IBM Global Technology Outlook: Rôle des données et transformation des systèmes de stockage

11.12.2015

Dr. Robert Haas
CTO Storage Europe, IBM

Topics

- What's cognitive business?
- Transformation of storage to new consumption models
- Selection of advanced innovation topics
  - Flash, Tape, hybrid cloud storage, hyper-convergence

Cognitive Computing: Technologies that will change the business and the needs in terms of data storage

#CognitiveEra

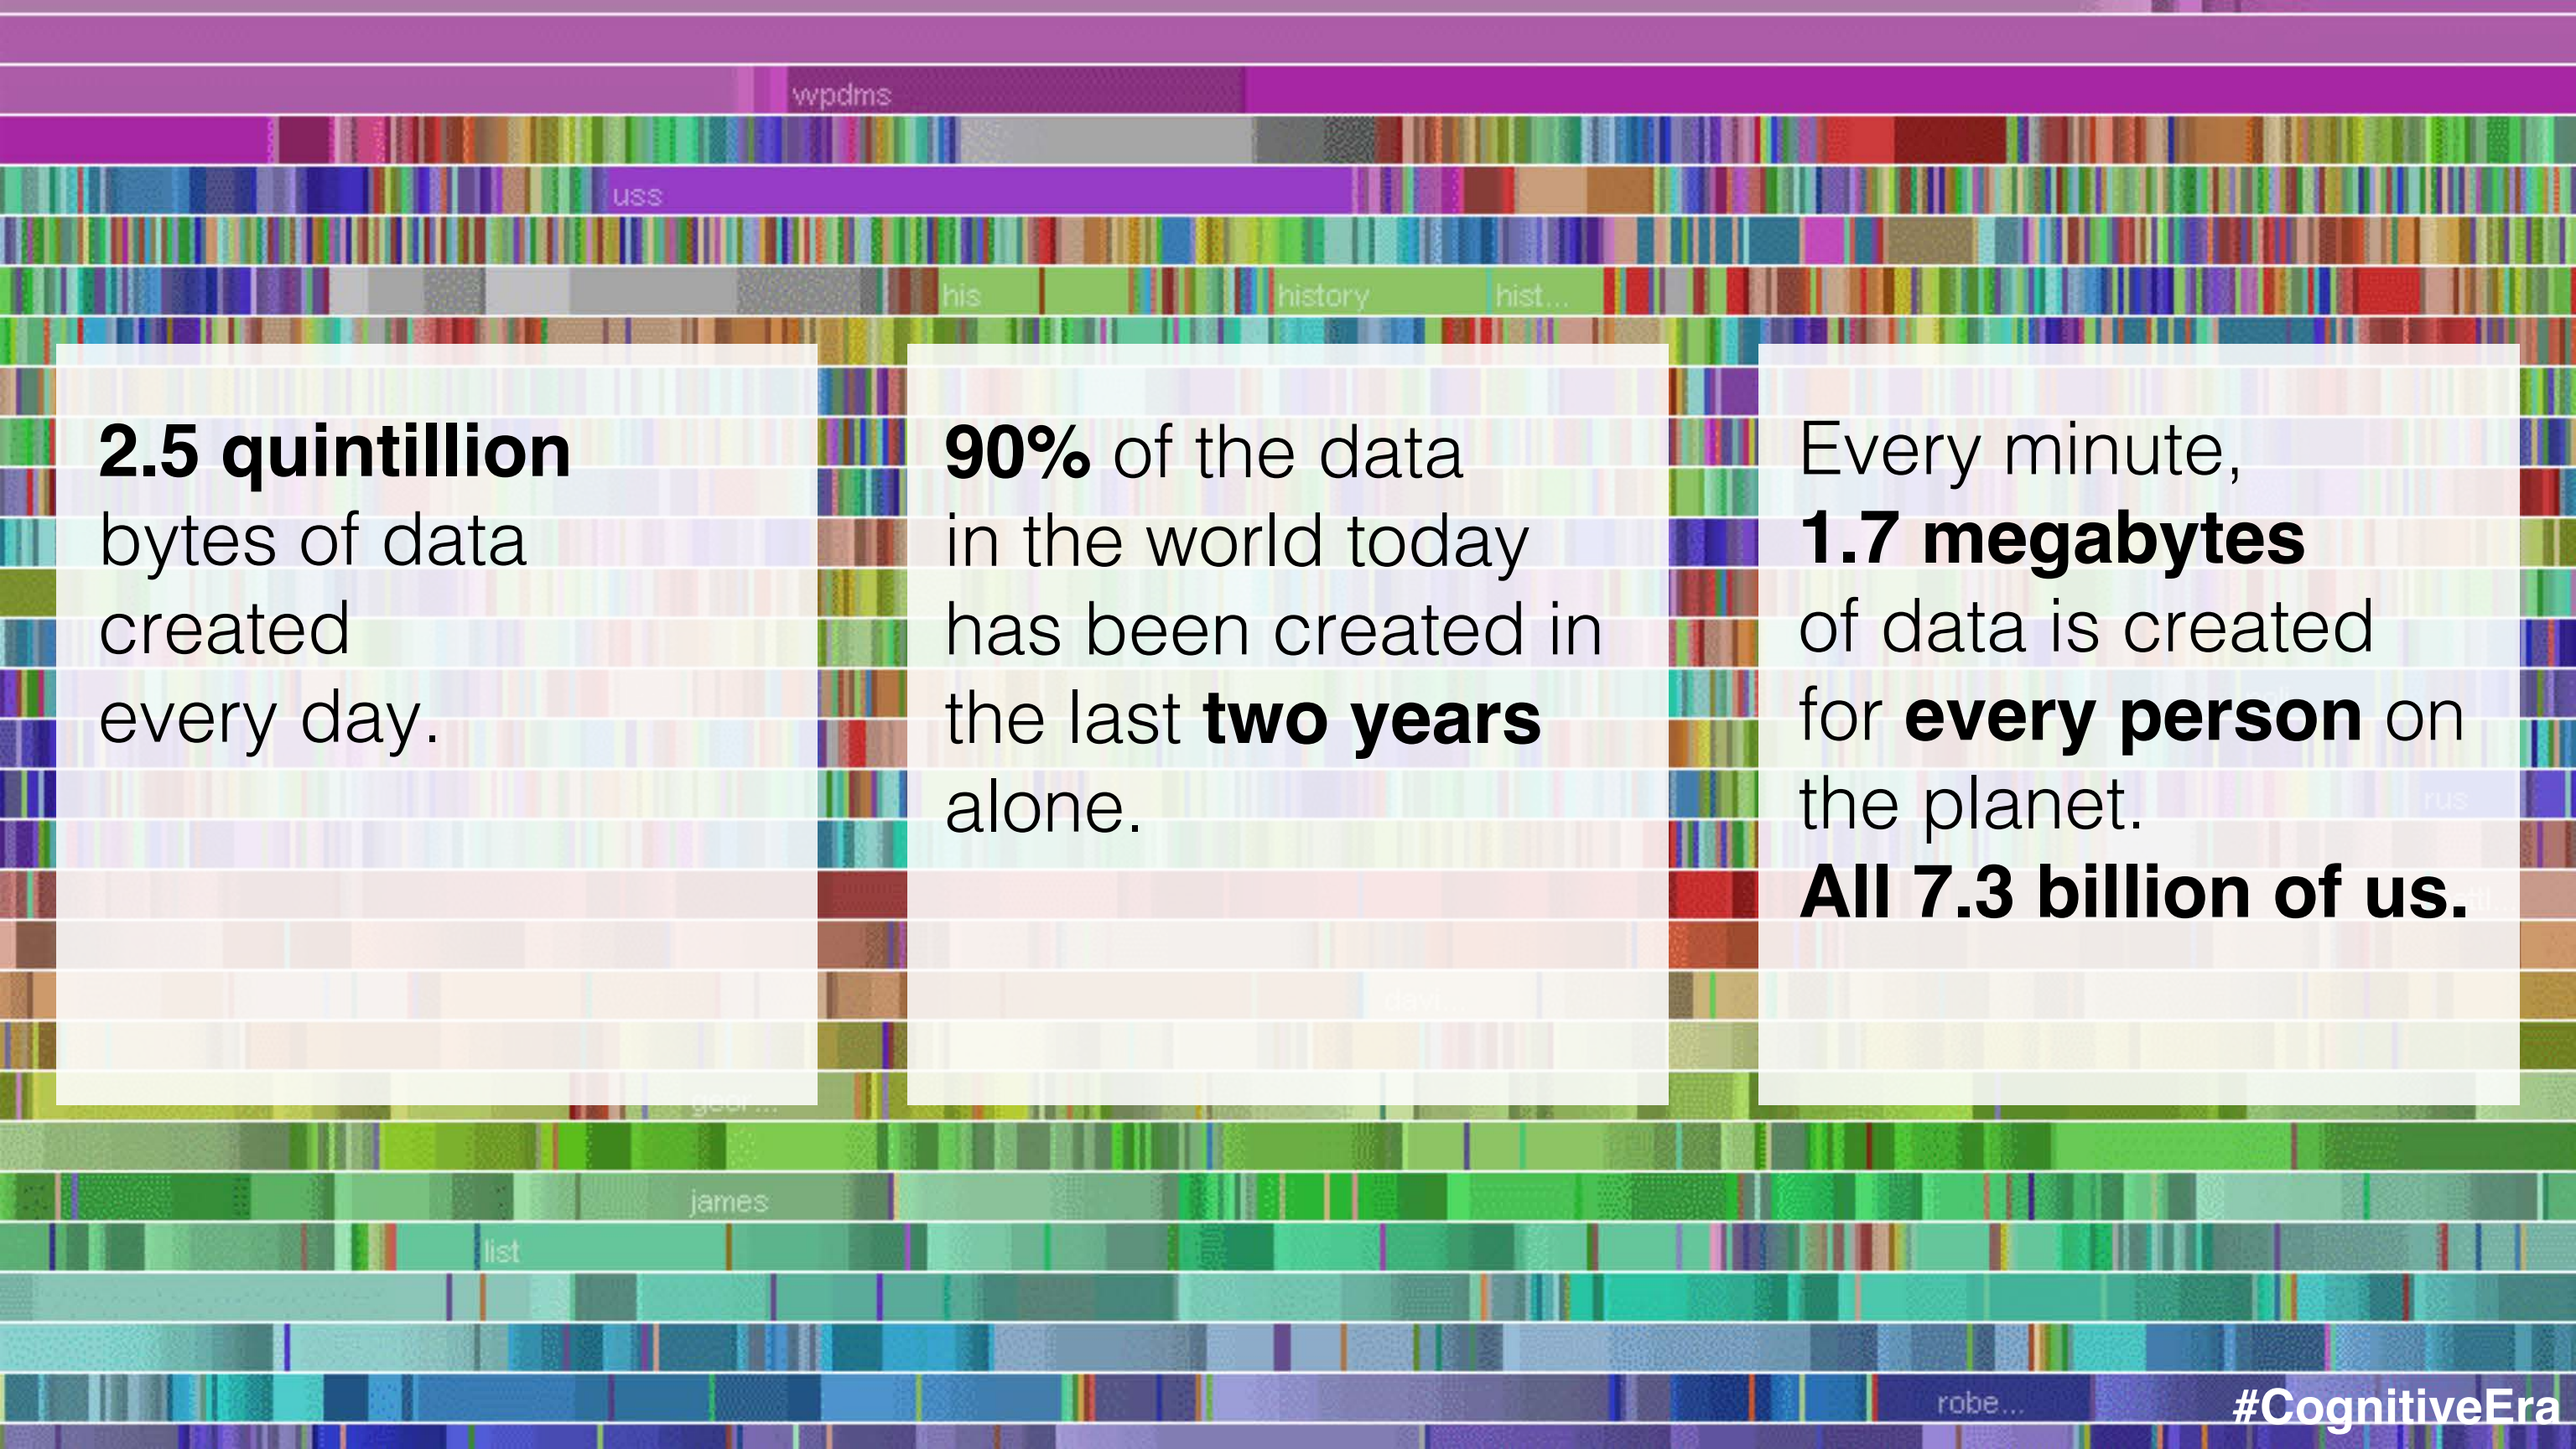

2.5 quintillion
bytes of data
created
every day.

90% of the data
in the world today
has been created in
the last **two years**
alone.

Every minute,
1.7 megabytes
of data is created
for **every person** on
the planet.
All 7.3 billion of us.

#CognitiveEra

Unstructured data — "dark data" — accounts for 80% of all data generated today.

#CognitiveEra

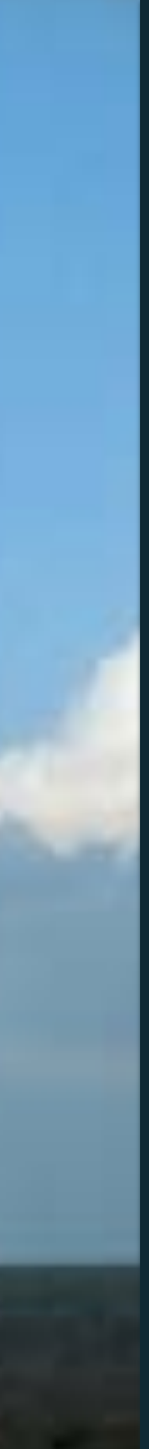

This is expect to grow to 93% by 2020.

