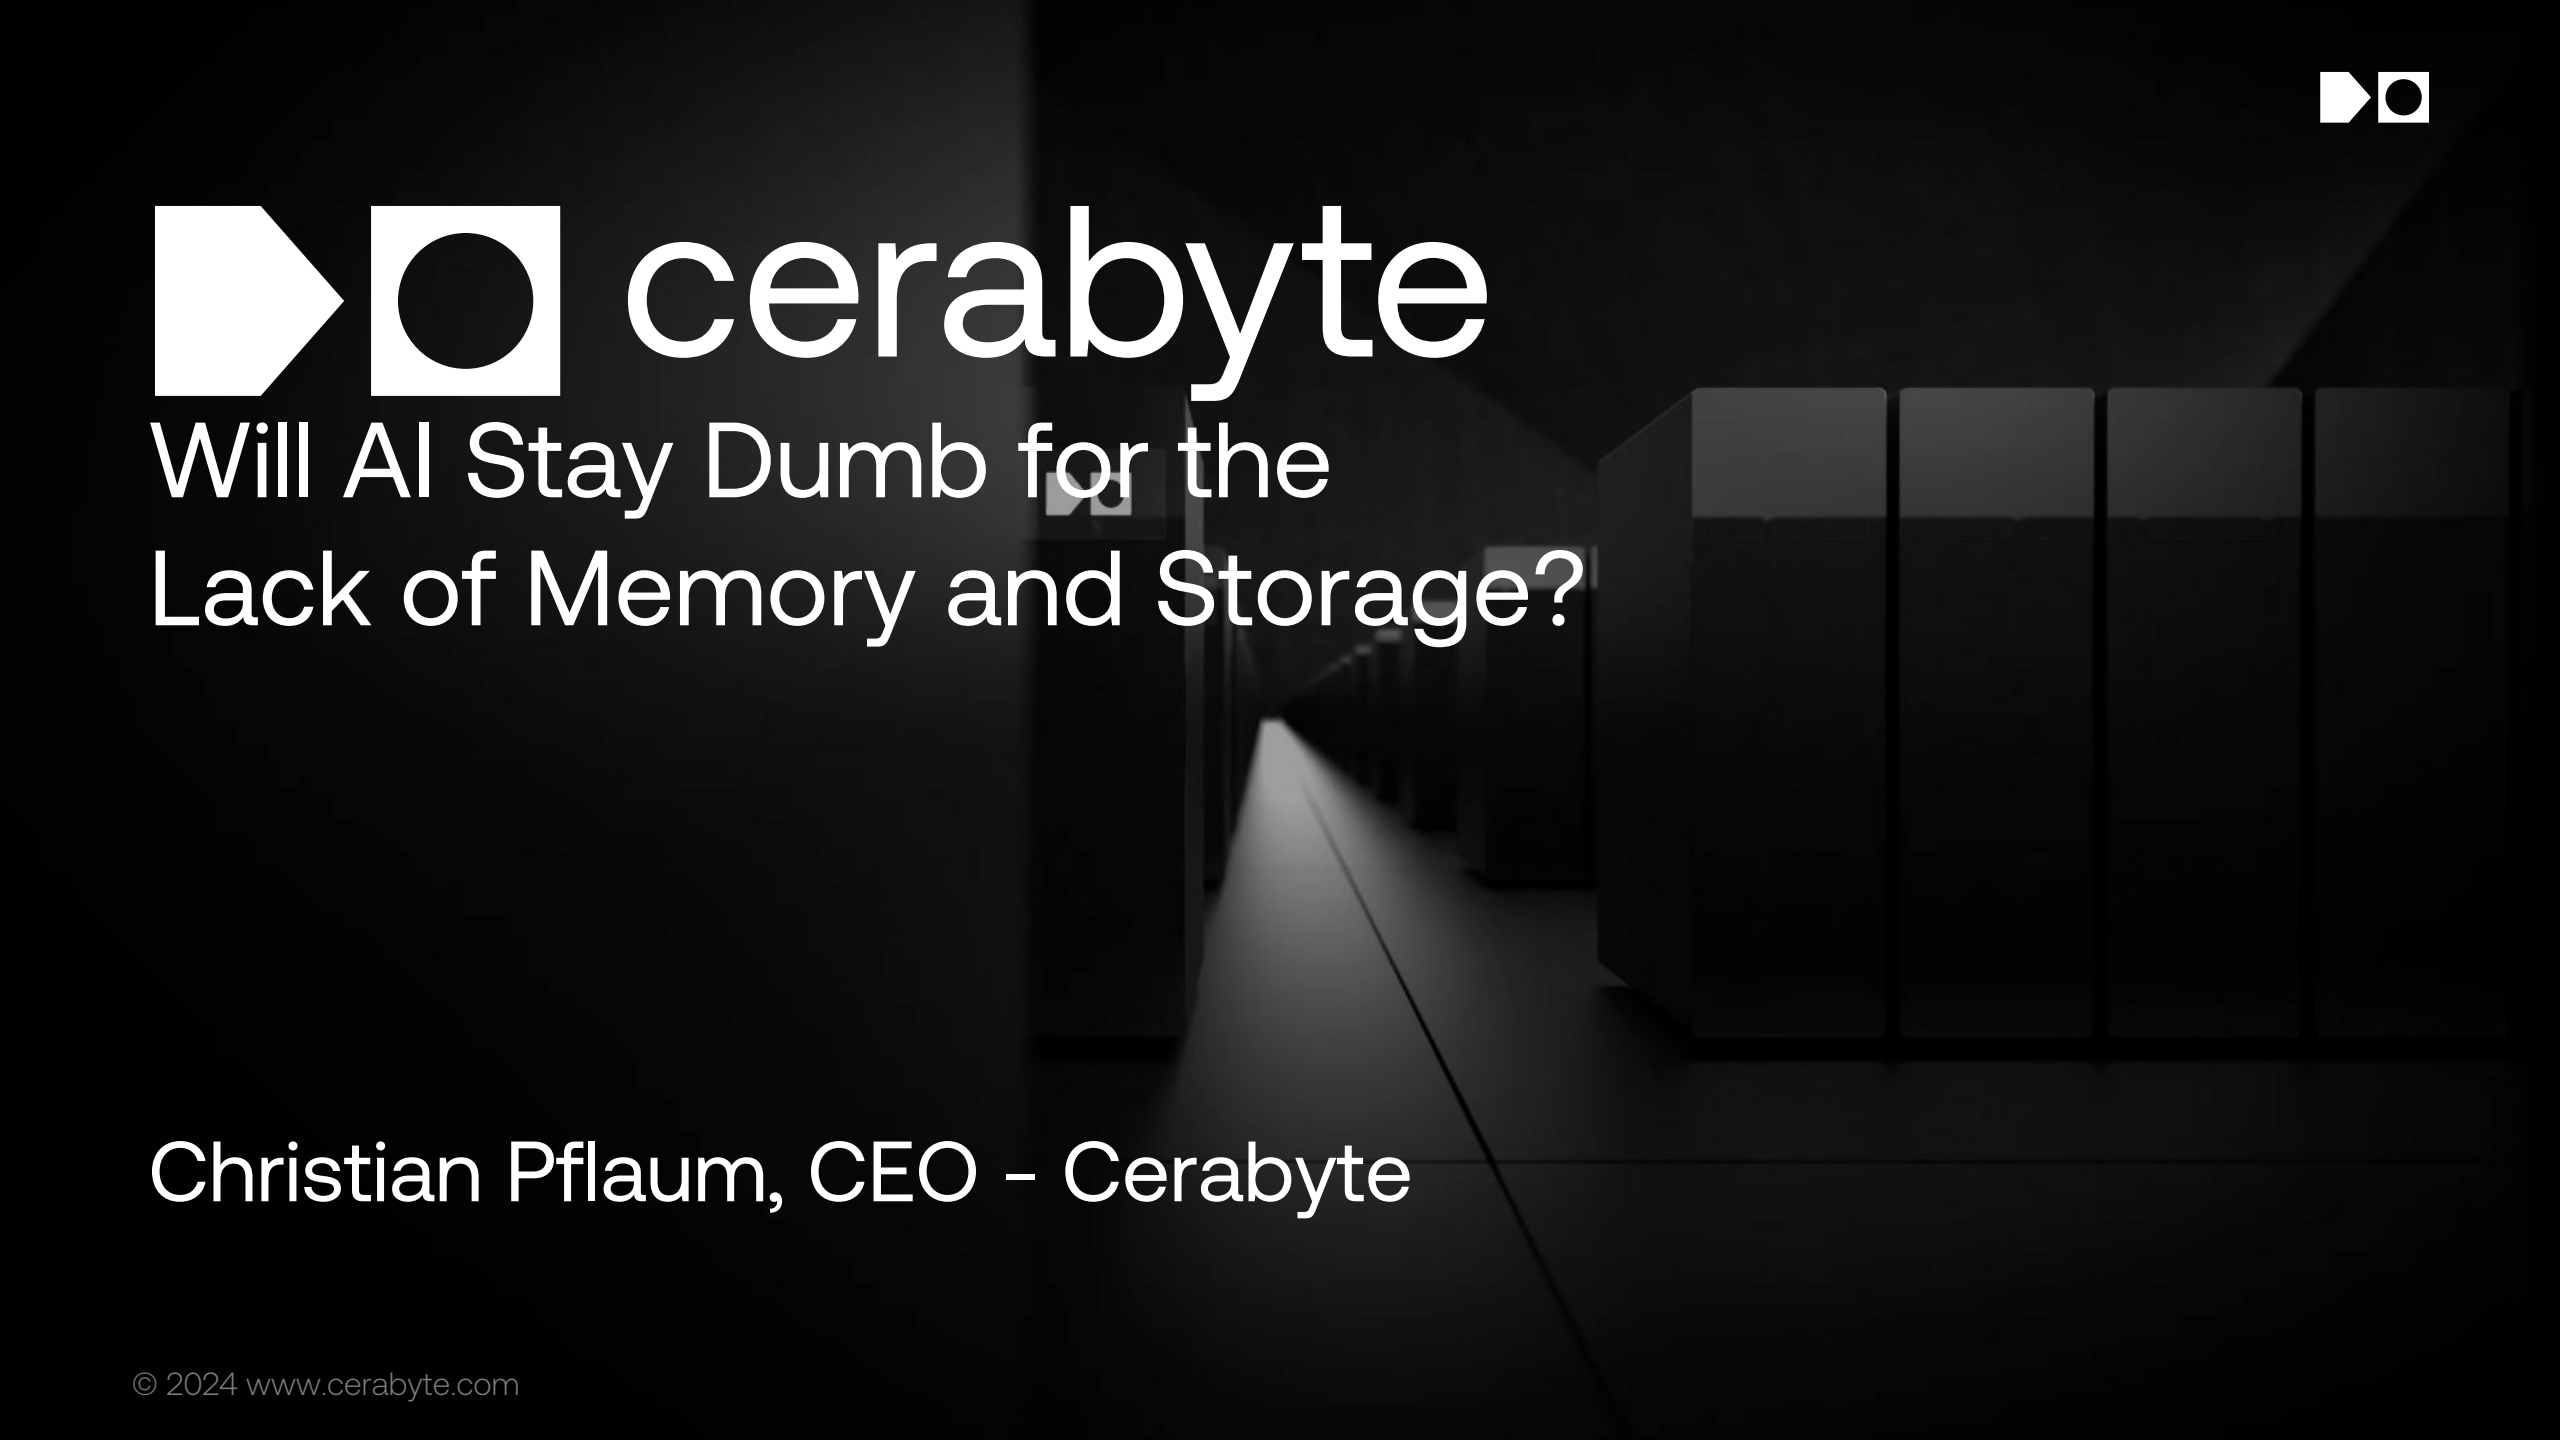

cerabyte

Will AI Stay Dumb for the Lack of Memory and Storage?

Christian Pflaum, CEO - Cerabyte

Cold data challenge

>70% of all data is **cold**
rarely retrieved after **90 days**
but stored for more than **a decade**

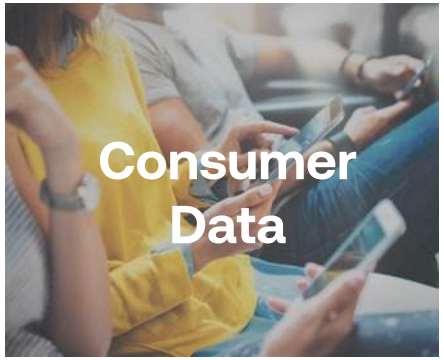

Cold data challenge

>70% of all data is **cold**
rarely retrieved after **90 days**
but stored for more than **a decade**

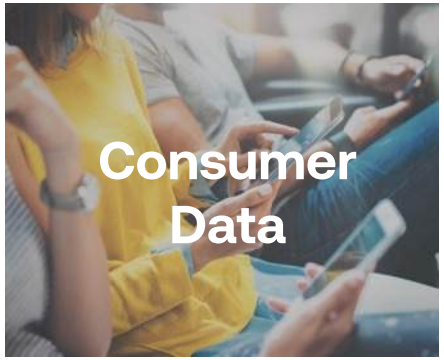

AI IMPACT

Will Stay AI dumb for lack of memory ?

Watch Video: https://vimeo.com/911929970/70e7294d04

Consumer 2030: 8 bn phones x 0.2 TB = 1.6 ZB for a lifetime

AI enhances smartphone experience

1. Camera Enhancements
2. Personalization
3. Voice Assistants
4. Predictive Text and Autocorrect
5. Health Monitoring
6. Augmented Reality (AR)
7. Virtual Reality (VR)
8. Natural Language Processing (NLP)

Source: AI generated content

Business 2030: 2.5 bn cars x 1 TB per year = 2.5 ZB per year

AI revolutionizes automotive

1. Autonomous Driving
2. Driver Assistance Systems
3. Predictive Maintenance
4. Natural Language Processing
5. Vehicle Design & Manufacturing
6. Personalized Driving Experience
7. Traffic Management:
8. Supply Chain Optimization

Source: AI generated content

Medicine 2030: 8.5 bn humans x 0.1 TB = 0.85 ZB for a lifetime

AI revolutionizes medicine & healthcare

1. Medical Imaging Analysis
2. Drug Discovery and Development
3. Personalized Medicine
4. Healthcare Operations
5. Remote Patient Monitoring
6. Medical Research and Clinical Trials
7. Healthcare Decision Support
8. Healthcare Robotics

Source: AI generated content

Energy 2030: $ 35 bn Oil&Gas Analytics (Oil&Gas = 8.5% GDP)

AI transforms oil & gas industry

1. Oil & Gas Exploration
2. Predictive Maintenance
3. Drilling Optimization
4. Reservoir Management
5. Production Optimization
6. Safety and Risk Management
7. Supply Chain Optimization
8. Environmental Monitoring

Source: AI generated content

Finance 2030: 8% of global GDP - IT spend 10% of revenues

AI impacts financial industry

1. Algorithmic Trading:
2. Risk Management
3. Fraud Detection and Prevention.
4. Customer Service and Chatbots
5. Credit Scoring and Underwriting
6. Portfolio Management
7. Regulatory Compliance
8. Quantitative Analysis

Source: AI generated content

Science 2030: Enhancing productivity (Science = 2% global GDP)

