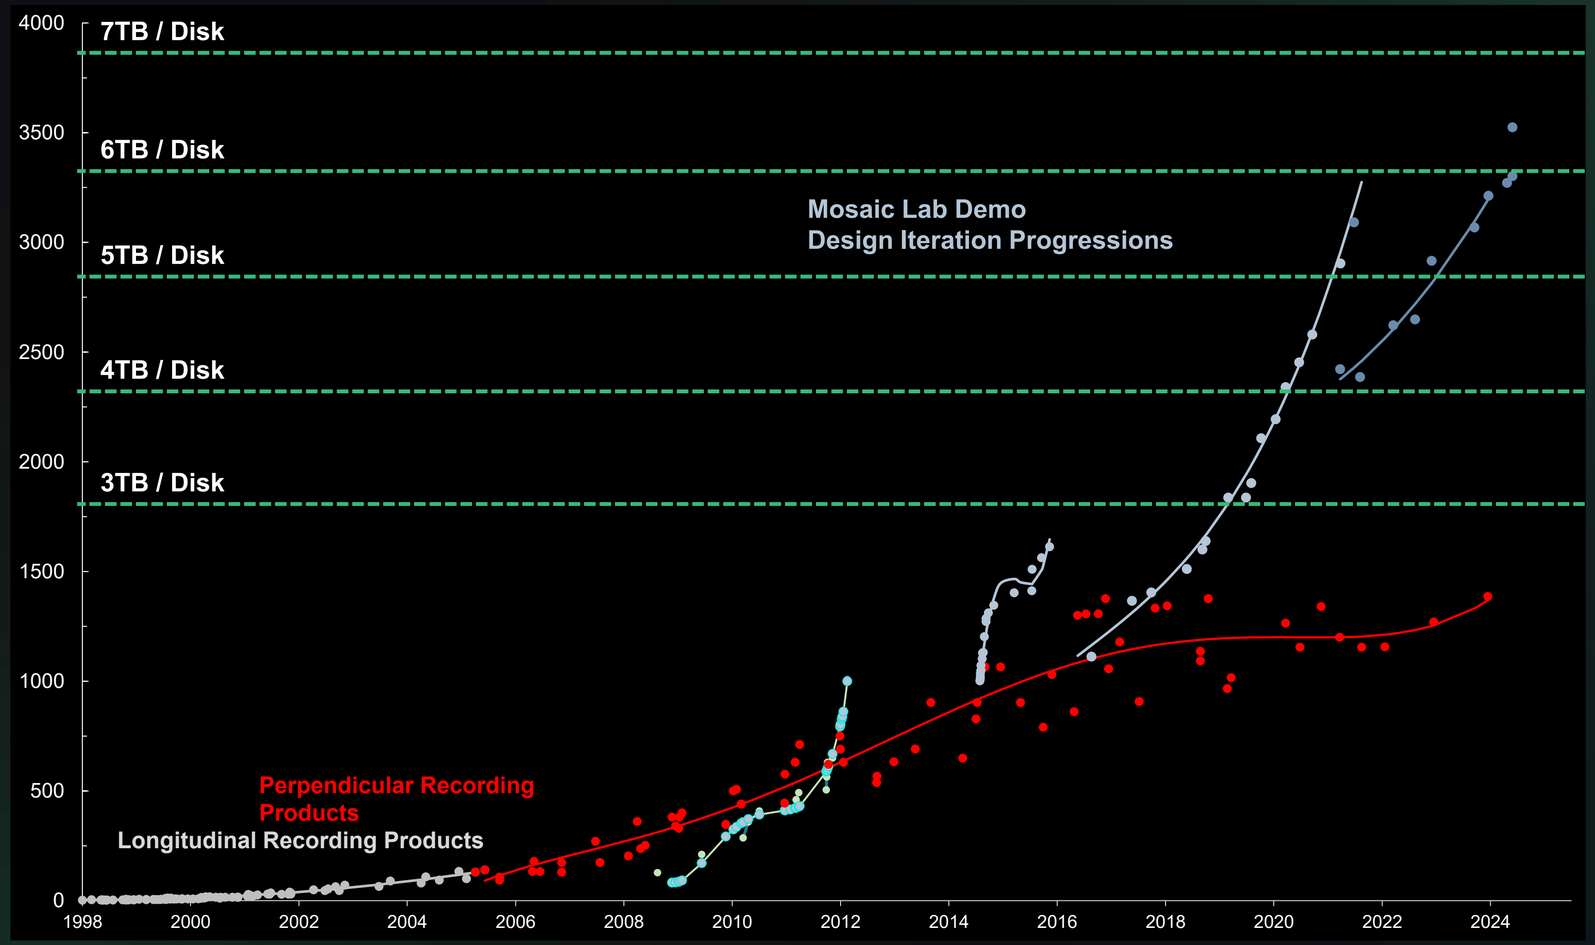

Source: HAMR ADC CMR and SMR with MSMR-2R/3R Gain: Linear Density vs Trackpitch
Steven Granz, Seagate Technology Research, et al., TMRC Magnetic Recording Conference, 2025

Opportunities for 10 TB/disk and beyond

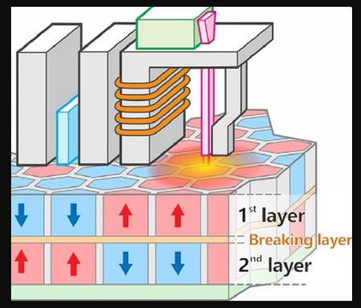

NIMS-Seagate Collaborations

Multi-level recording:

Recording three or more states in out of plane direction.

- Relaxes geometry.
- Challenges include multi-layer media growth, sensing, write process.

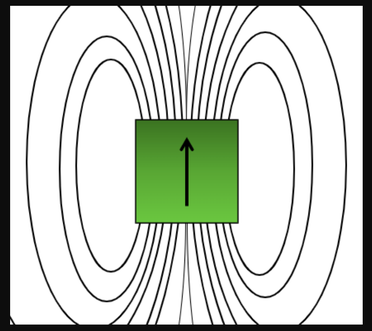

Multi-field component sensing:

Readers designed to sense all field components.

- Relaxes reader geometry.
- Challenges include channel and sensor design.

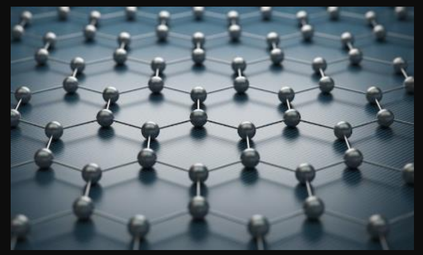

Novel materials for sensors:

2D materials may exhibit magnetoresistive properties.

- Suitability for HAMR needs demonstrating.
- Challenges include manufacturing process integration.

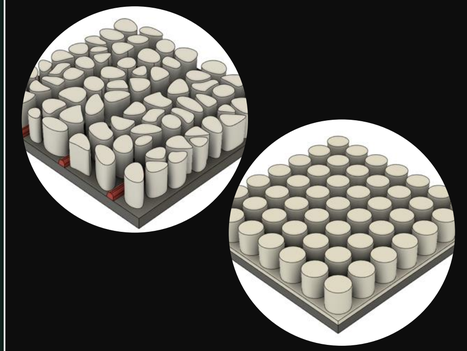

Novel media structures:

Improved media ordering enables high ADC through reduced jitter.

- Ordered granular media may be stepping stone to patterned media
- Low cost process needs innovation.