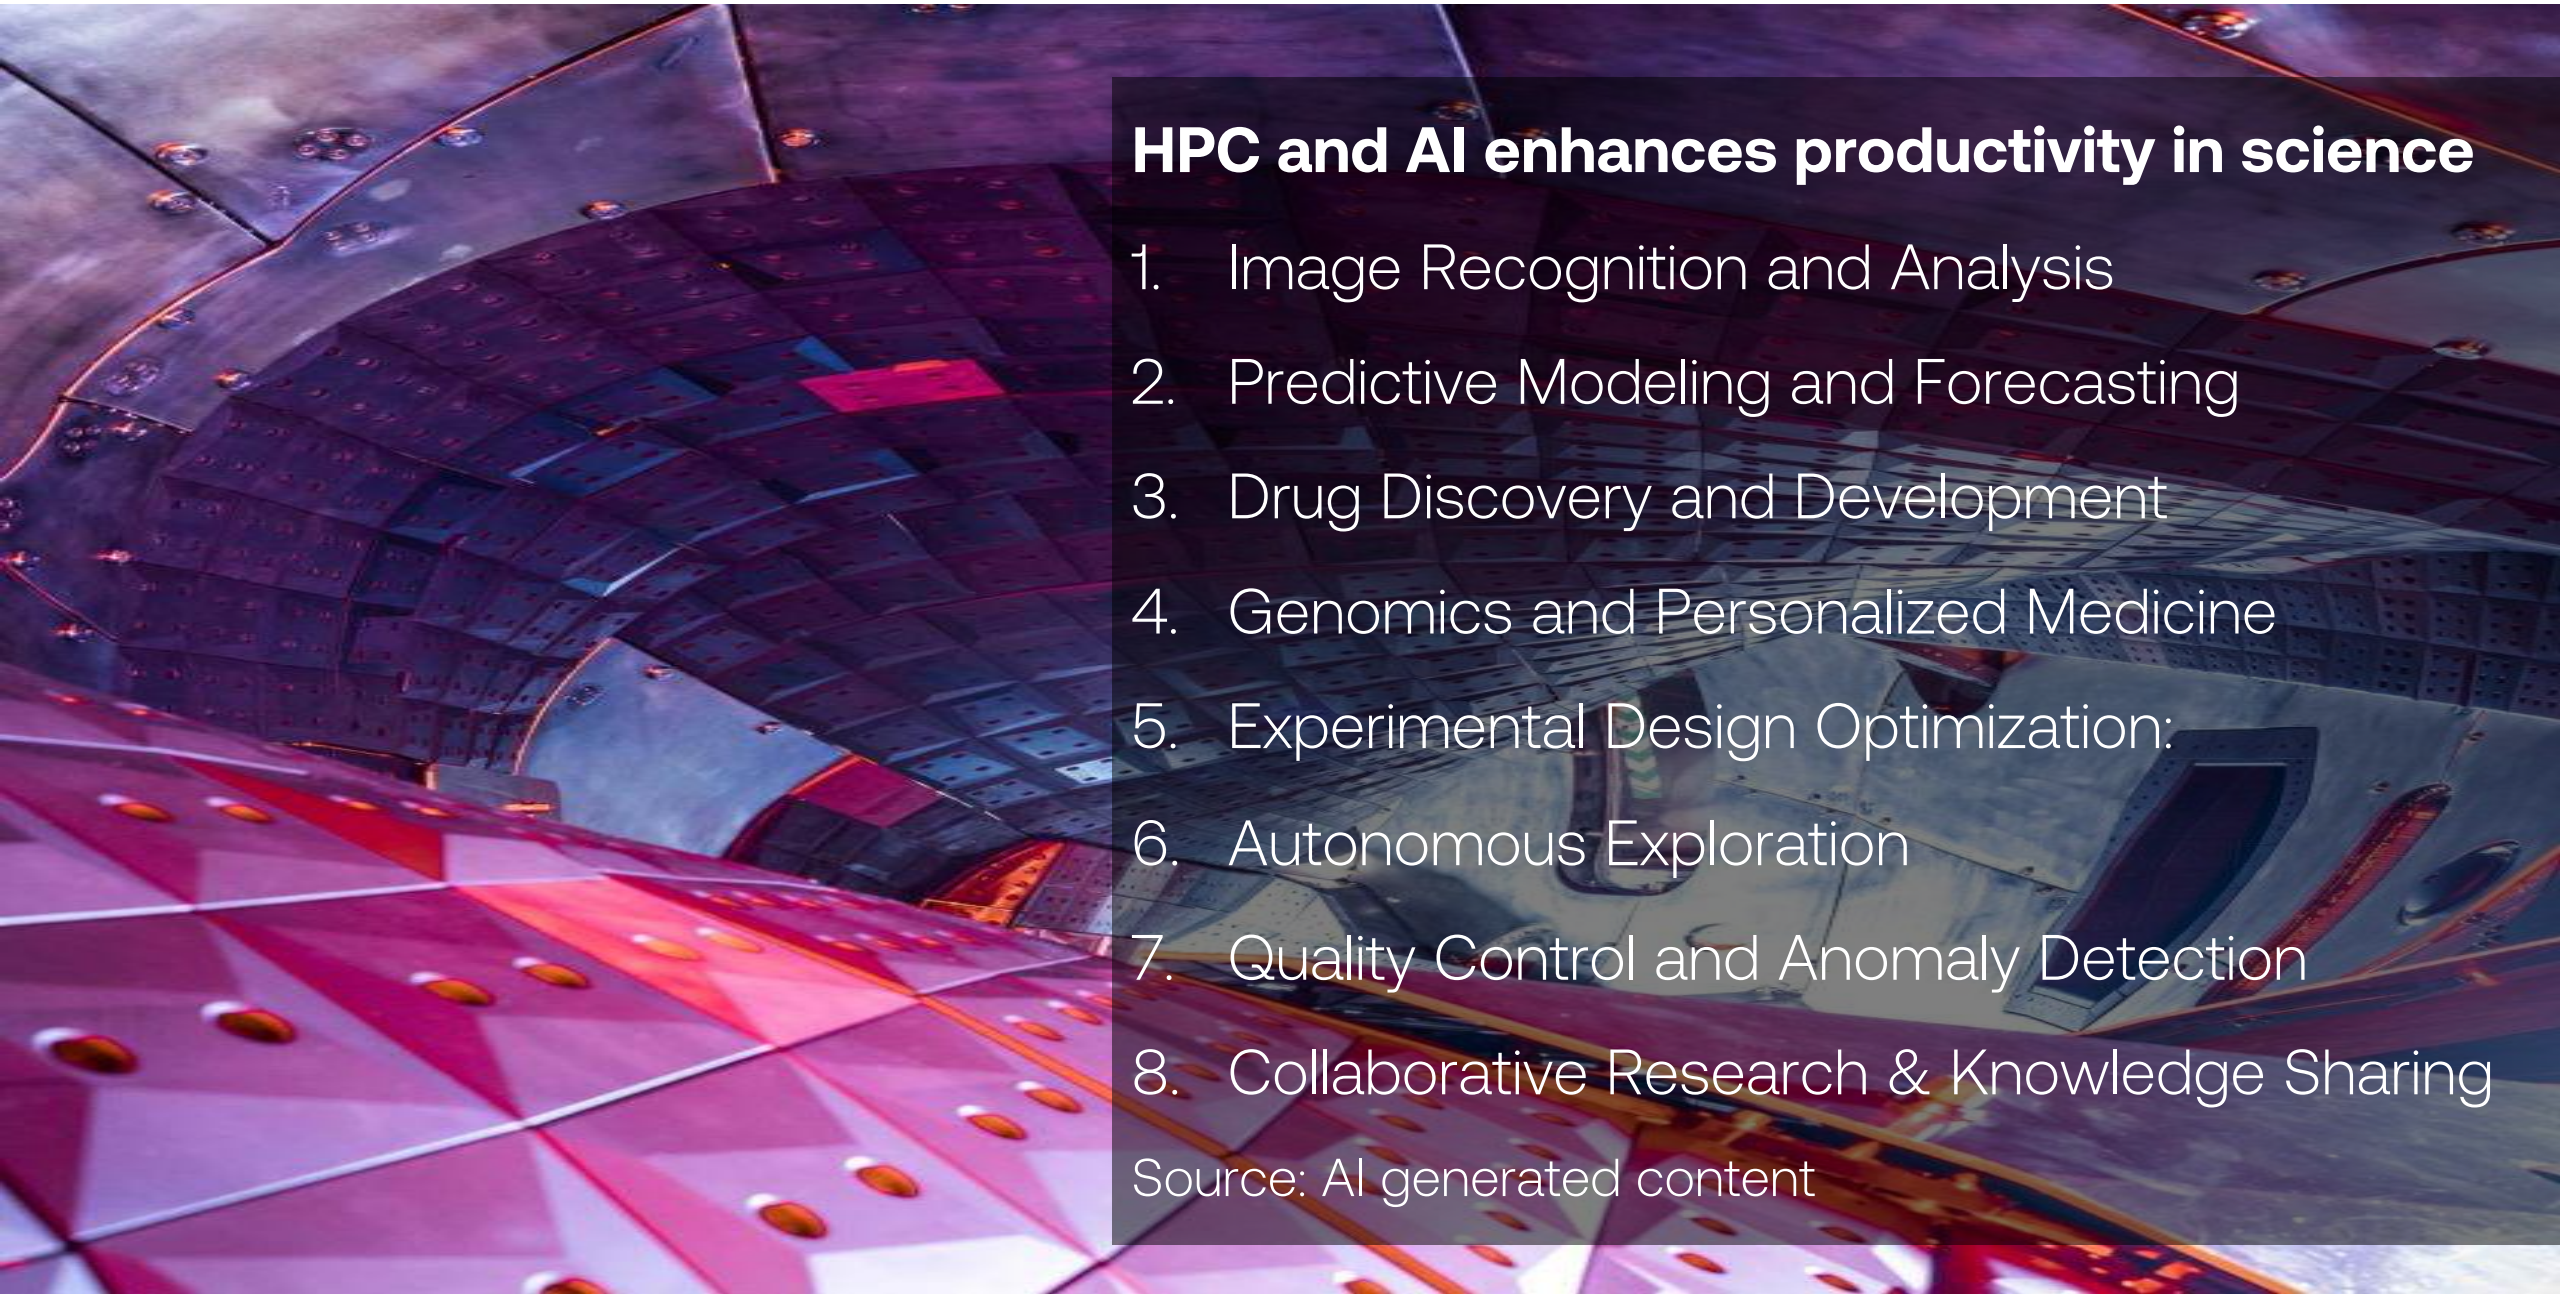

HPC and AI enhances productivity in science

1. Image Recognition and Analysis
2. Predictive Modeling and Forecasting
3. Drug Discovery and Development
4. Genomics and Personalized Medicine
5. Experimental Design Optimization:
6. Autonomous Exploration
7. Quality Control and Anomaly Detection
8. Collaborative Research & Knowledge Sharing

Source: AI generated content

Entertainment 2030: 2,5% of global GDP - 80% digital

AI revolutionize Entertainment

1. Content Creation
2. Content Recommendation
3. Predictive Analytics
4. Enhanced Visual Effects
5 .Personalized Marketing
6. Virtual Performers
7. Content Moderation
8 AI Influencers

Source: AI generated content

Superhuman AI-Cluster double compute every few months

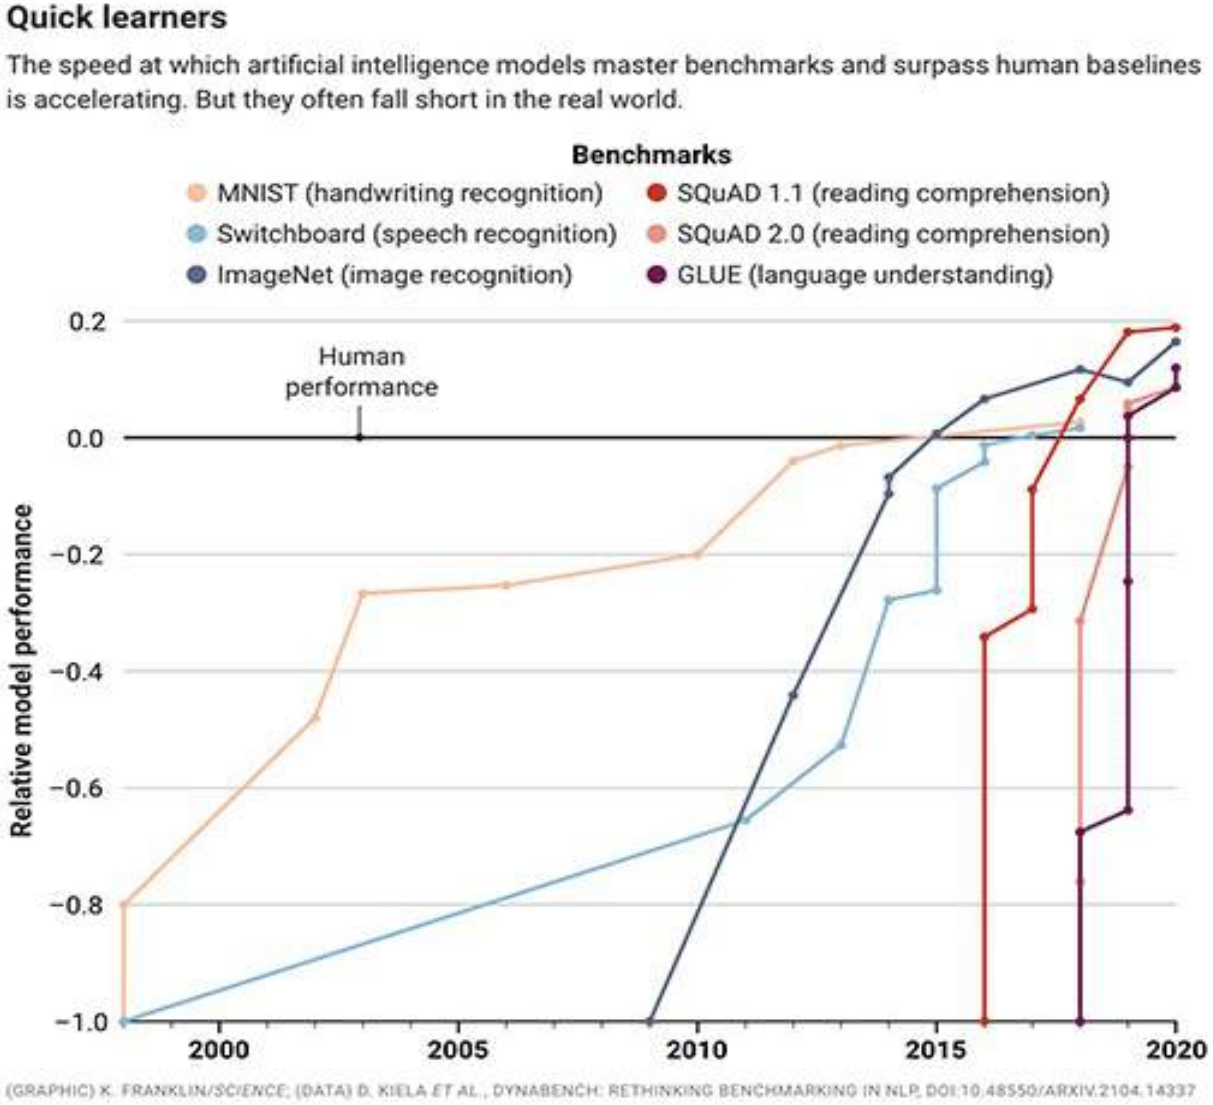

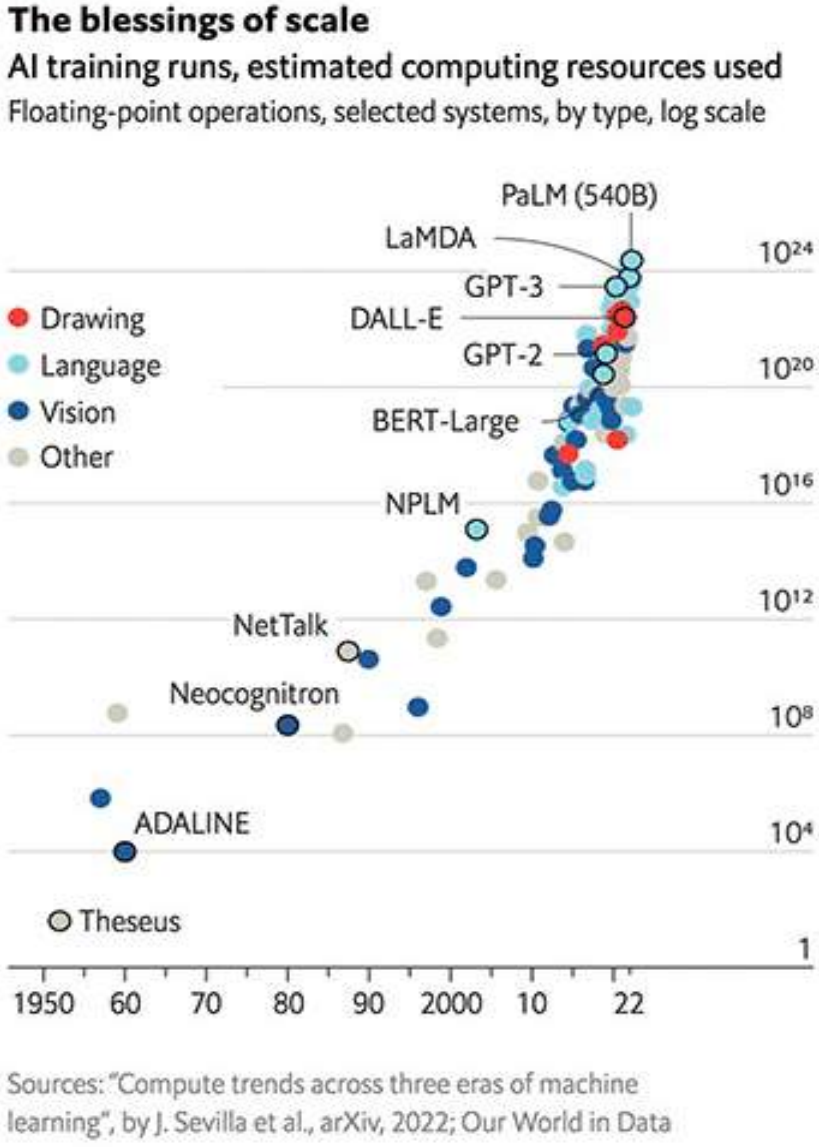

AI - Impact

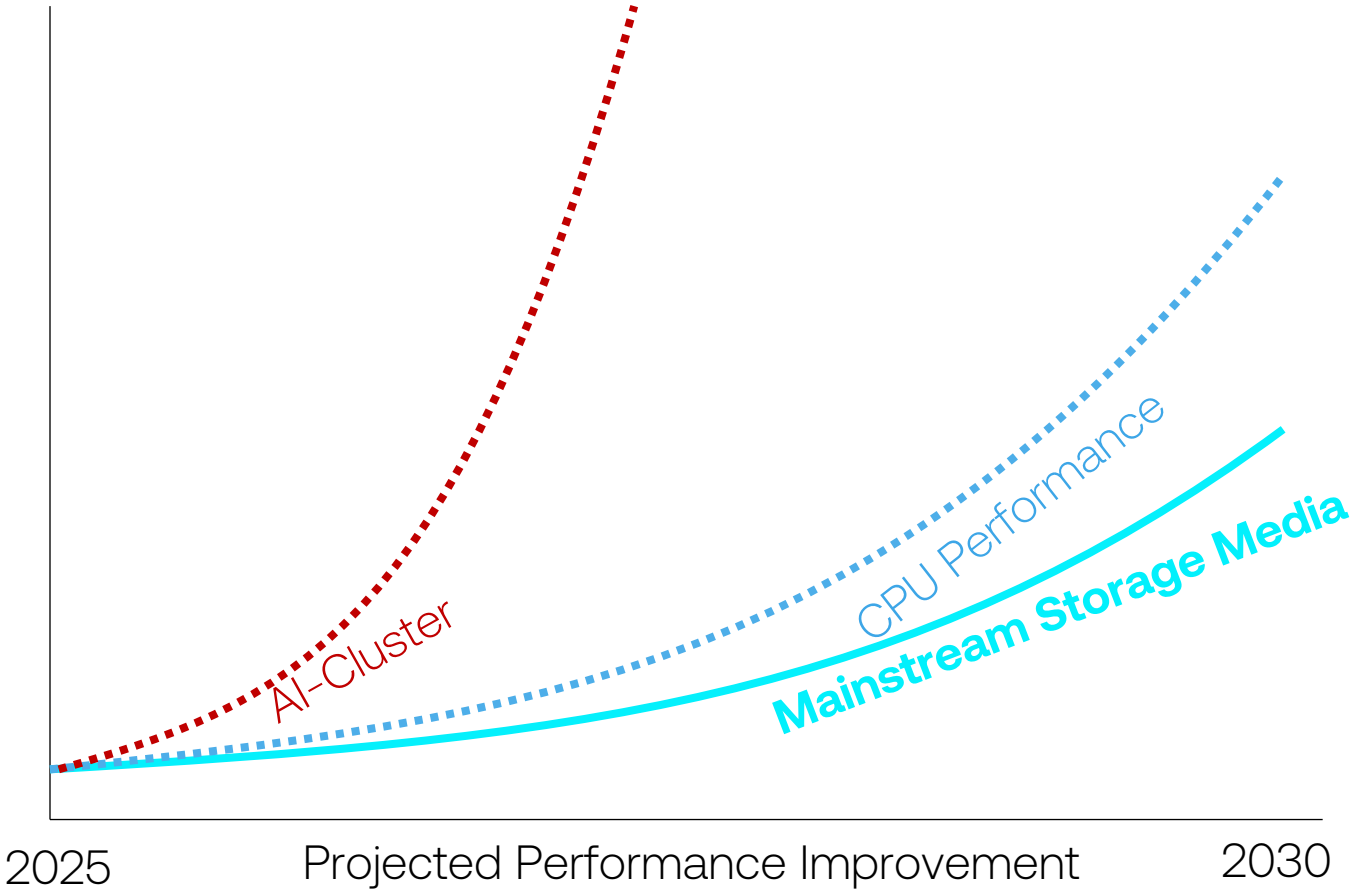

AI-Cluster performance improvement outpaces CPU and mainstream storage media

New S - Curve

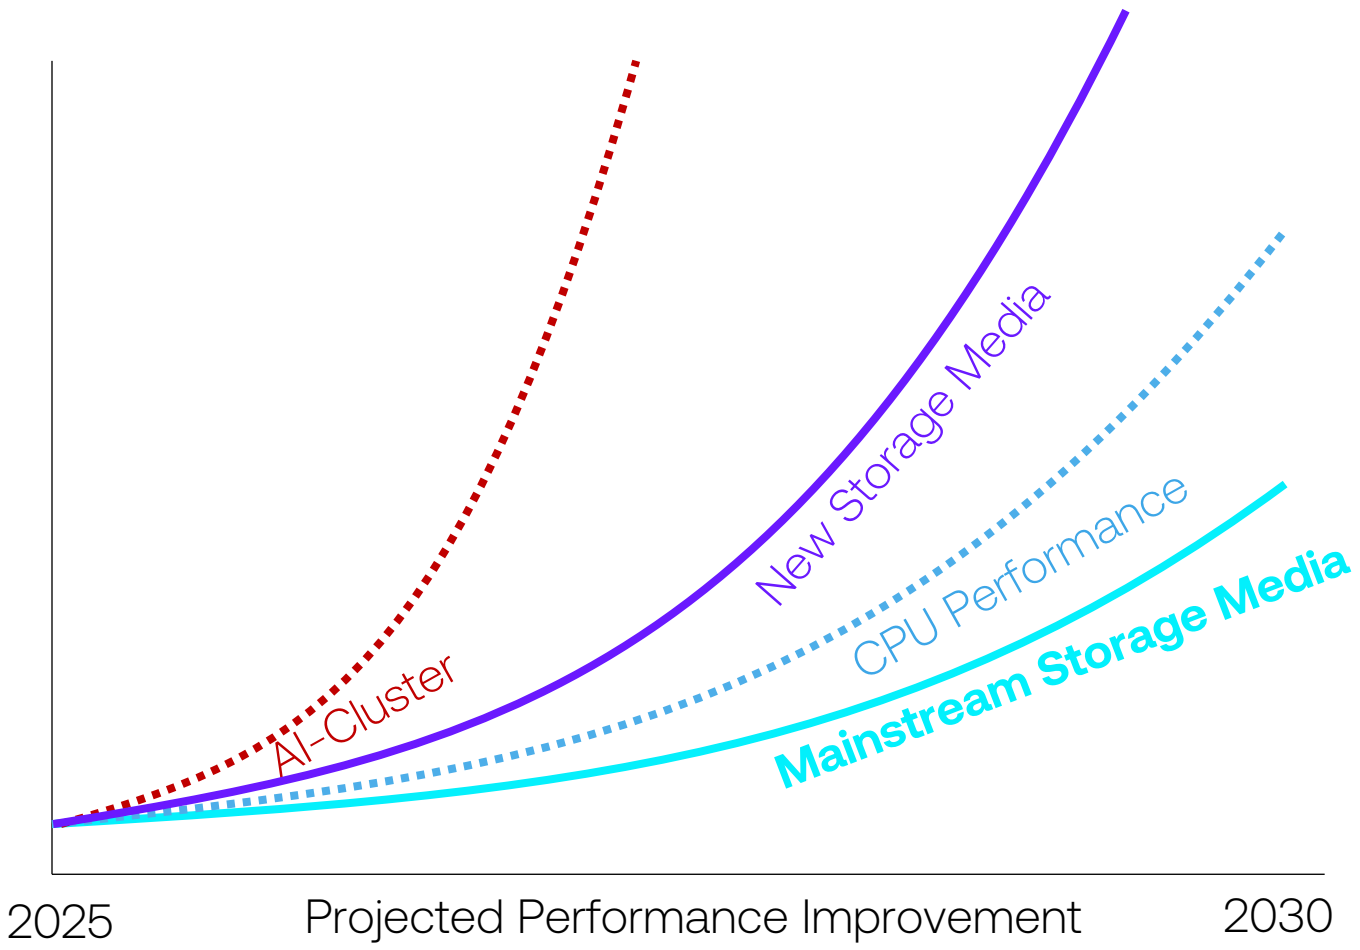

New Storage Media required for cold data storage & enabling a new technology S-curve