A clear innovation roadmap for the next decade and beyond

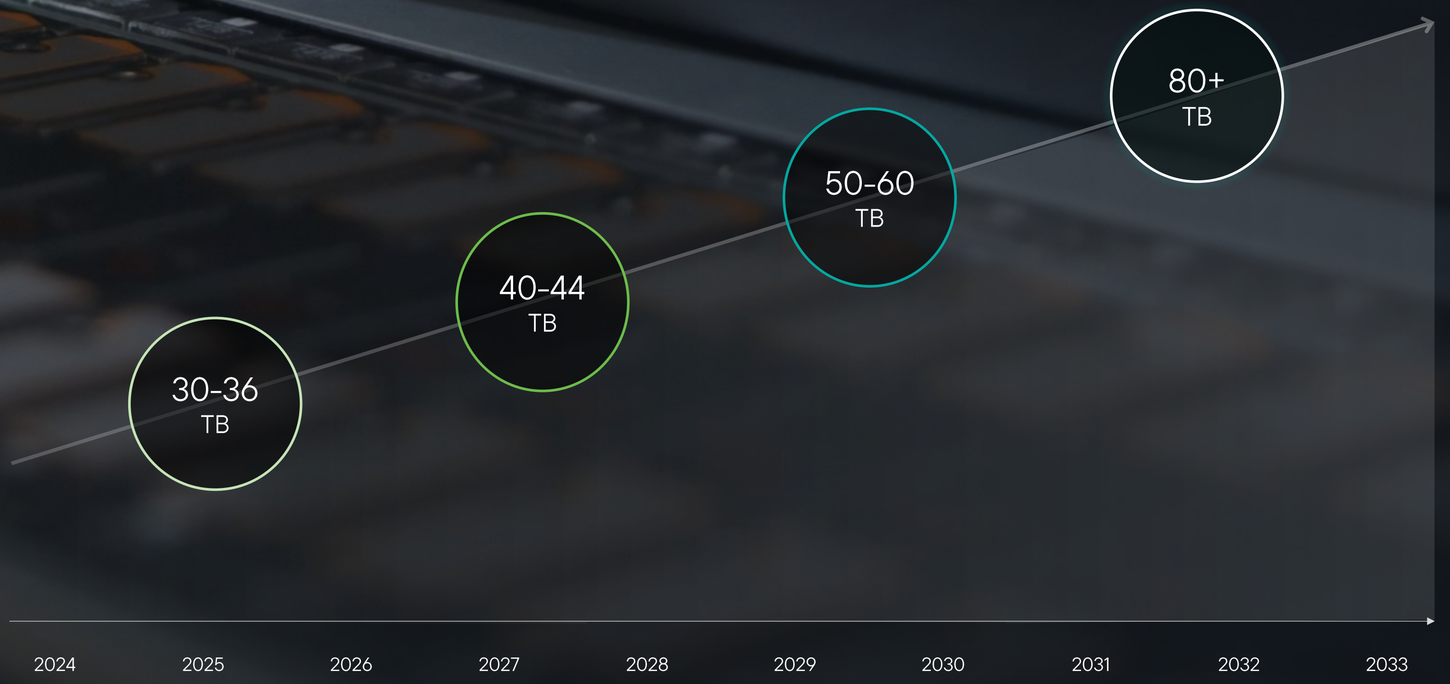

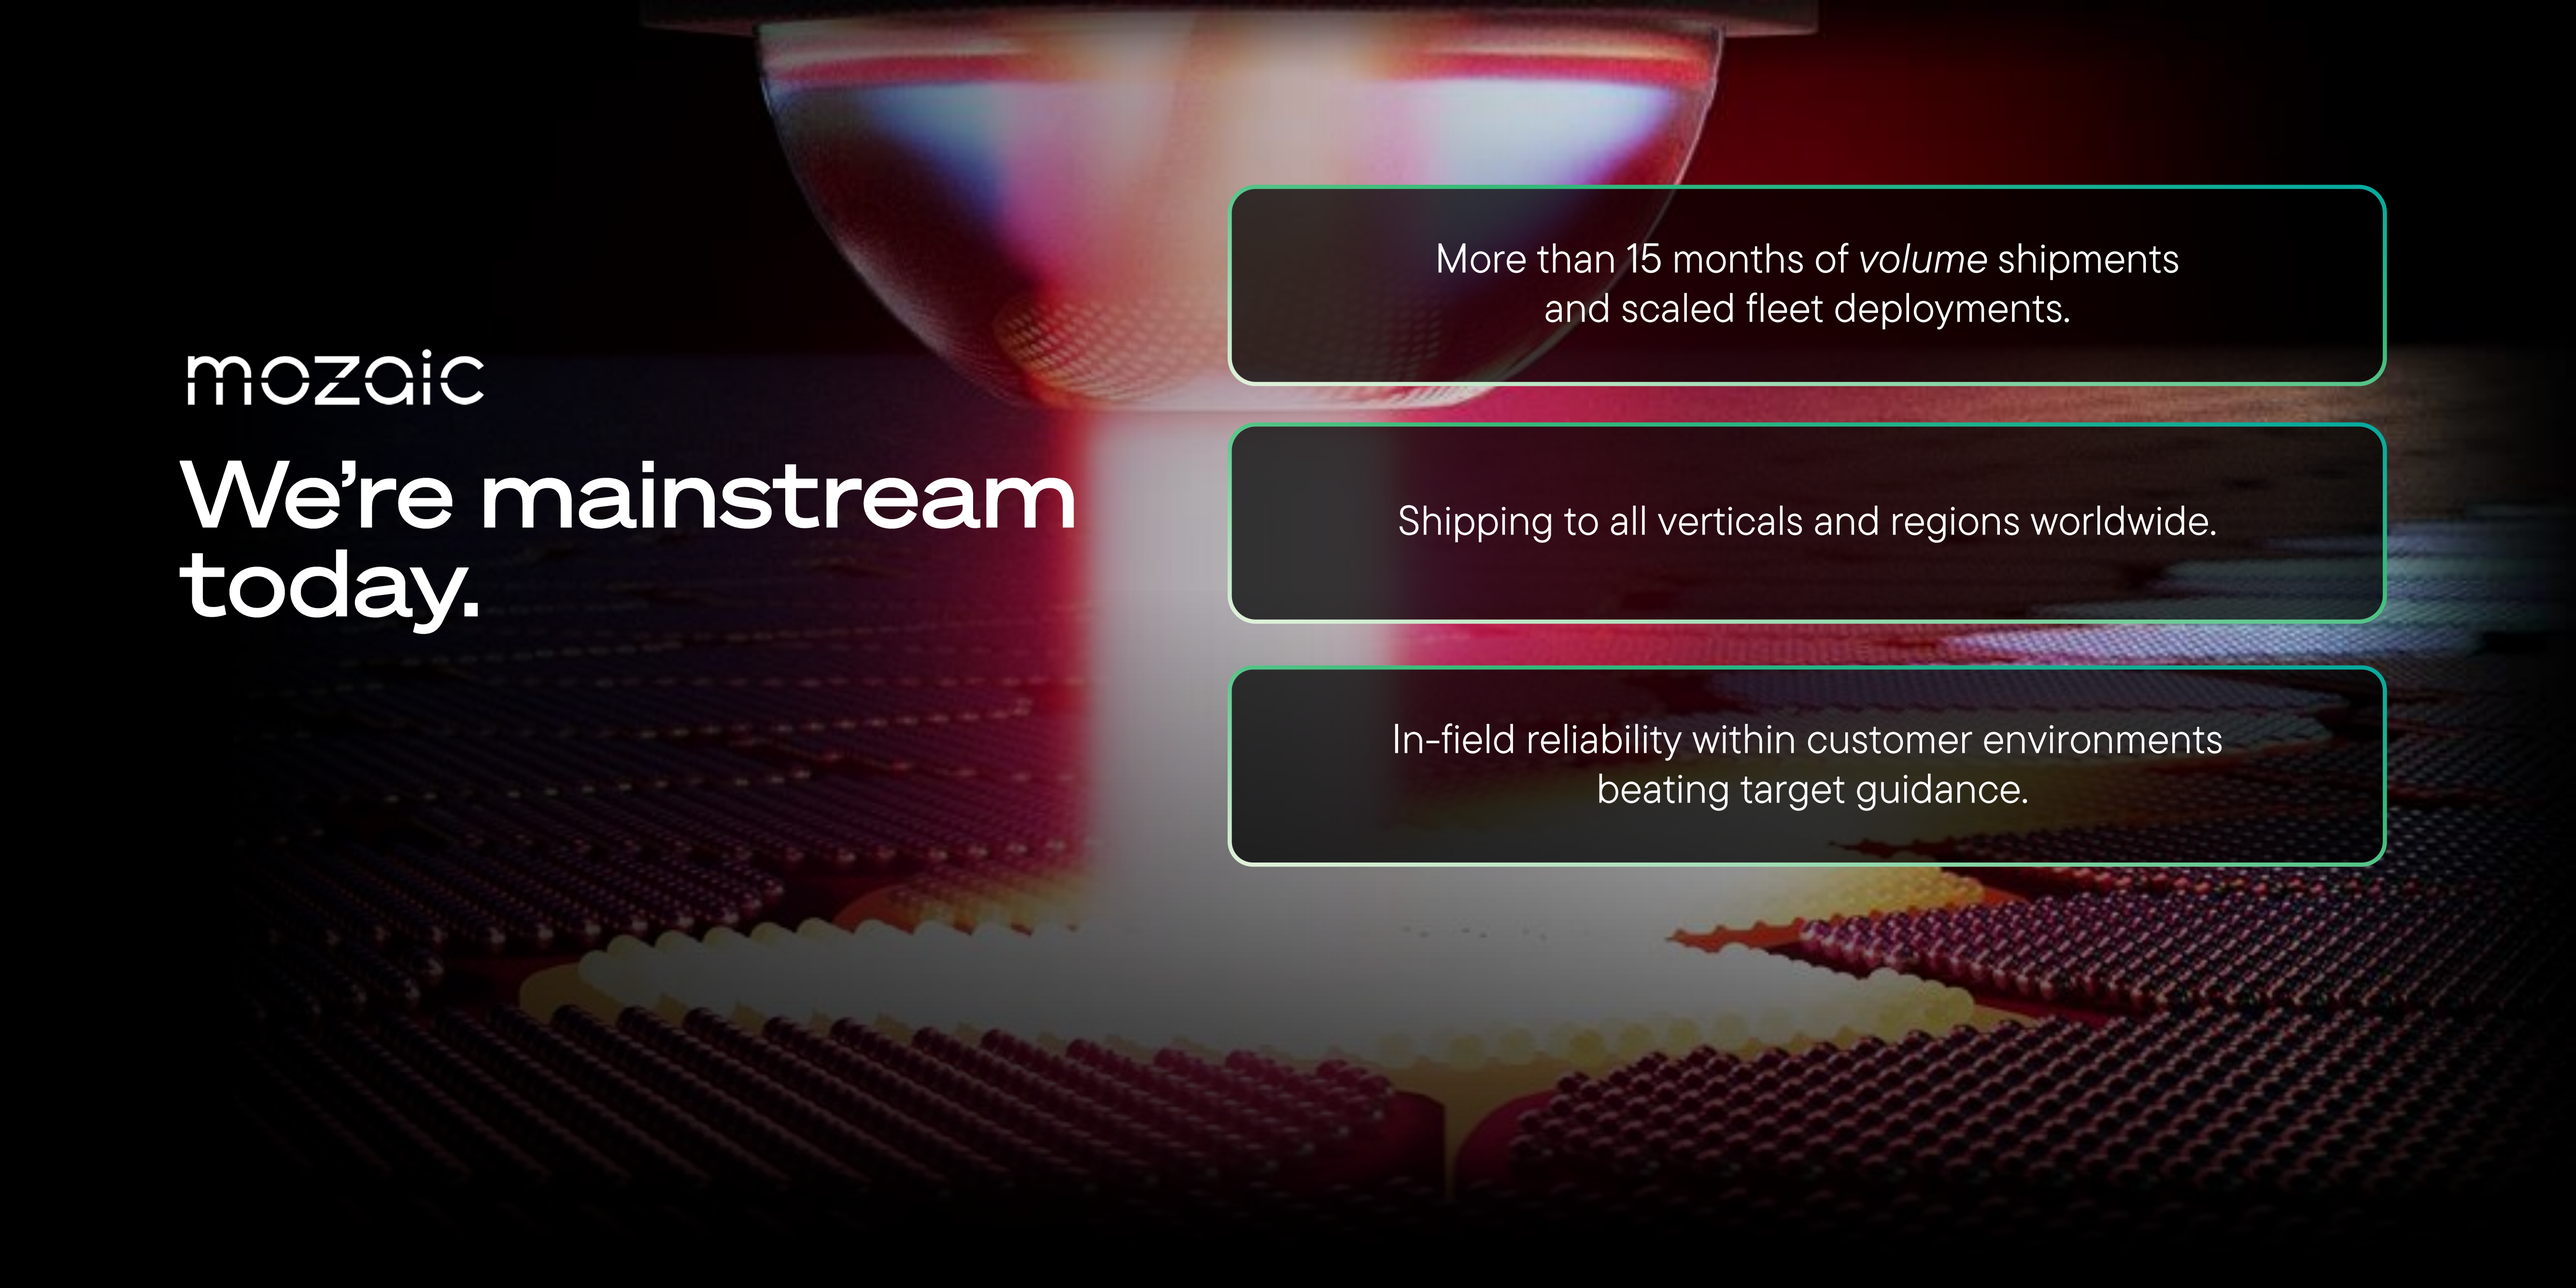
mozaic
We’re mainstream today.
More than 15 months of *volume* shipments and scaled fleet deployments.
Shipping to all verticals and regions worldwide.
In-field reliability within customer environments beating target guidance.