Development of AI

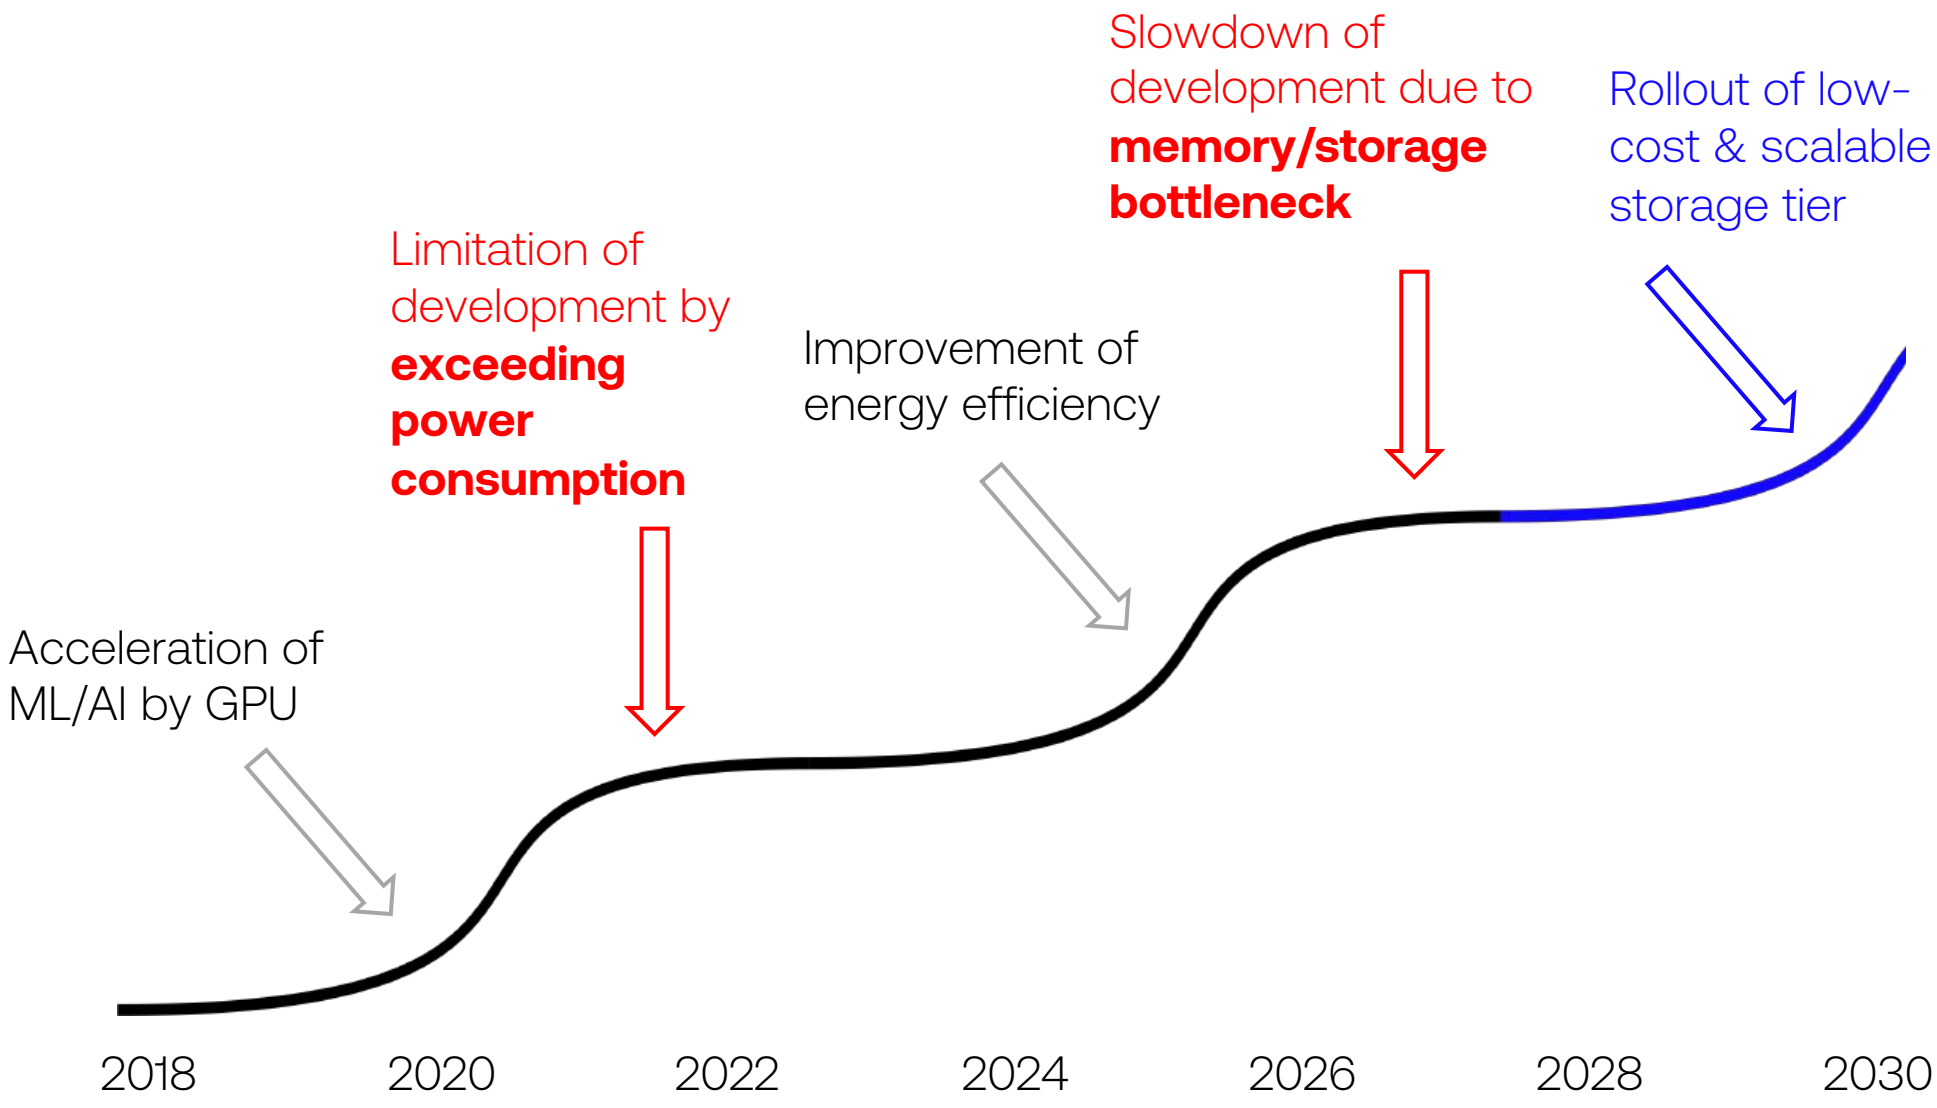

Global Installed Data Storage Base & Forecast

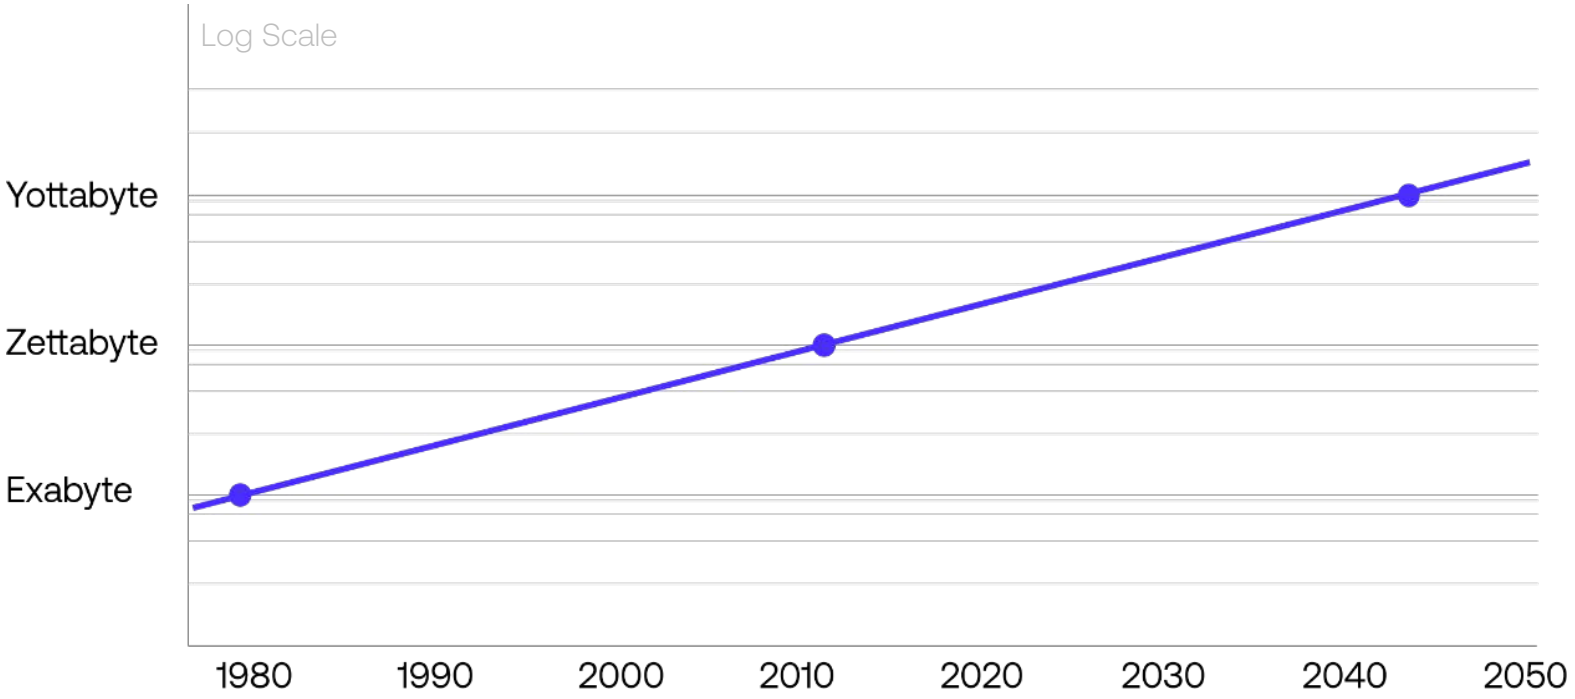

Source: Horison Information Strategies

Every ~30 years:

1000x growth

Requirements for AI Storage Dataverse in the next 20 years

Storage Density

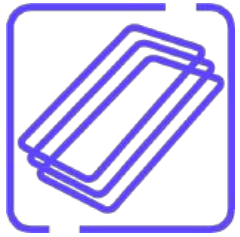

100x

Growth

Energy Efficiency

100x

Improvement

Cost Scaling

$1/PB

per month

Emerging Technologies - Accessible Permanent Storage Tier

Durable, Fast Write/Read, Low Latency, Low Cost

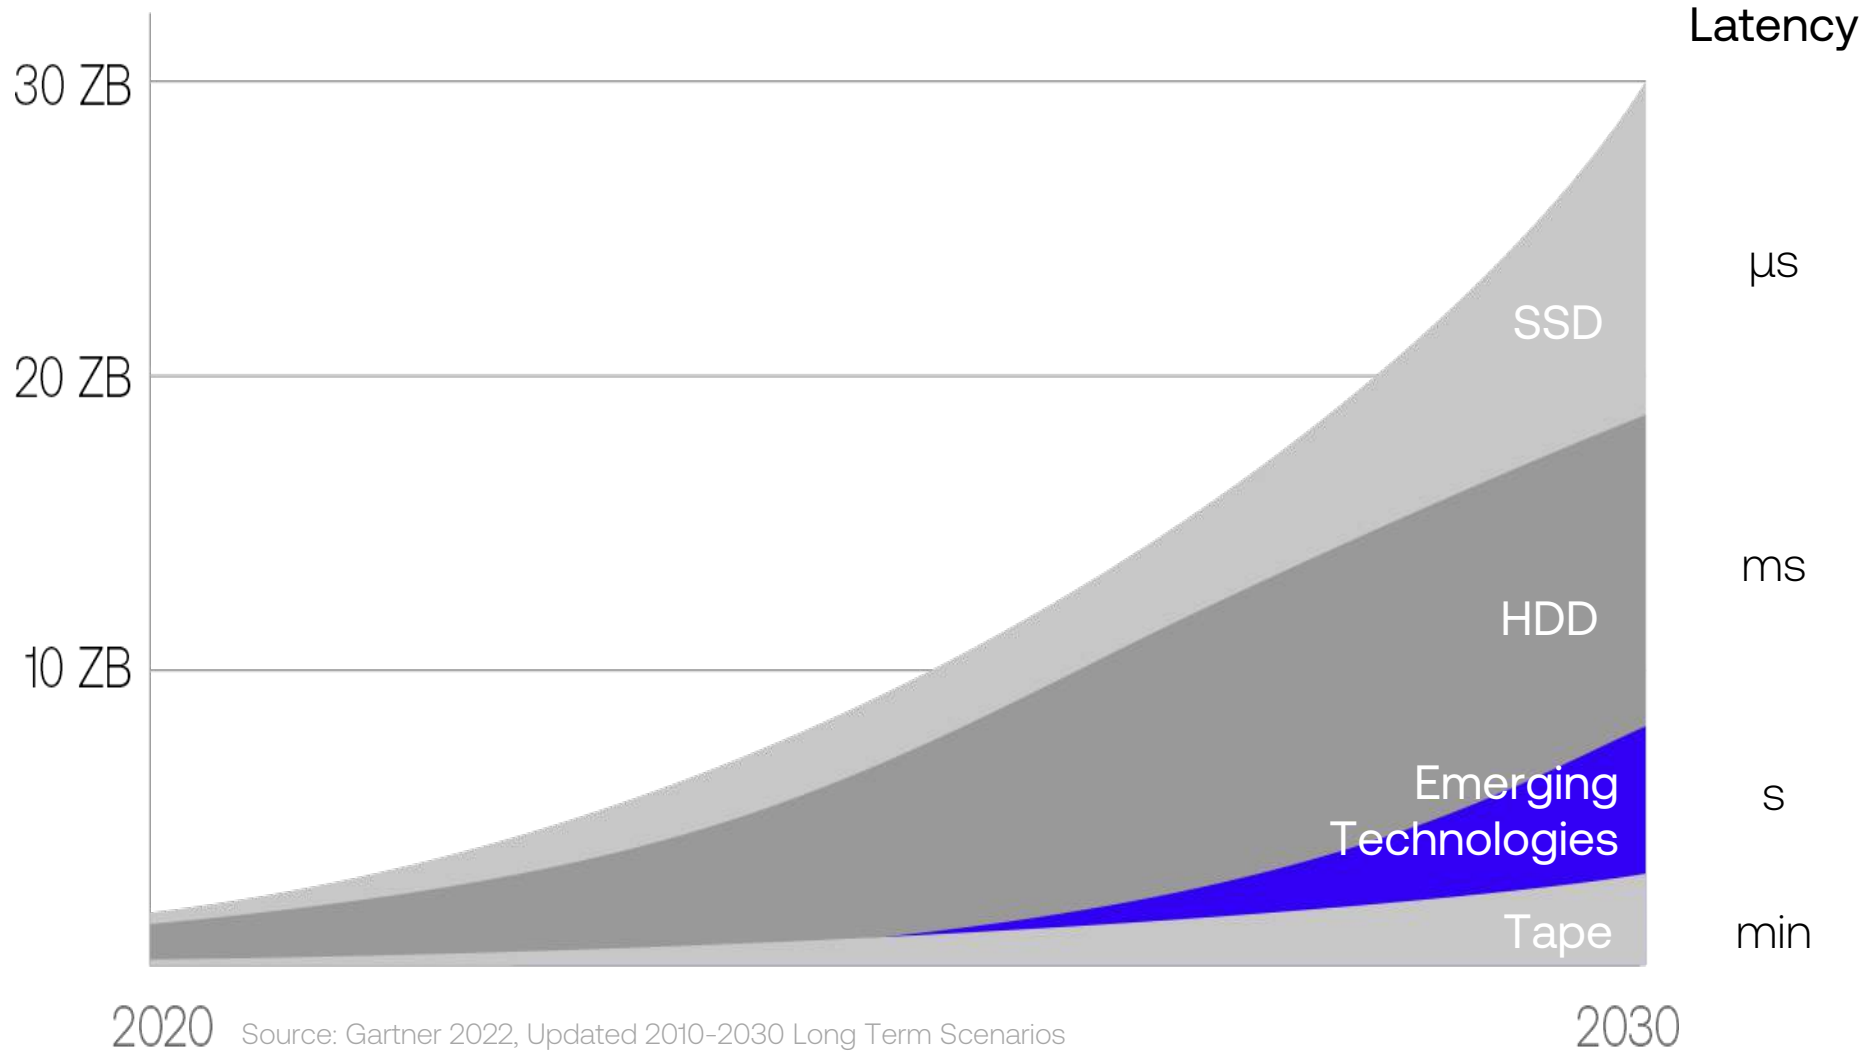

Source: Gartner 2022, Updated 2010-2030 Long Term Scenarios

Vision to Reality

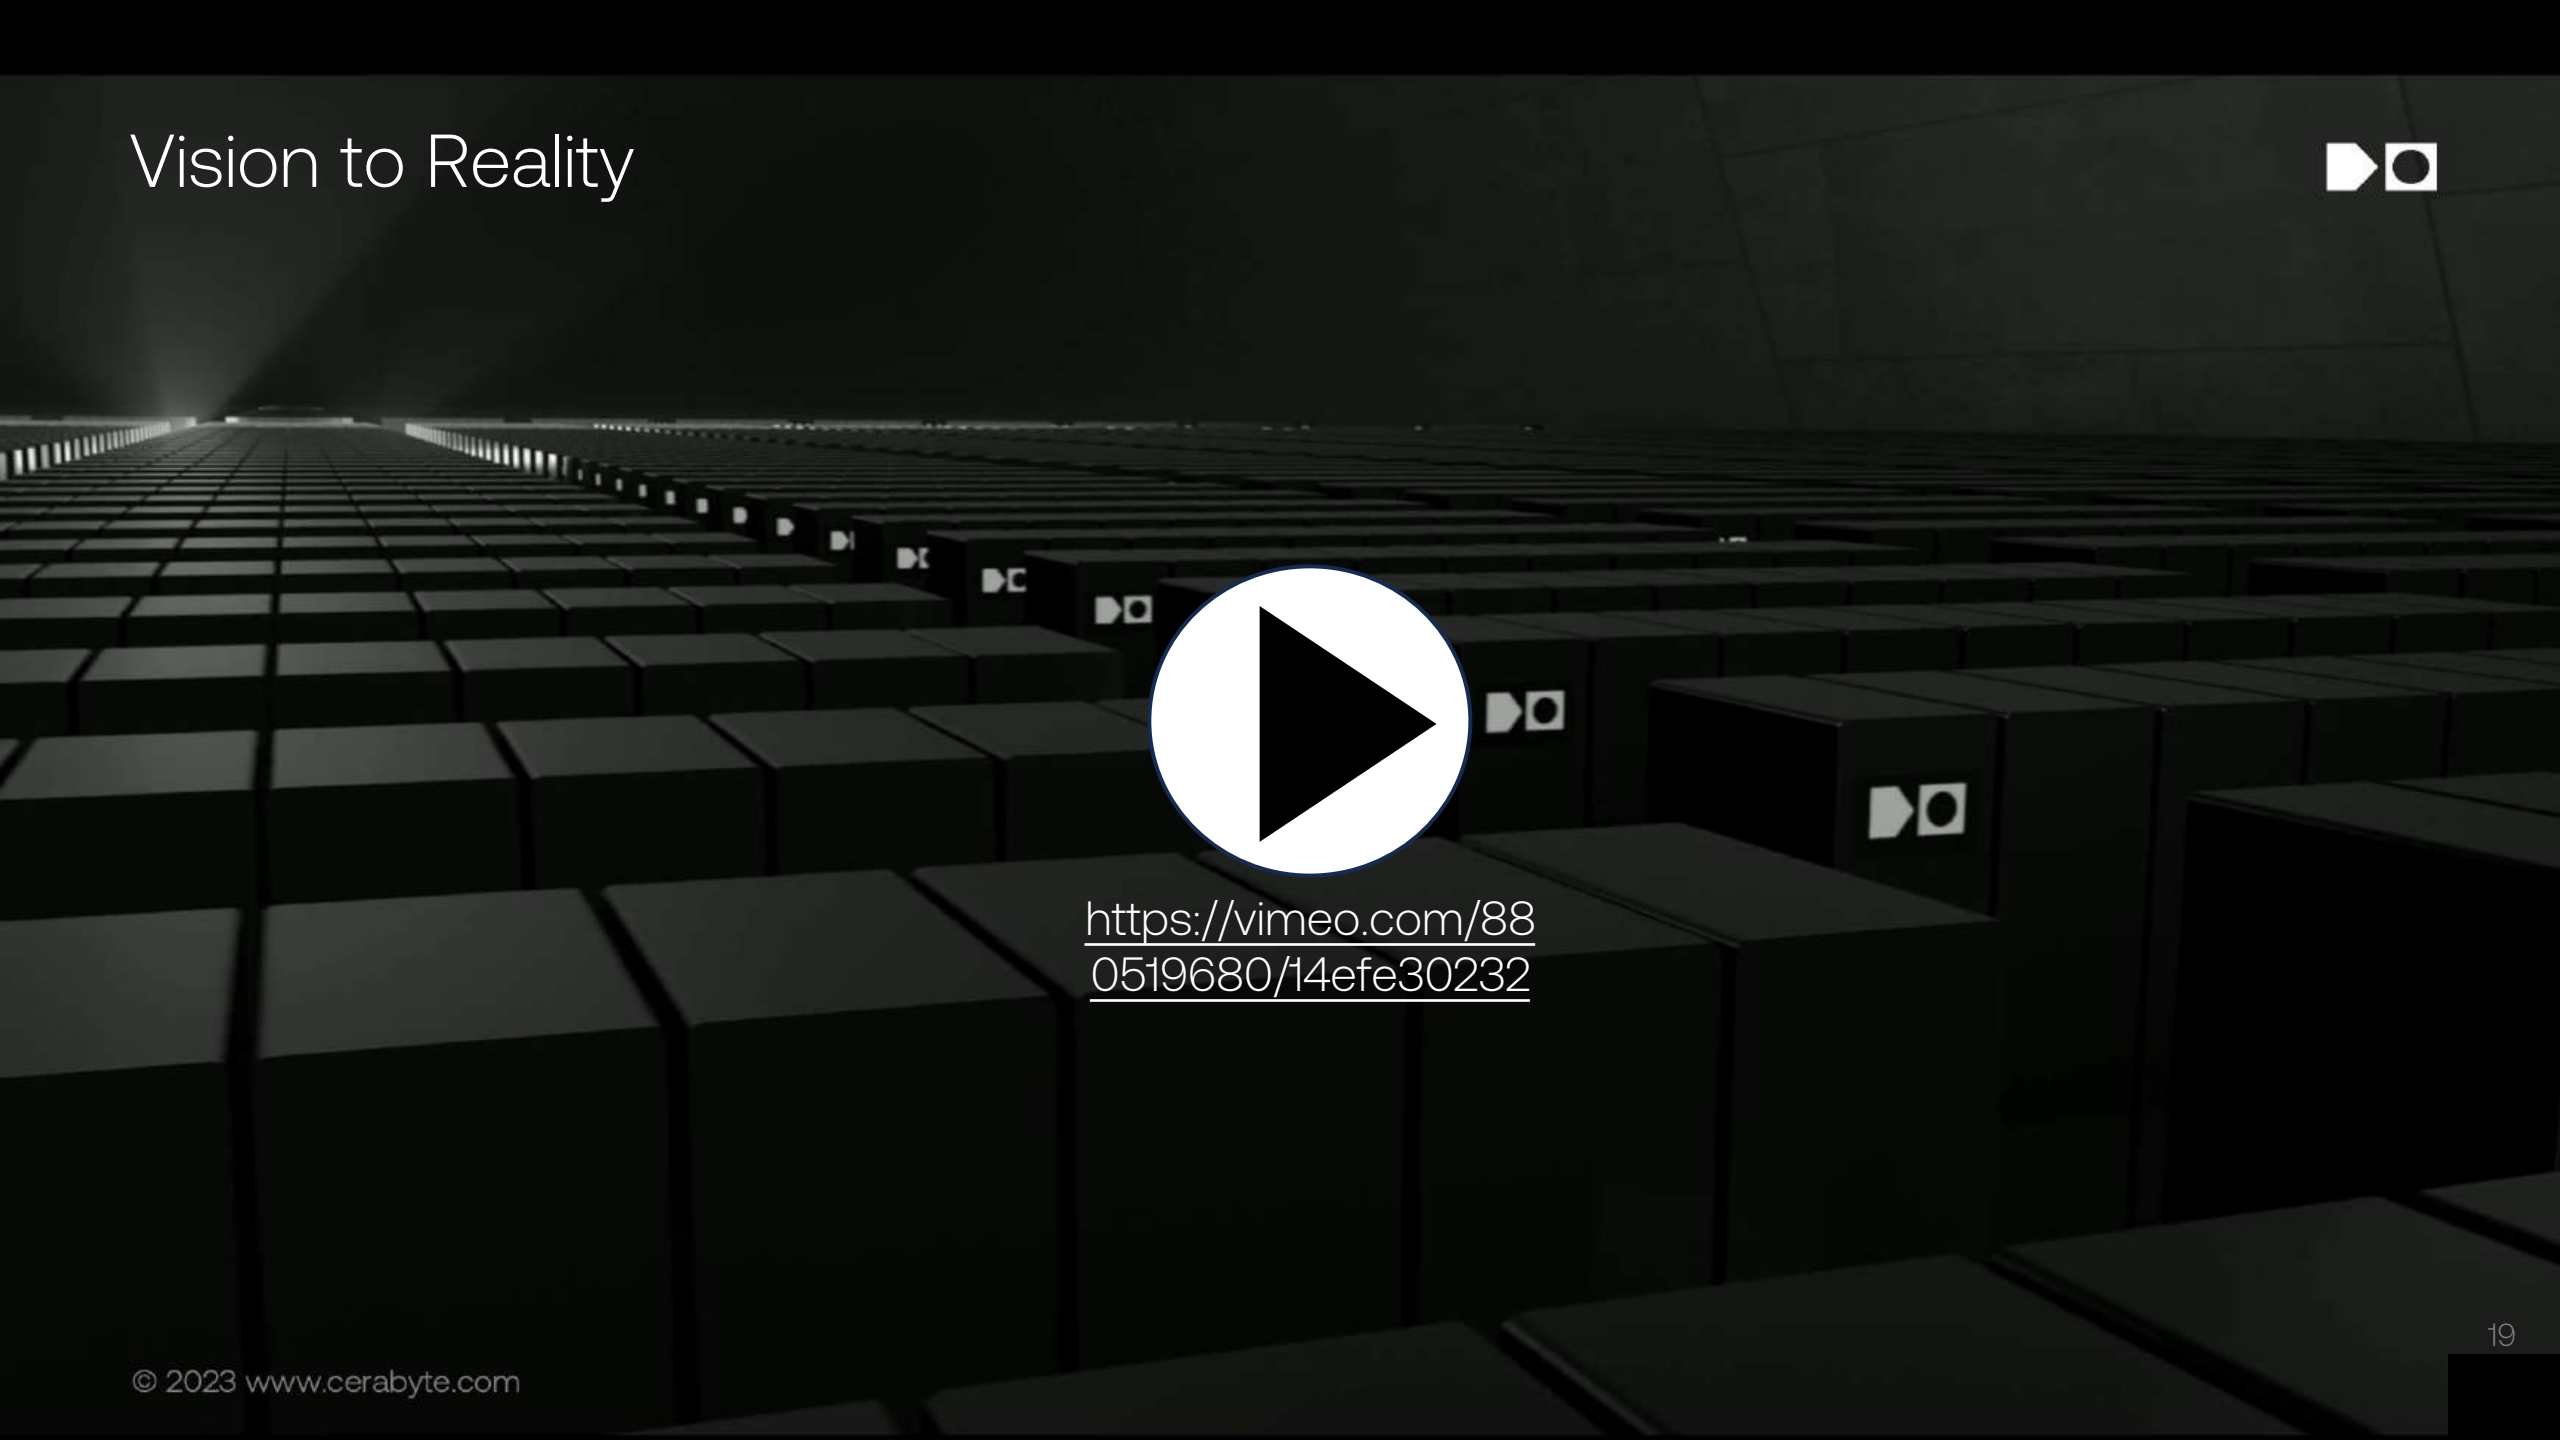

https://vimeo.com/880519680/14efe30232

Principle

Ceramic-on-glass Punch Cards with Laser Nanoscale Matrices

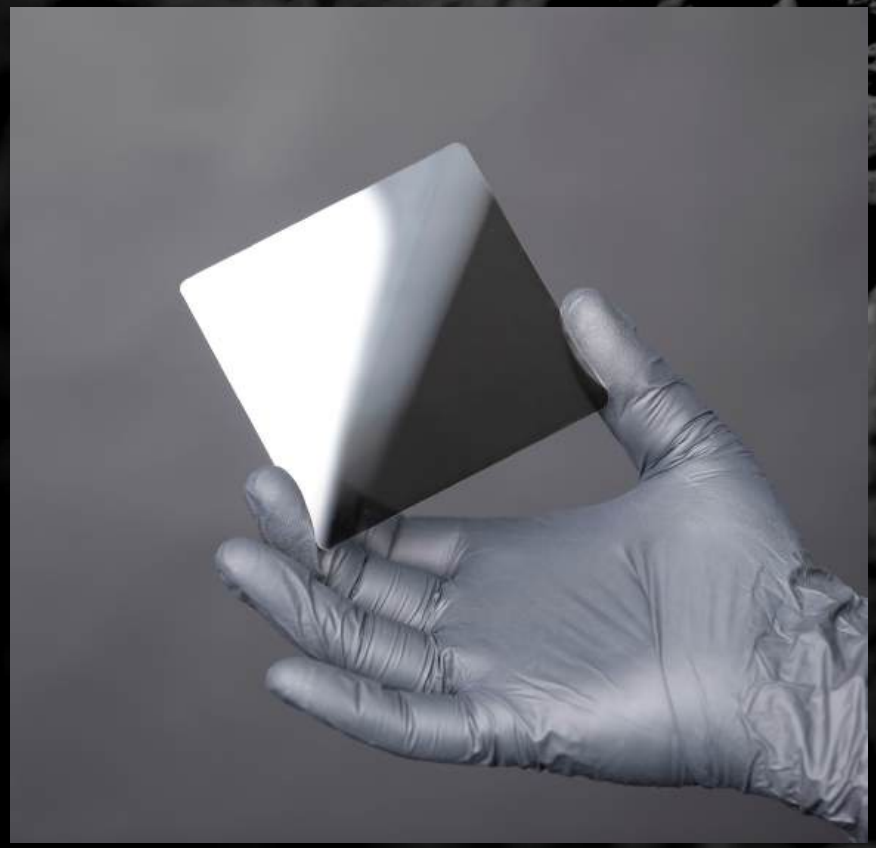

+

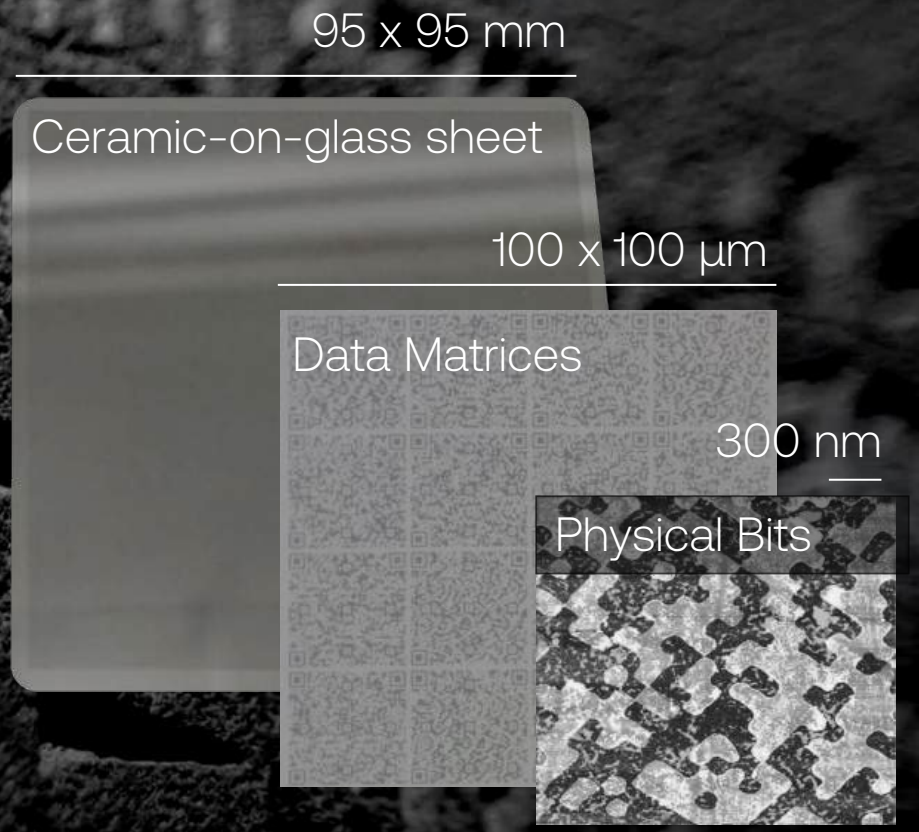

Ceramic-on-glass sheets

Laser nanoscale matrices

Summary

Display glass, ceramic coating, laser, and standard industry robotics

Storage media stacked in cartridges

End-to-end system using existing robotic libraries

Watch prototype Video: https://vimeo.com/880519680/14efe30232

Squared glass carrier with dark ceramic nano coating

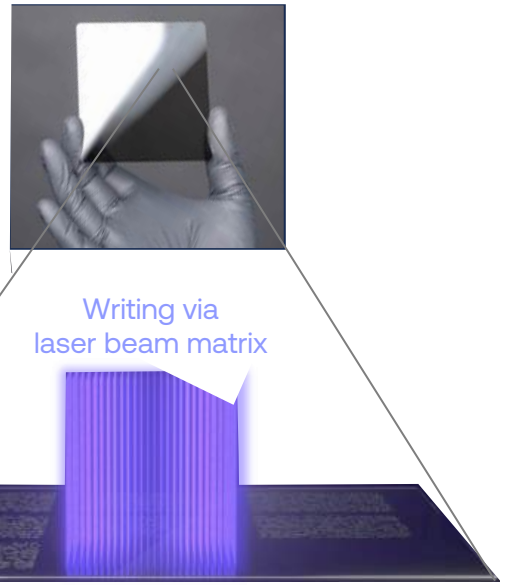

Writing via laser beam matrix

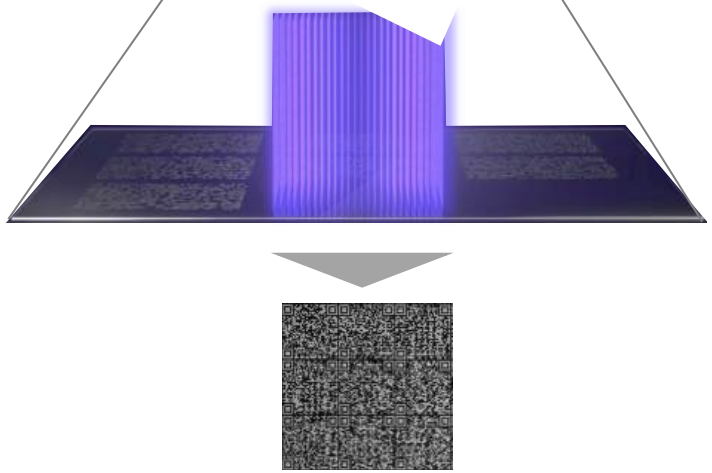

Encoded data stored in a matrix code

- **Ultra low-cost display glass** with **PVD[1] coated ceramic nano-layers**
- **Fast writing using laser beam matrix** with up to 2 Mb per shot
- **Fast reading** using microscope with **ultra-fast image sensor**
- **Fast random access to first bit** avoiding spooling time of tape
- **Stacking media in cartridges** using existing robotic libraries
- **S3 -** Standard Software Interfaces

1) PVD - Physical Vapor Deposition

Cerabyte scales from on-prem, cloud system to hyperscale system

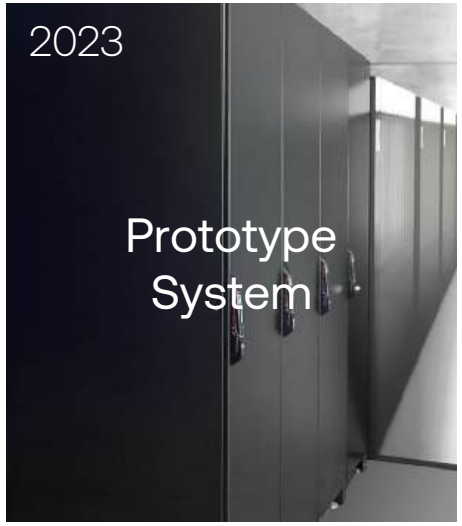

	2024-25	2025-26	2026-27	2028-30
Data capacity Write/read speed Access time	1 PB/rack 100 MB/s <90 sec to first byte	5 PB/rack 500 MB/s <30 sec to first byte	10-30 PB/rack 1+ GB/s <15 sec to first byte	60-100 PB/rack 2+ GB/s <10 sec to first byte

Cerabyte positions between magnetic tape & HDD, due to ...