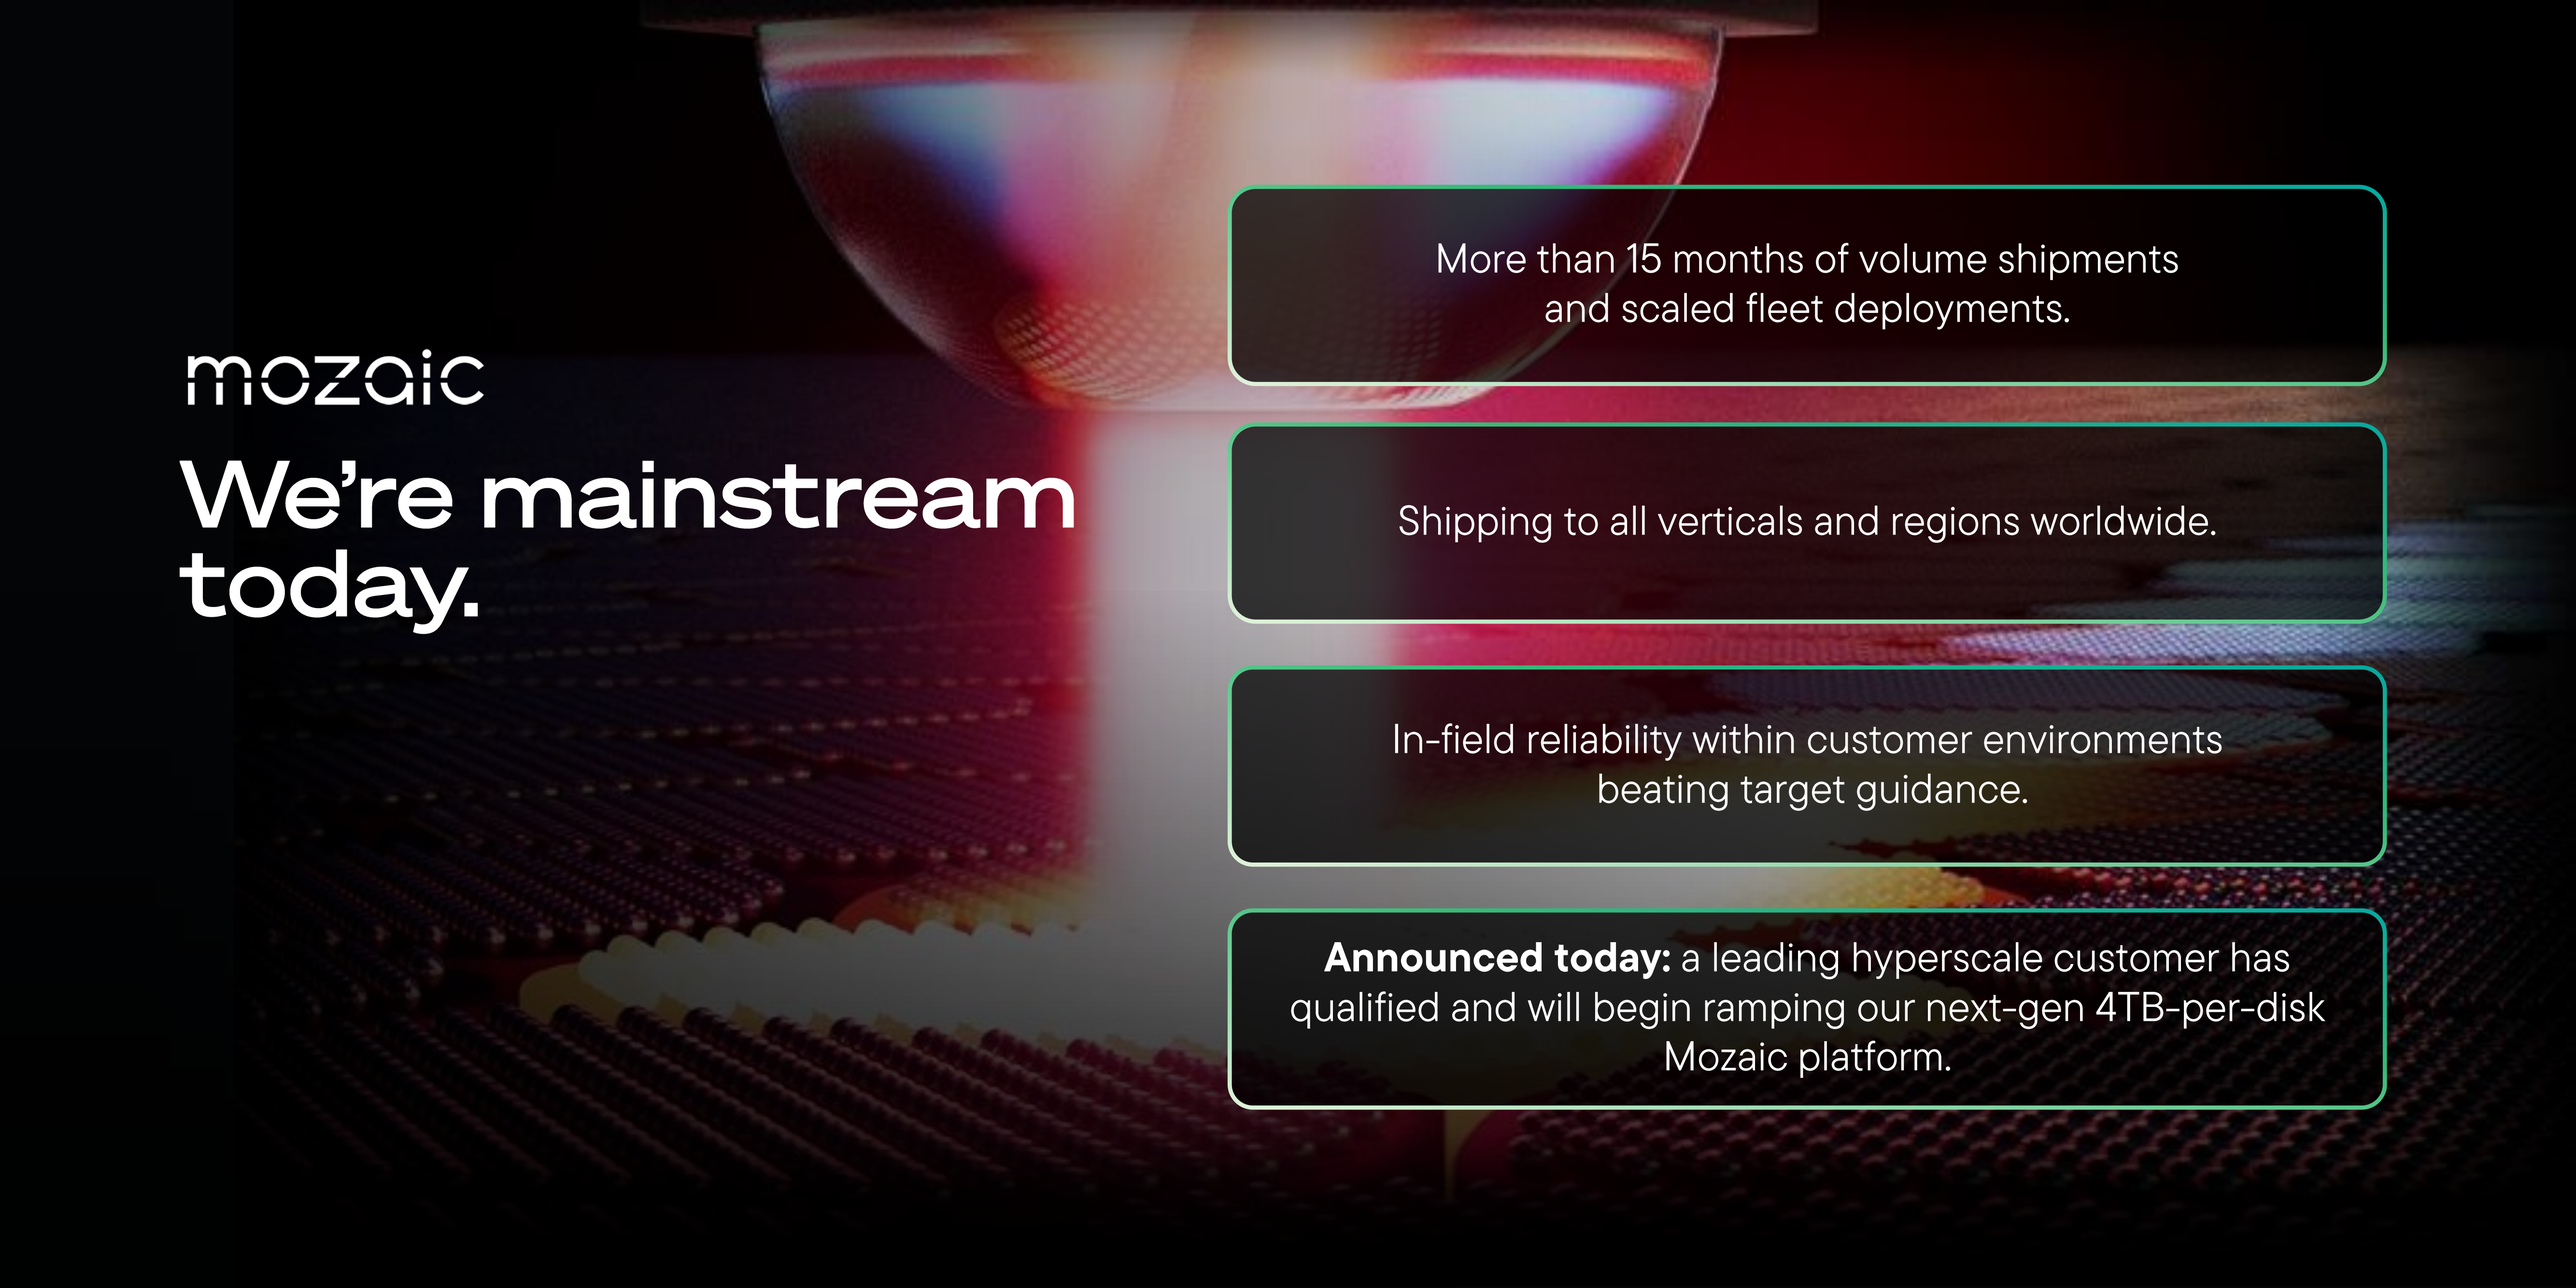
mozaic
We’re mainstream today.
More than 15 months of volume shipments and scaled fleet deployments.
Shipping to all verticals and regions worldwide.
In-field reliability within customer environments beating target guidance.
Announced today: a leading hyperscale customer has qualified and will begin ramping our next-gen 4TB-per-disk Mozaic platform.

Seagate stores infinite potential.