Fulfillment of customer needs in 2028 (schematic)

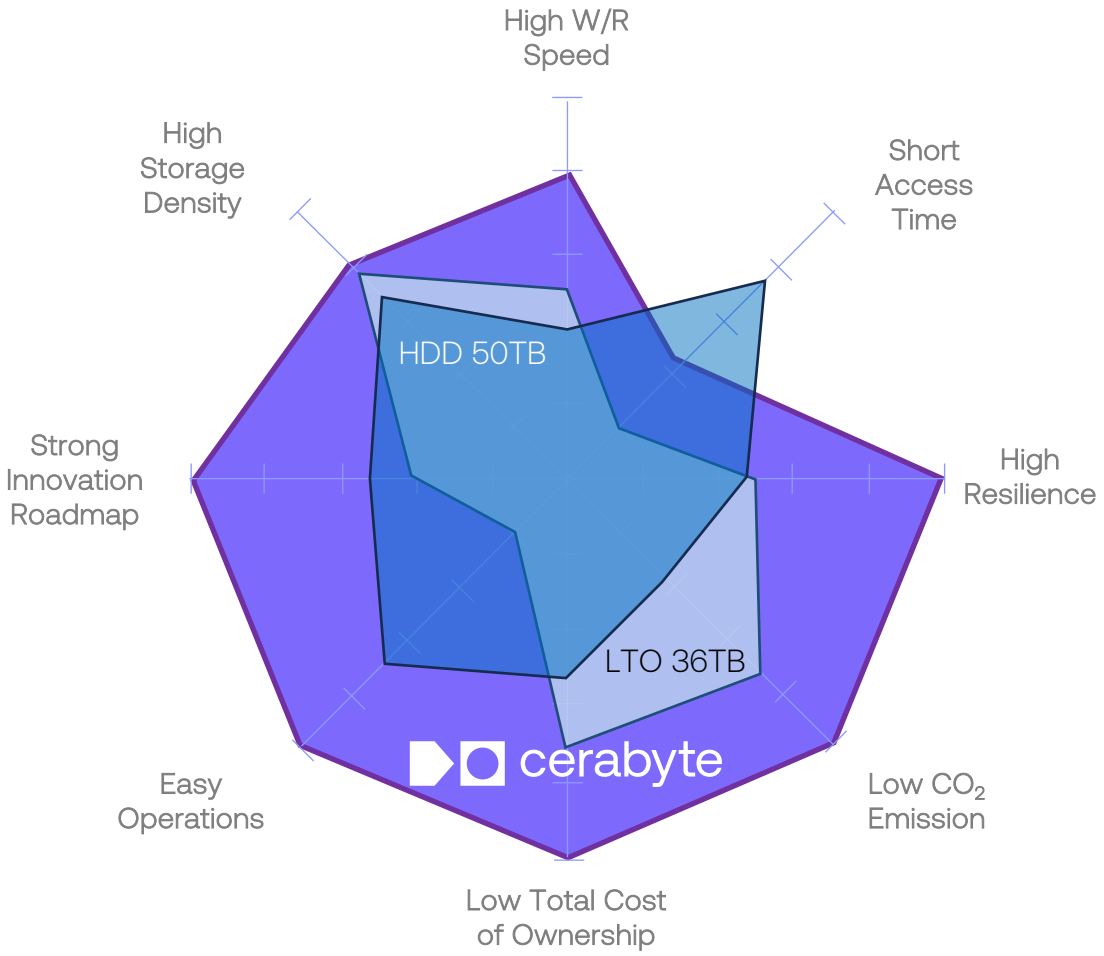

- **Extreme resilience**: High lifespan – temperature & humidity resistant etc.
- **Strong performance:** Fast W/R speed, reasonable access time
- **Easy operations:** Standard data center environment, no climatization
- **High sustainability:** 99% reduction of CO_2 vs. HDD & 100% recyclable
- **Low TCO:** Reduction of 95% vs. HDD, 75% vs. tape.

Reduction of Total Cost of Ownership by **95%** vs HDD

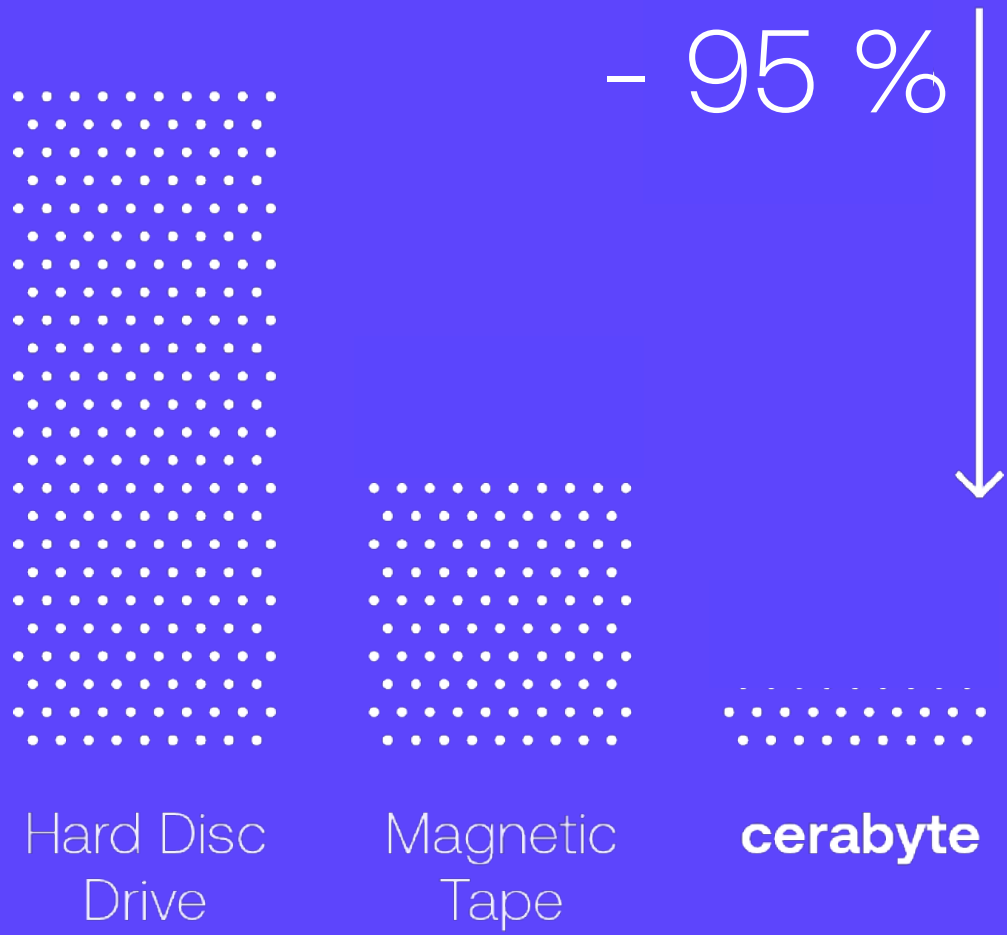

Cost Scaling

Orders of Magnitude Lower TCO

Key drivers:

- **Lower media cost**
- **Less frequent HW replacement**
- **No data migration**
- **Less energy; no A/C**

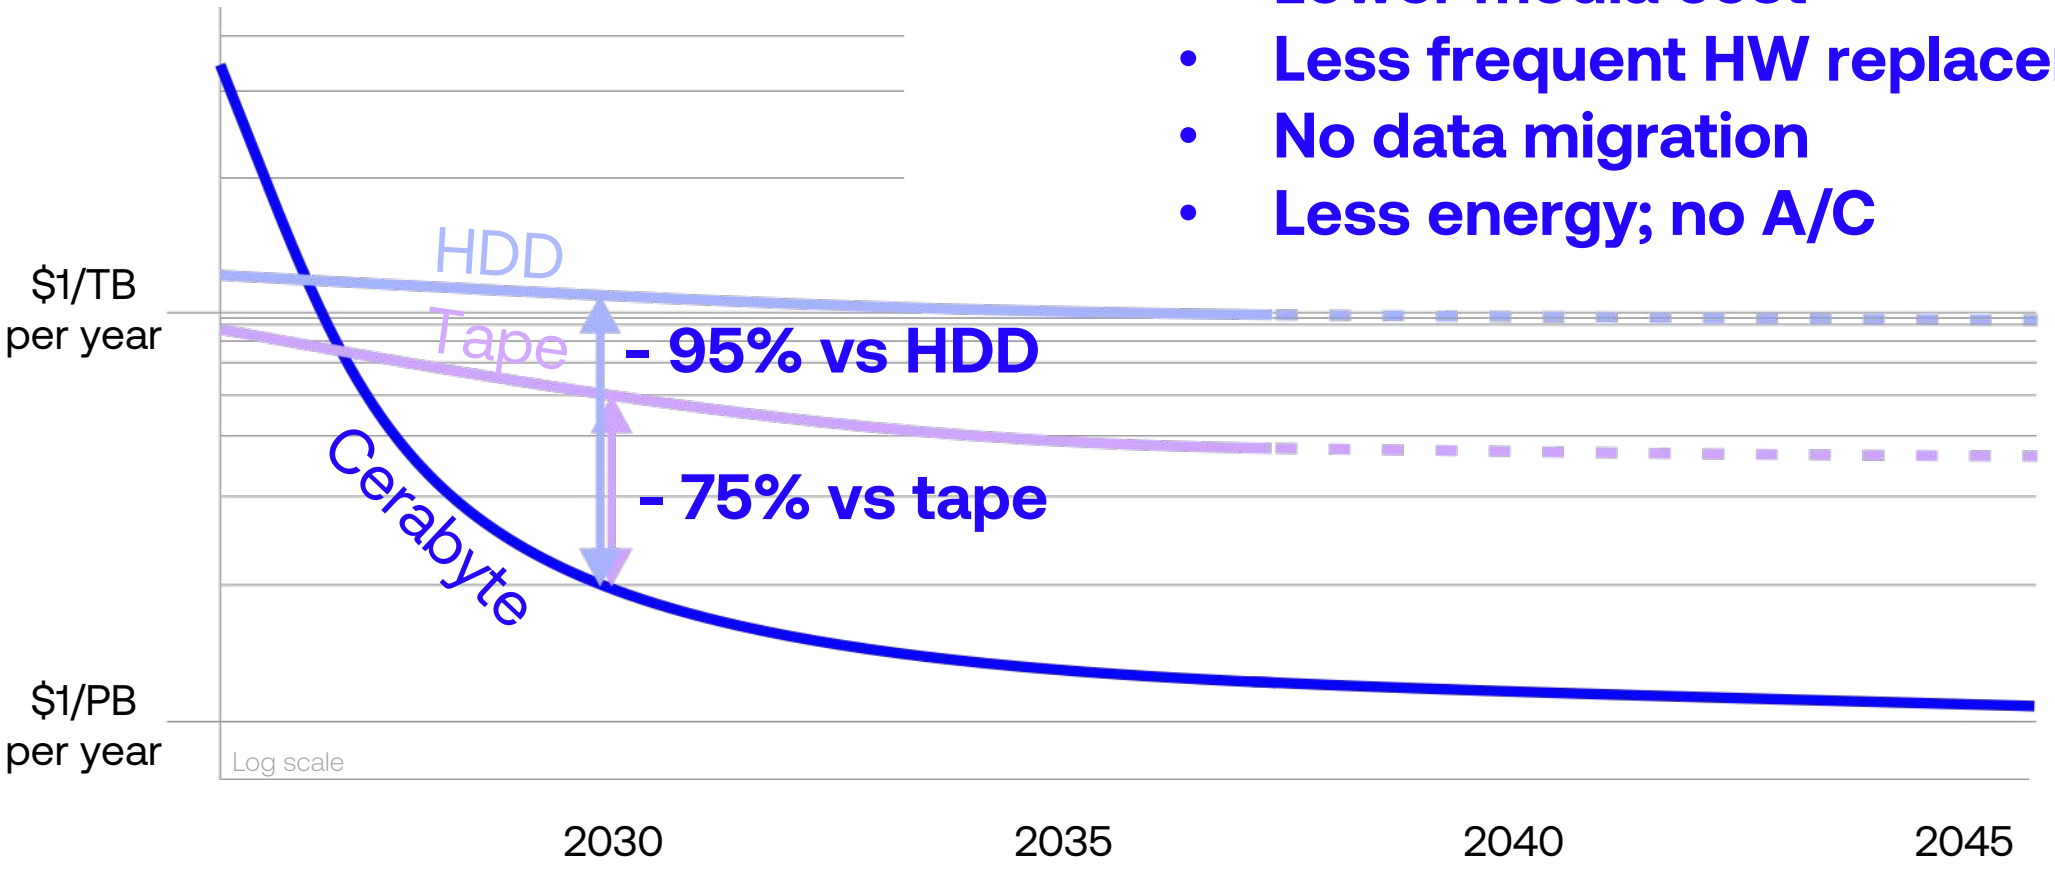

Zero energy for storing data

99%
less energy

100%
recyclable

99%
CO_2 reduction

100%
less waste

Installed Base & Shipments as base for CO_2 emission savings

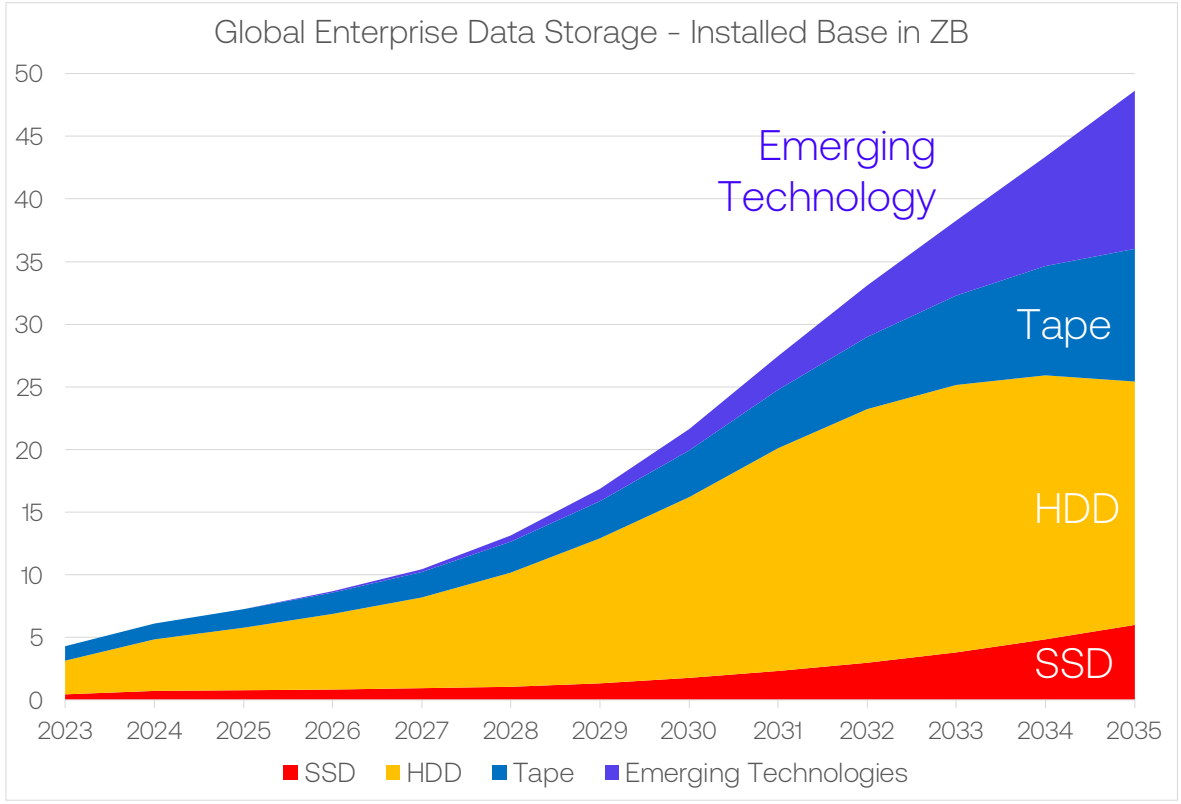

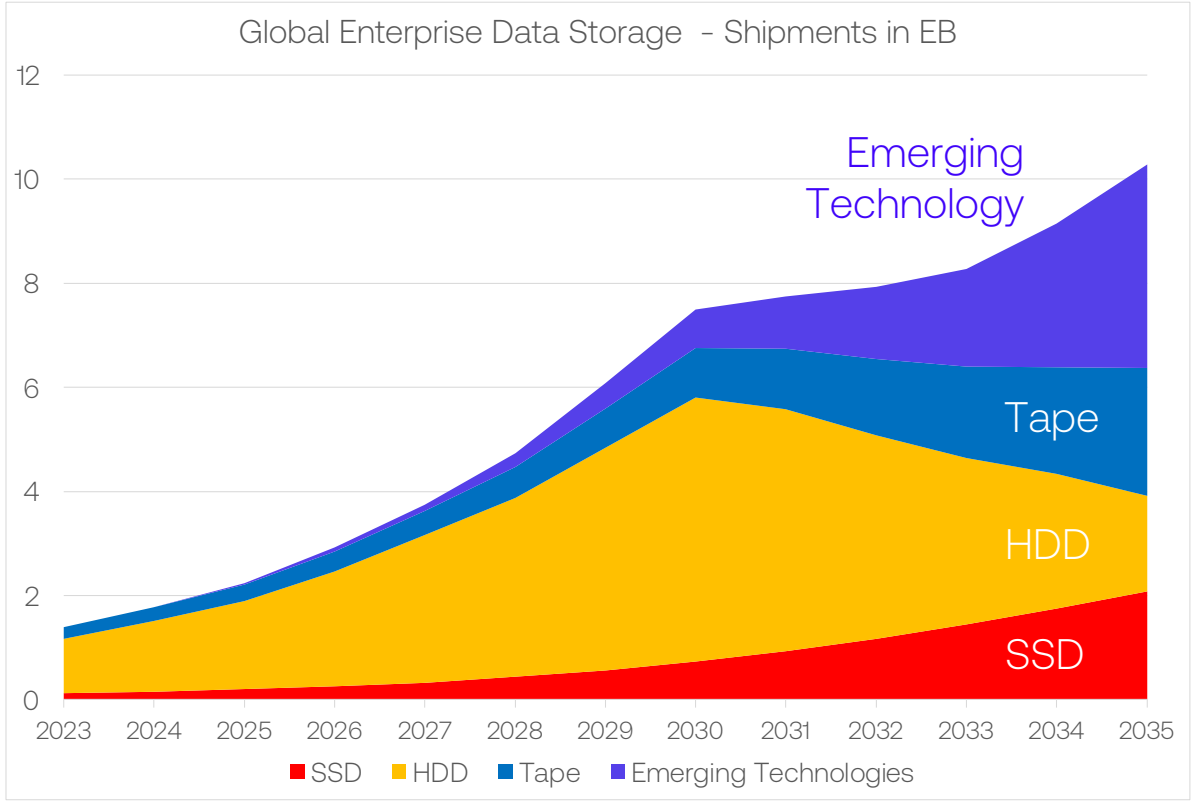

OPEX CO_2 Emissions (operation)

CAPEX CO_2 Emissions (embodied)

Source: The Sustainable Preservation of Enterprise Data
John Monroe, Brad Johns - Further Market Research - http://furthurdata.com/

CO_2 – Emission Estimate

Active Archives (Tape & Emerging) contribute significantly to CO_2 savings

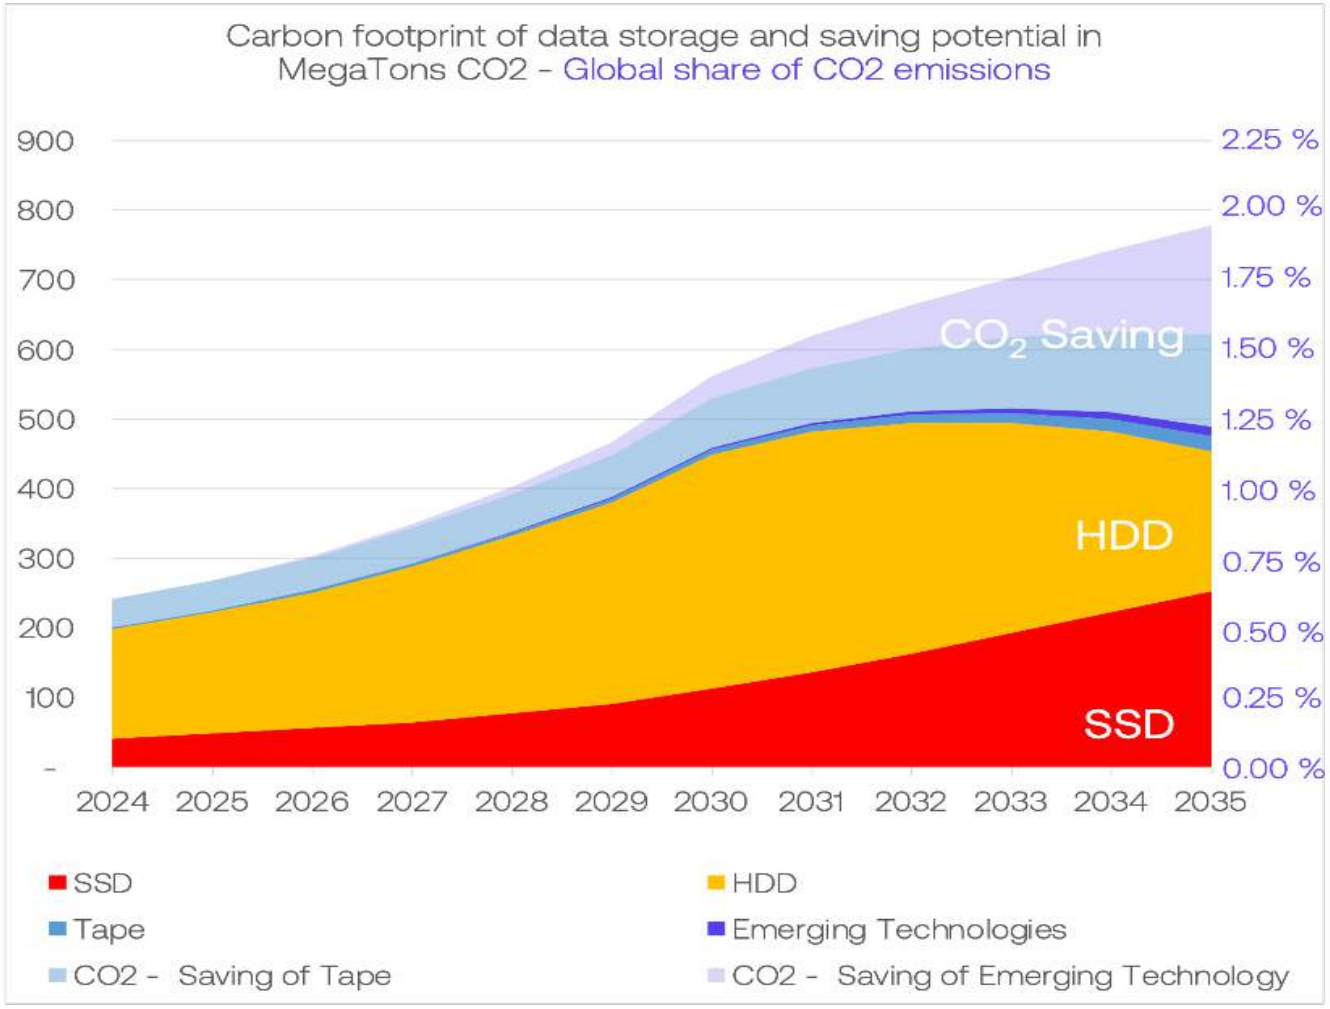

CO_2 OPEX & CO_2 CAPEX for SSD & HDD

Storage	Energy (KWh)		OPEX CO2e (Kg)		CAPEX CO2e (Kg)		Total CO2e (Kg)	
Exp. Life	5yr	10yr	5yr	10yr	5yr	10yr	5yr	10yr
HDD (1TB)	183.9	367.9	79.6	159	20	40	99.6	199
SSD (1TB)	56.9	113.8	24.6	49.2	160	320	184	369.2

Source: The Dirty Secret of SSDs: Embodied Carbon – Swamit S Tannu, Prashant J. Nair - ACM SIGEnergy Energy Informatics ReviewVolume 3Issue 3pp 4–9
https://arxiv.org/pdf/2207.10793.pdf

CO_2 Emissions in Data Centers

	DC 1	DC 2	DC 3	DC 4	Average
Server	17%	6%	1%	8%	6%
Storage	55%	51%	65%	63%	62%
Network	2%	8%	4%	1%	3%
Infrastructure	27%	36%	29%	28%	29%
	100%	100%	100%	100%	100%

Source: Green Cloud Computing – German Federal Environment Agency - ens Gröger, Ran Liu, Öko-Institut e.V., Berlin, Dr. Lutz Stobbe, Jan Druschke, Nikolai Richter, Fraunhofer-Institut für Zuverlässigkeit und Mikrointegration (IZM), Berlin
https://www.umweltbundesamt.de/sites/default/files/medien/5750/publikationen/2021-06-17_texte_94-2021_green-cloud-computing.pdf

Today - Data Storage Pyramid

Use case versus Retention, Latency & Media

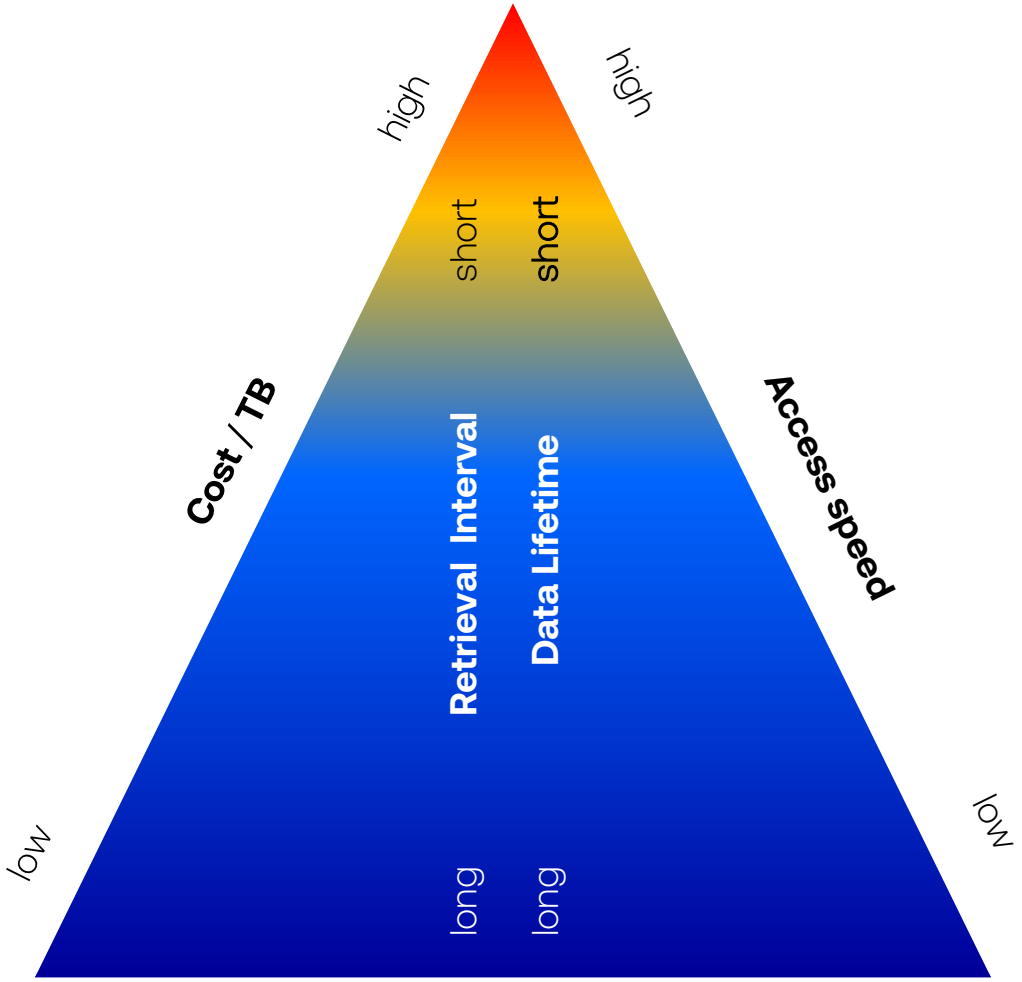

Use case	*Retention*	Latency	Media
Process data	*ps - ns*		
	ns - ms	µs	SSD
Surveillance data	*h - d*		
	weeks		
Intermediate data		ms	HDD
	years		
Health data Governmental data Financial data Liability protection Consumer archives Media & entertainment Scientific data	*decades*		
	century	min	Tape

Hypothesis 2030+: Data Storage Barbell ?

Emerging Technologies with access time of a few seconds

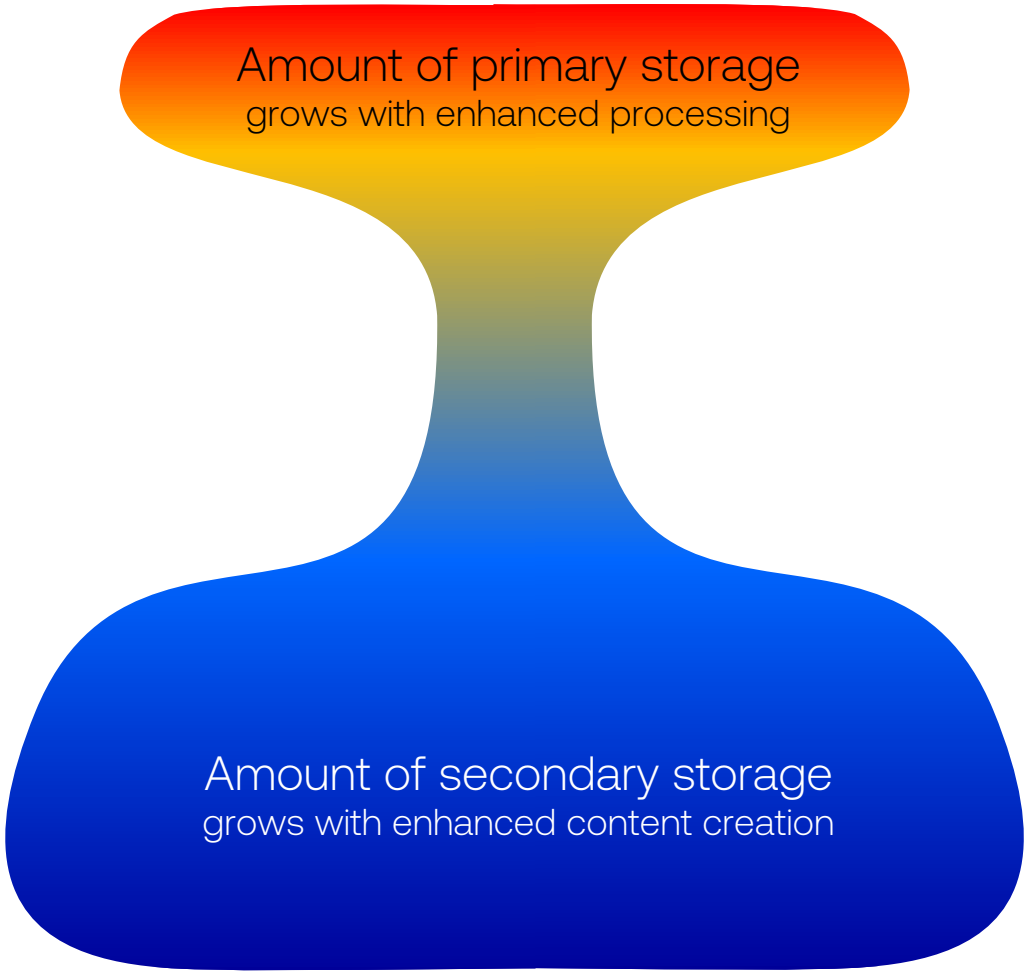

Use case	*Retention*	Latency	Media
Process data	*ps - ns*		
	ns - ms	µs	SSD
Surveillance data	*h - d*		
	weeks		
Intermediate data		ms	HDD
	years		
Health data Governmental data Financial data Liability protection Consumer archives	*decades*	**s**	**Emerging Technologies**
Media & entertainment Scientific data	*century*	min	Tape

Cerabyte & Silica: Emerging Technologies for Archival Storage in 2030s

NATIONAL ACADEMIES Sciences Engineering Medicine

NATIONAL ACADEMIES PRESS Washington, DC

This PDF is available at http://nap.nationalacademies.org/27445

Rapid Expert Consultation on Archival Data Storage Technologies for the Intelligence Community (2024)

DETAILS

27 pages | 8.5 x 11 | PAPERBACK

ISBN 978-0-309-71460-0 | DOI 10.17226/27445

CONTRIBUTORS

National Academies of Sciences, Engineering, and Medicine

BUY THIS BOOK

FIND RELATED TITLES

SUGGESTED CITATION

National Academies of Sciences, Engineering, and Medicine. 2024. *Rapid Expert Consultation on Archival Data Storage Technologies for the Intelligence Community*. Washington, DC: The National Academies Press. https://doi.org/10.17226/27445.

SSD, HDD & Tape (commercially available)

- Replacement every 5-10 years
- No long-term storage media for 25-50 years
- Capacity increase limited to 10-25% p.a.

Emerging Technologies for Archival Storage

Cerabyte - TRL 6 (Technical Readiness Level 6)

- Cerabyte will be available at scale in 2025+
- 100 PB rack system until 2030
- Exa-scale racks with particle beam 2035+

Microsoft Silica - TRL 6

- Silica will be available at scale in 2030+
- Limiting factor are laser sources until 2030
- Only available via Microsoft Azure

DNA Data Storage (TRL 4)

- DNA will be available at scale only in 2040+
- Access time in minutes rather than seconds
- Sensitive to oxidation, UV, radiation and temperature

Leveraging Semiconductor Fab Tool Technology

to scale density and access speed

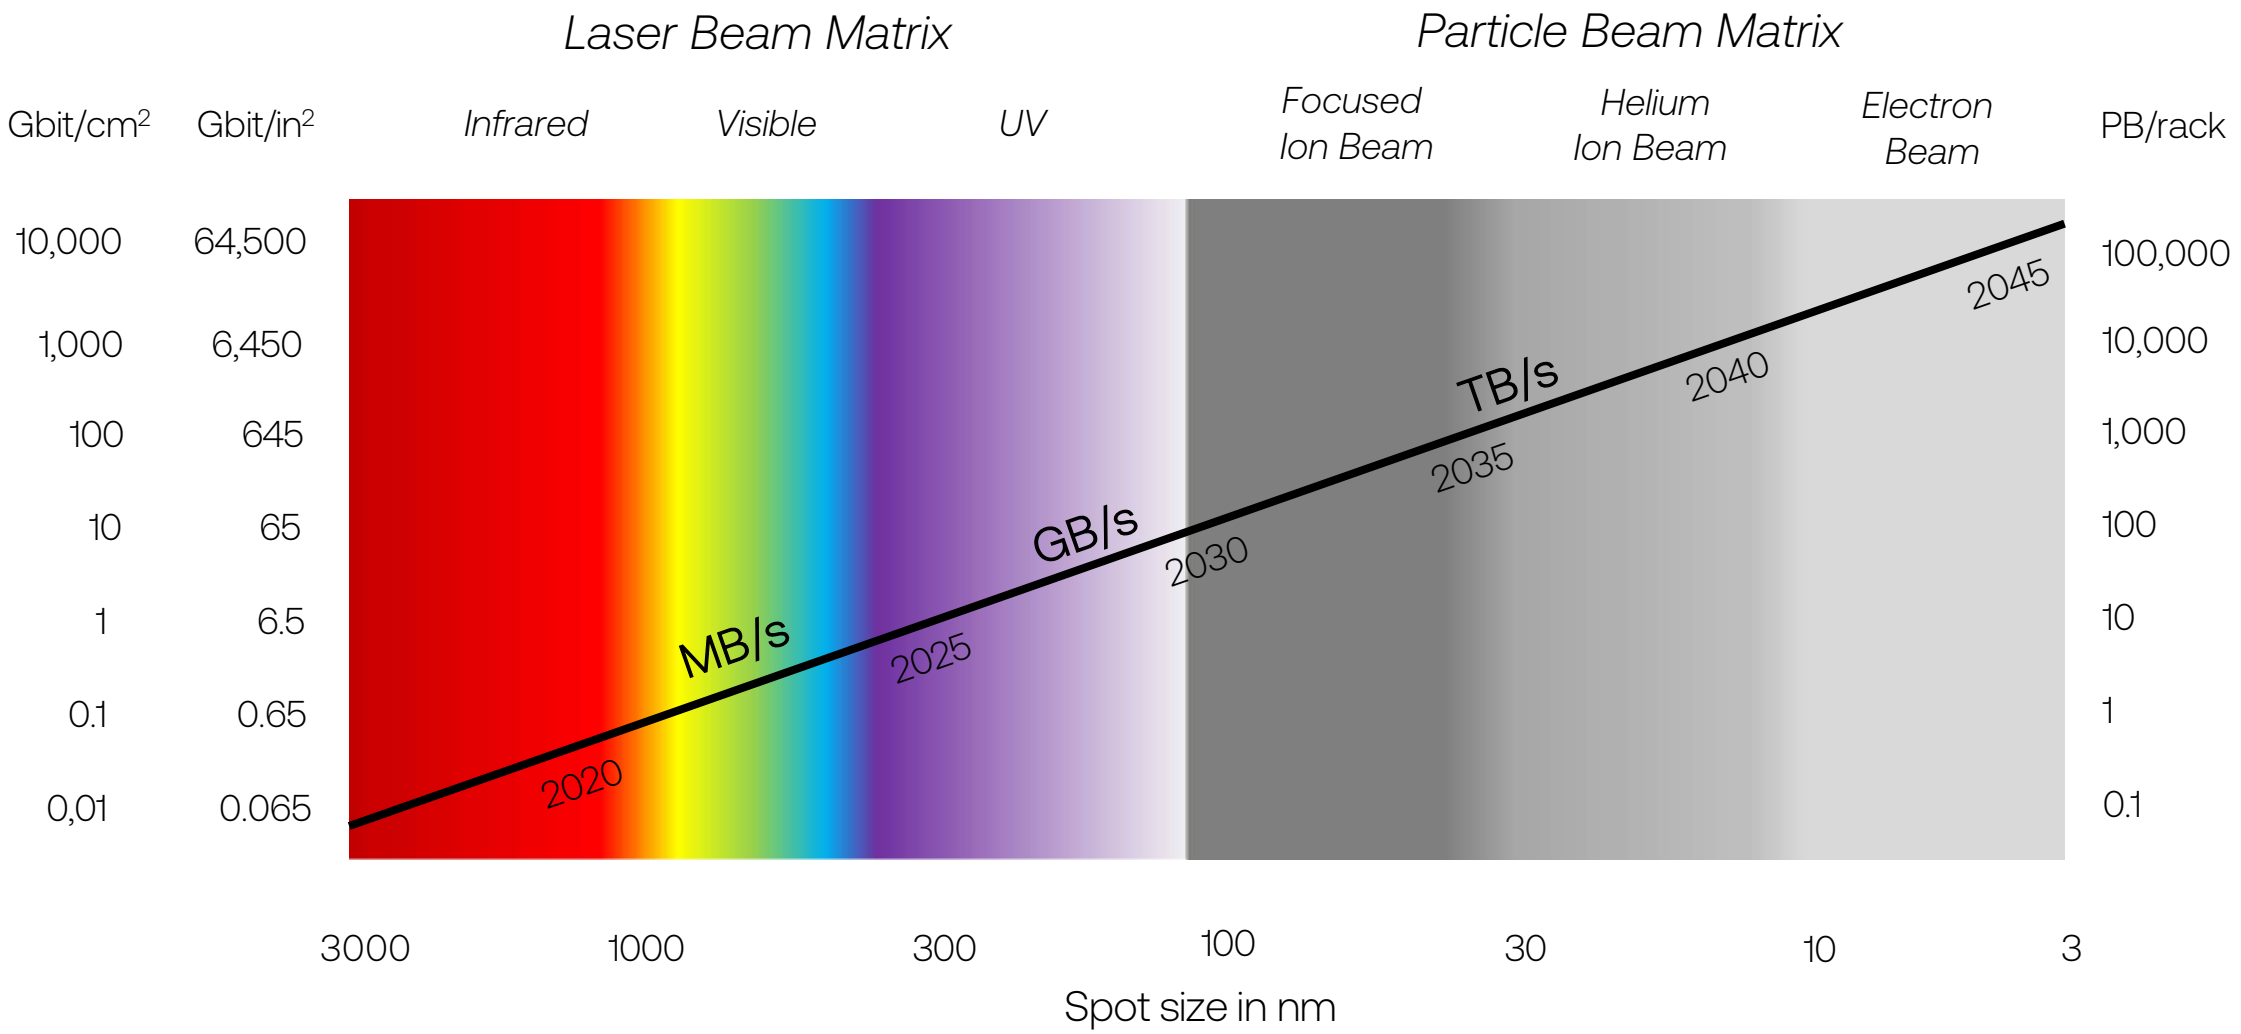

Technology Outlook

Electron-beam is pushing pixel size to **1,5 nm or 400+ TB per 4 x 4 in²**

	Laser beam matrix		**Particle beam matrix**	
		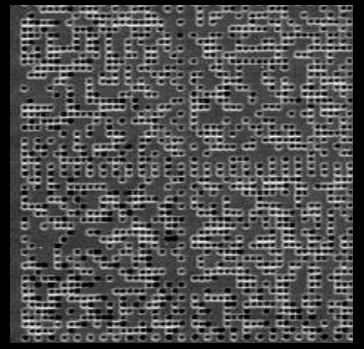		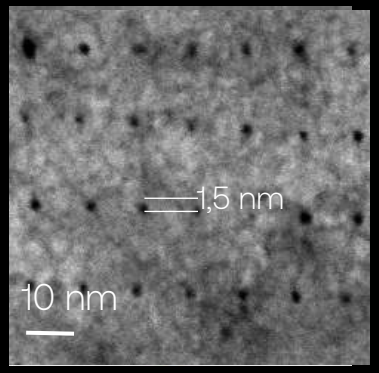
	ps Laser	fs Laser	Ion Beam	Electron Beam
Pixel size	1,0 µm	100 nm	30 nm	1,5 nm
Data capacity per 10x10 cm² or 4x4 in²	1,25 GB	125 GB	1+ TB	400 TB

ICSD laboratory - Innovative Coatings for Sustainable Data Storage

Cerabyte & Leading Research Partners are exploring ICSD Laboratory

Ceramic Nano Coating	Laser Beam	Particle Beams

Fluorescent materials

Reflective Ceramics

Fluorescent Ceramics

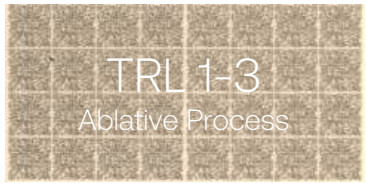

Research Partner: Recendt (AT)

UV – laser (100 nm - W/R)

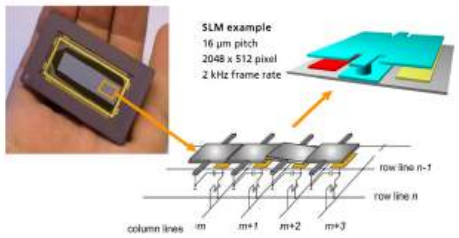

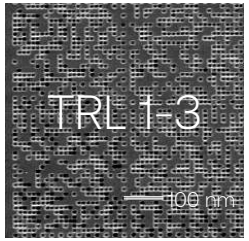

Research partner: Fraunhofer IPMS – Dresden (DE)

Helium Ion Beam (30 nm W/R)

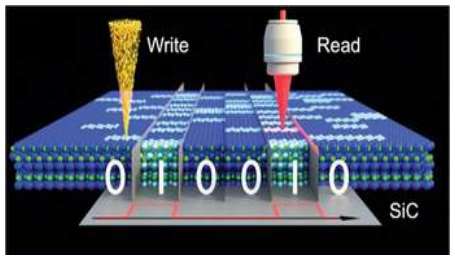

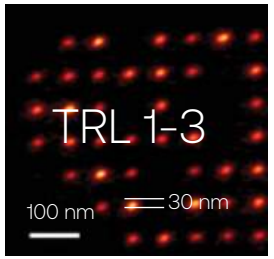

Research Partner: Helmholtz HZDR Dresden (DE)

Magnetic materials

Conductive Ceramics

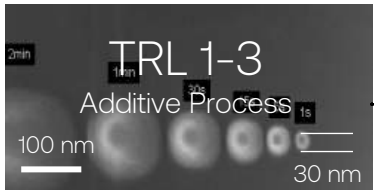

Magnetic Materials

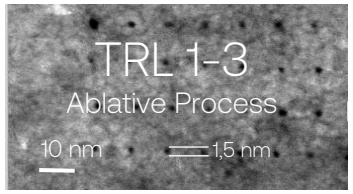

Research partner: IMEC (BE)

Multi-Layer & Overlapping Matrix

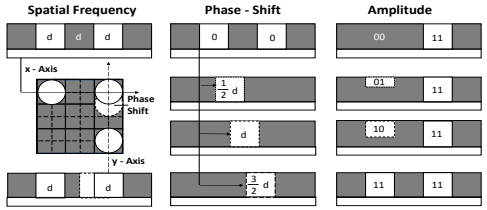

Research partner: FH Bern (CH)

Electron Beam Matrix (1,5 nm W/R)

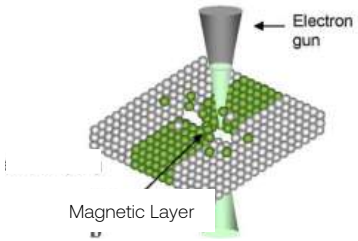

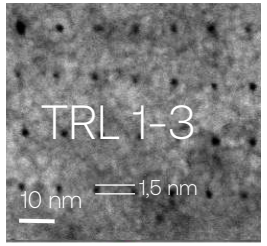

Research partner: TU Vienna (AT)

Building the Ceramic Data Storage Eco-System

USA Technology Cluster

- Pure Storage
- UC Santa Cruz
- University of Colorado Boulder
- **YOUR ORGANISATION**

European Technology Cluster

- Technical University Vienna
- CERN – Geneva
- IMEC – Loewen
- Helmholtz HZDR – Dresden

YOUR PARTICIPATION

- Contribute to Ceramic Data Storage architecture.
- Define Ceramic Data Storage hardware and software platform.
- Input from the SNIA Community is highly welcome!

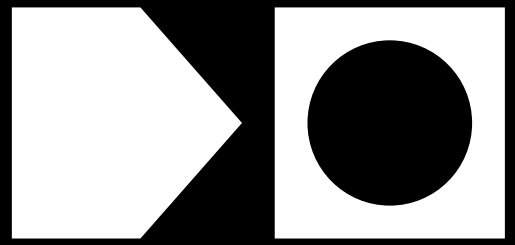

cerabyte

Store all data forever

Christian Pflaum, CEO & Founder
christian.pflaum@cerabyte.com

Steffen Hellmold, Director
steffen.hellmold@cerabyte.com

Cerabyte GmbH
Rundfunkplatz 2
D 80335 Munich
Germany

Cerabyte Inc.
110 Del Oro Lagoon
CA 94949 Novato
USA