#CognitiveEra


#CognitiveEra

Oil & Gas

Modern facilities have more than **80,000 sensors** in place, and a single reservoir will produce more than **15 petabytes** of data in its lifetime.

#CognitiveEra

Retail

Consumers post **500 million** tweets and **55 million** Facebook updates each day.

#CognitiveEra

Public Safety

New York City surveillance cameras and sensors generate **520 TB** of data per day, largely unstructured.

#CognitiveEra

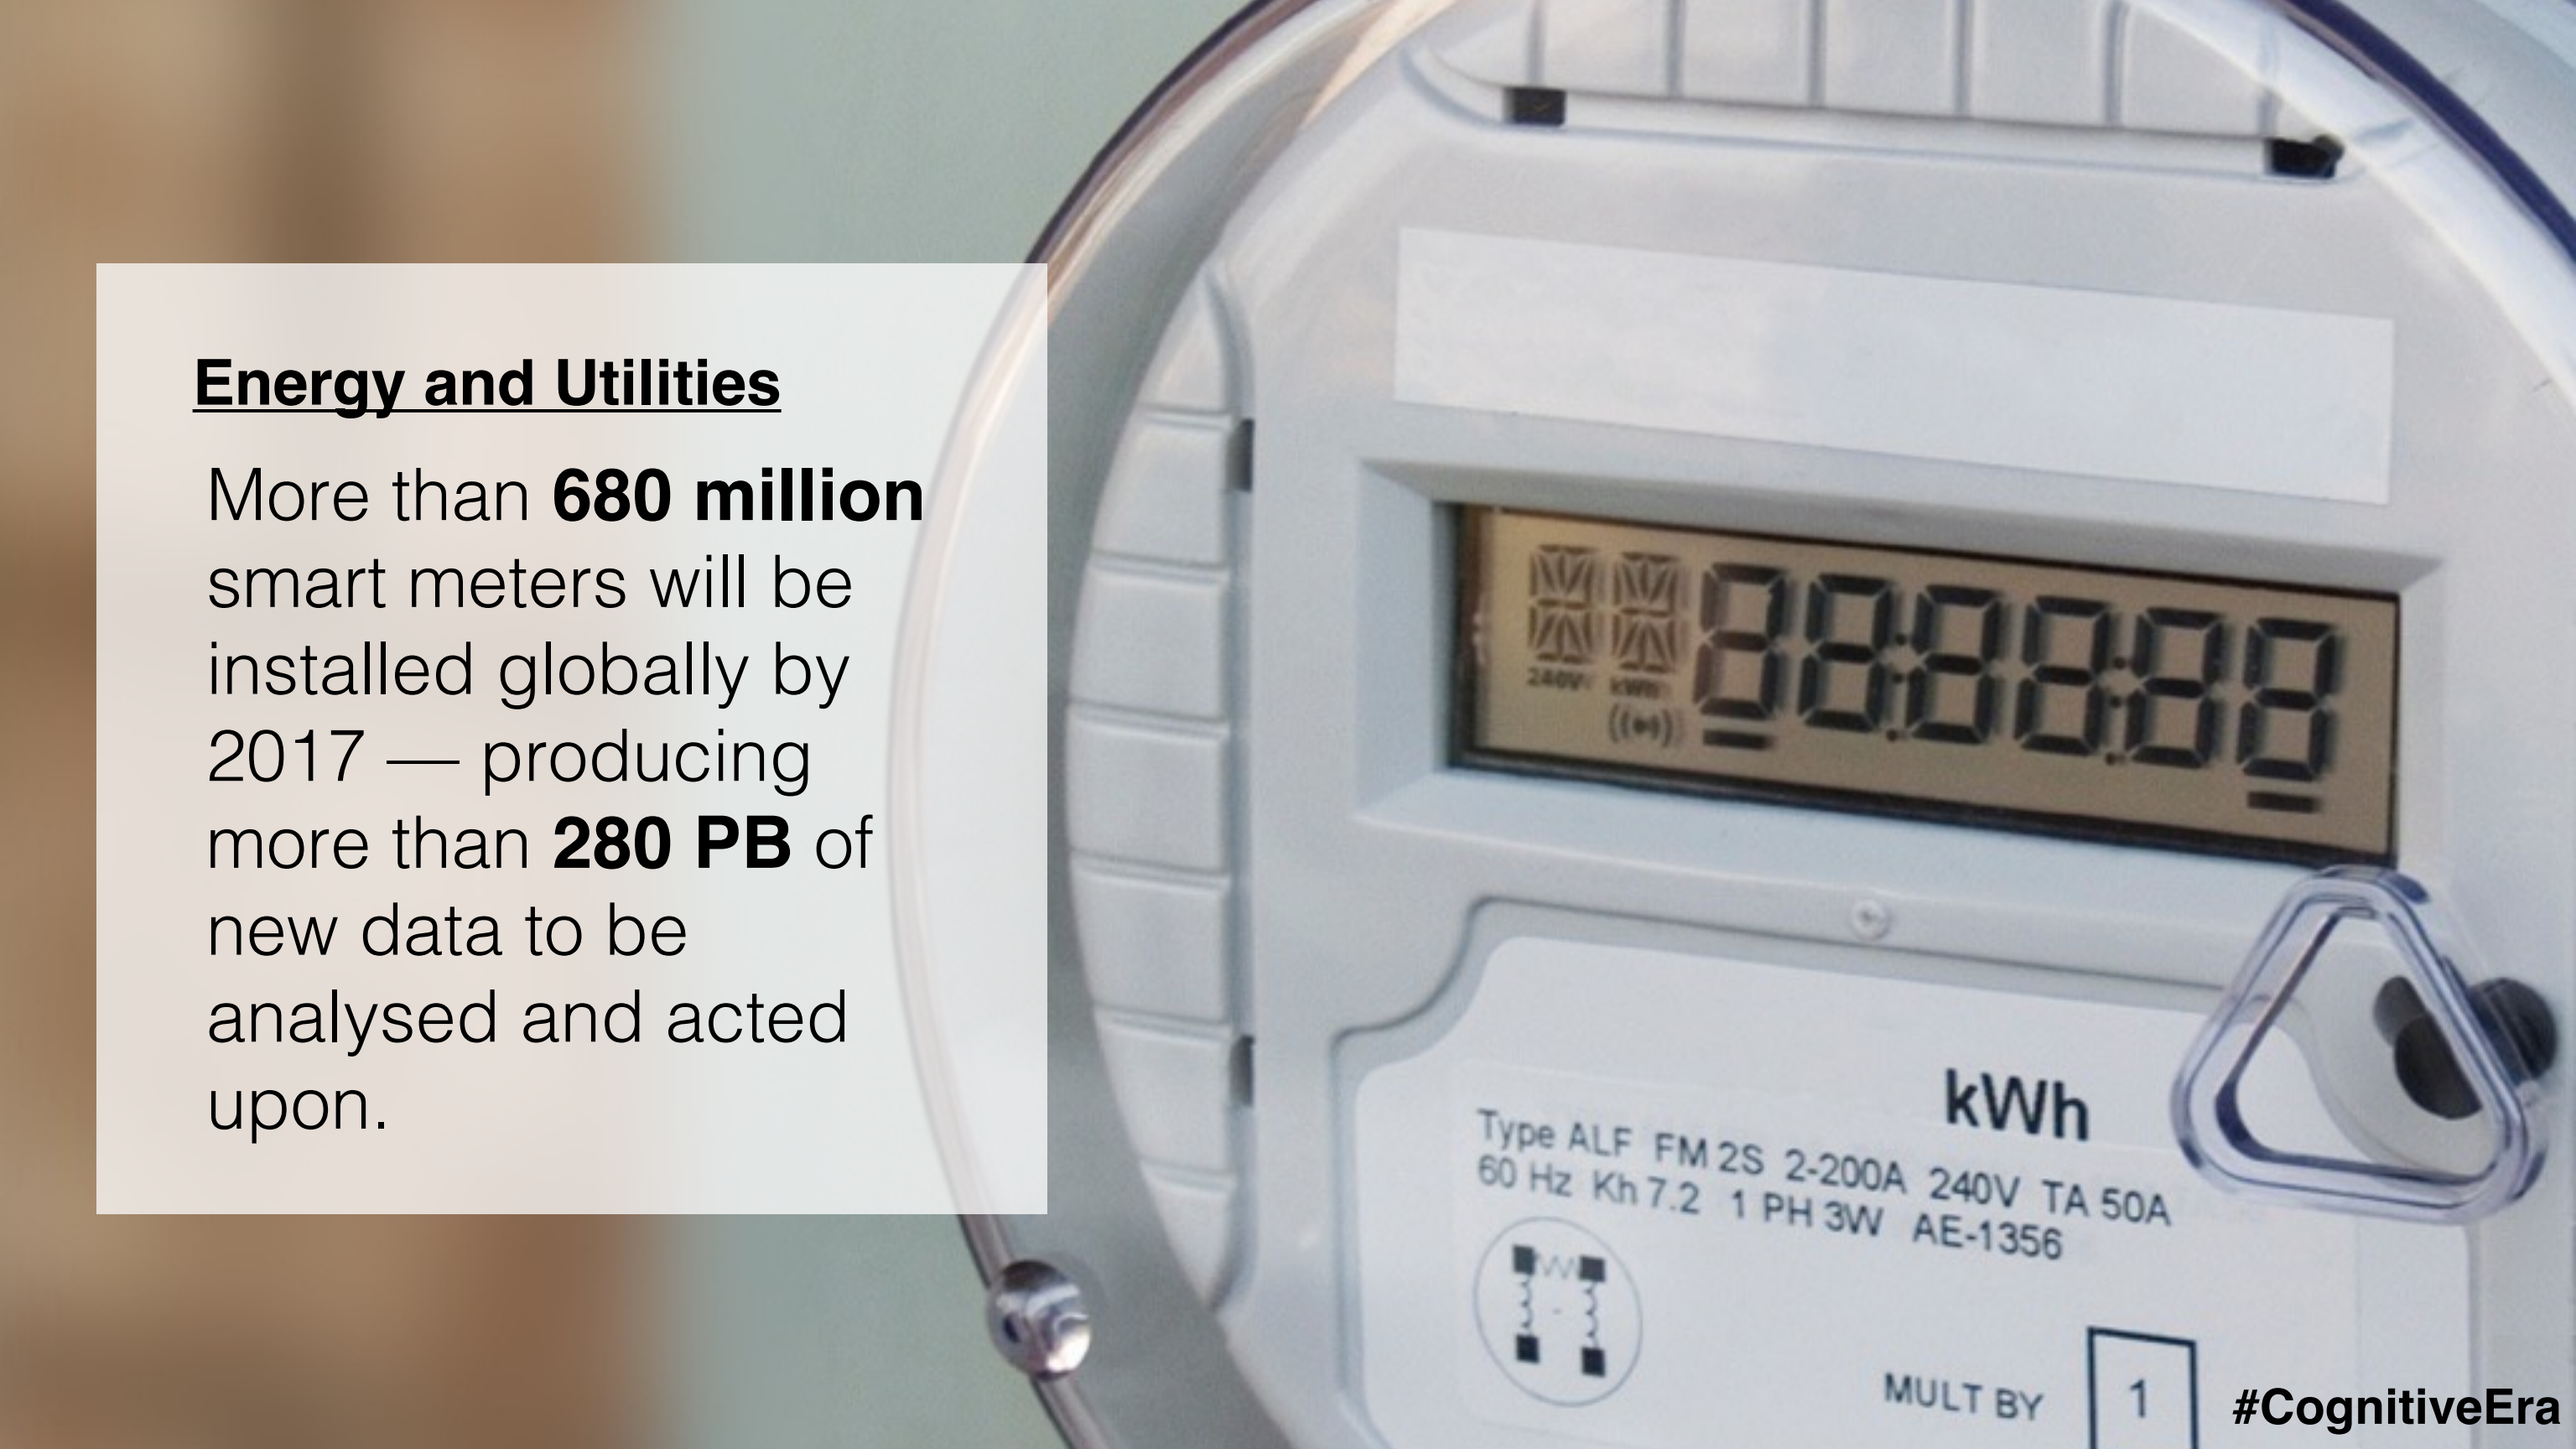

Energy and Utilities

More than **680 million** smart meters will be installed globally by 2017 — producing more than **280 PB** of new data to be analysed and acted upon.

#CognitiveEra

Healthcare

Each person will generate **1 million GB** of health-related data in their lifetime — equivalent to about **300 million** books

#CognitiveEra

Transportation

By 2020, **75%** of the world's cars will be connected…and they will produce **350 MB** of data per second to be assessed and acted upon.

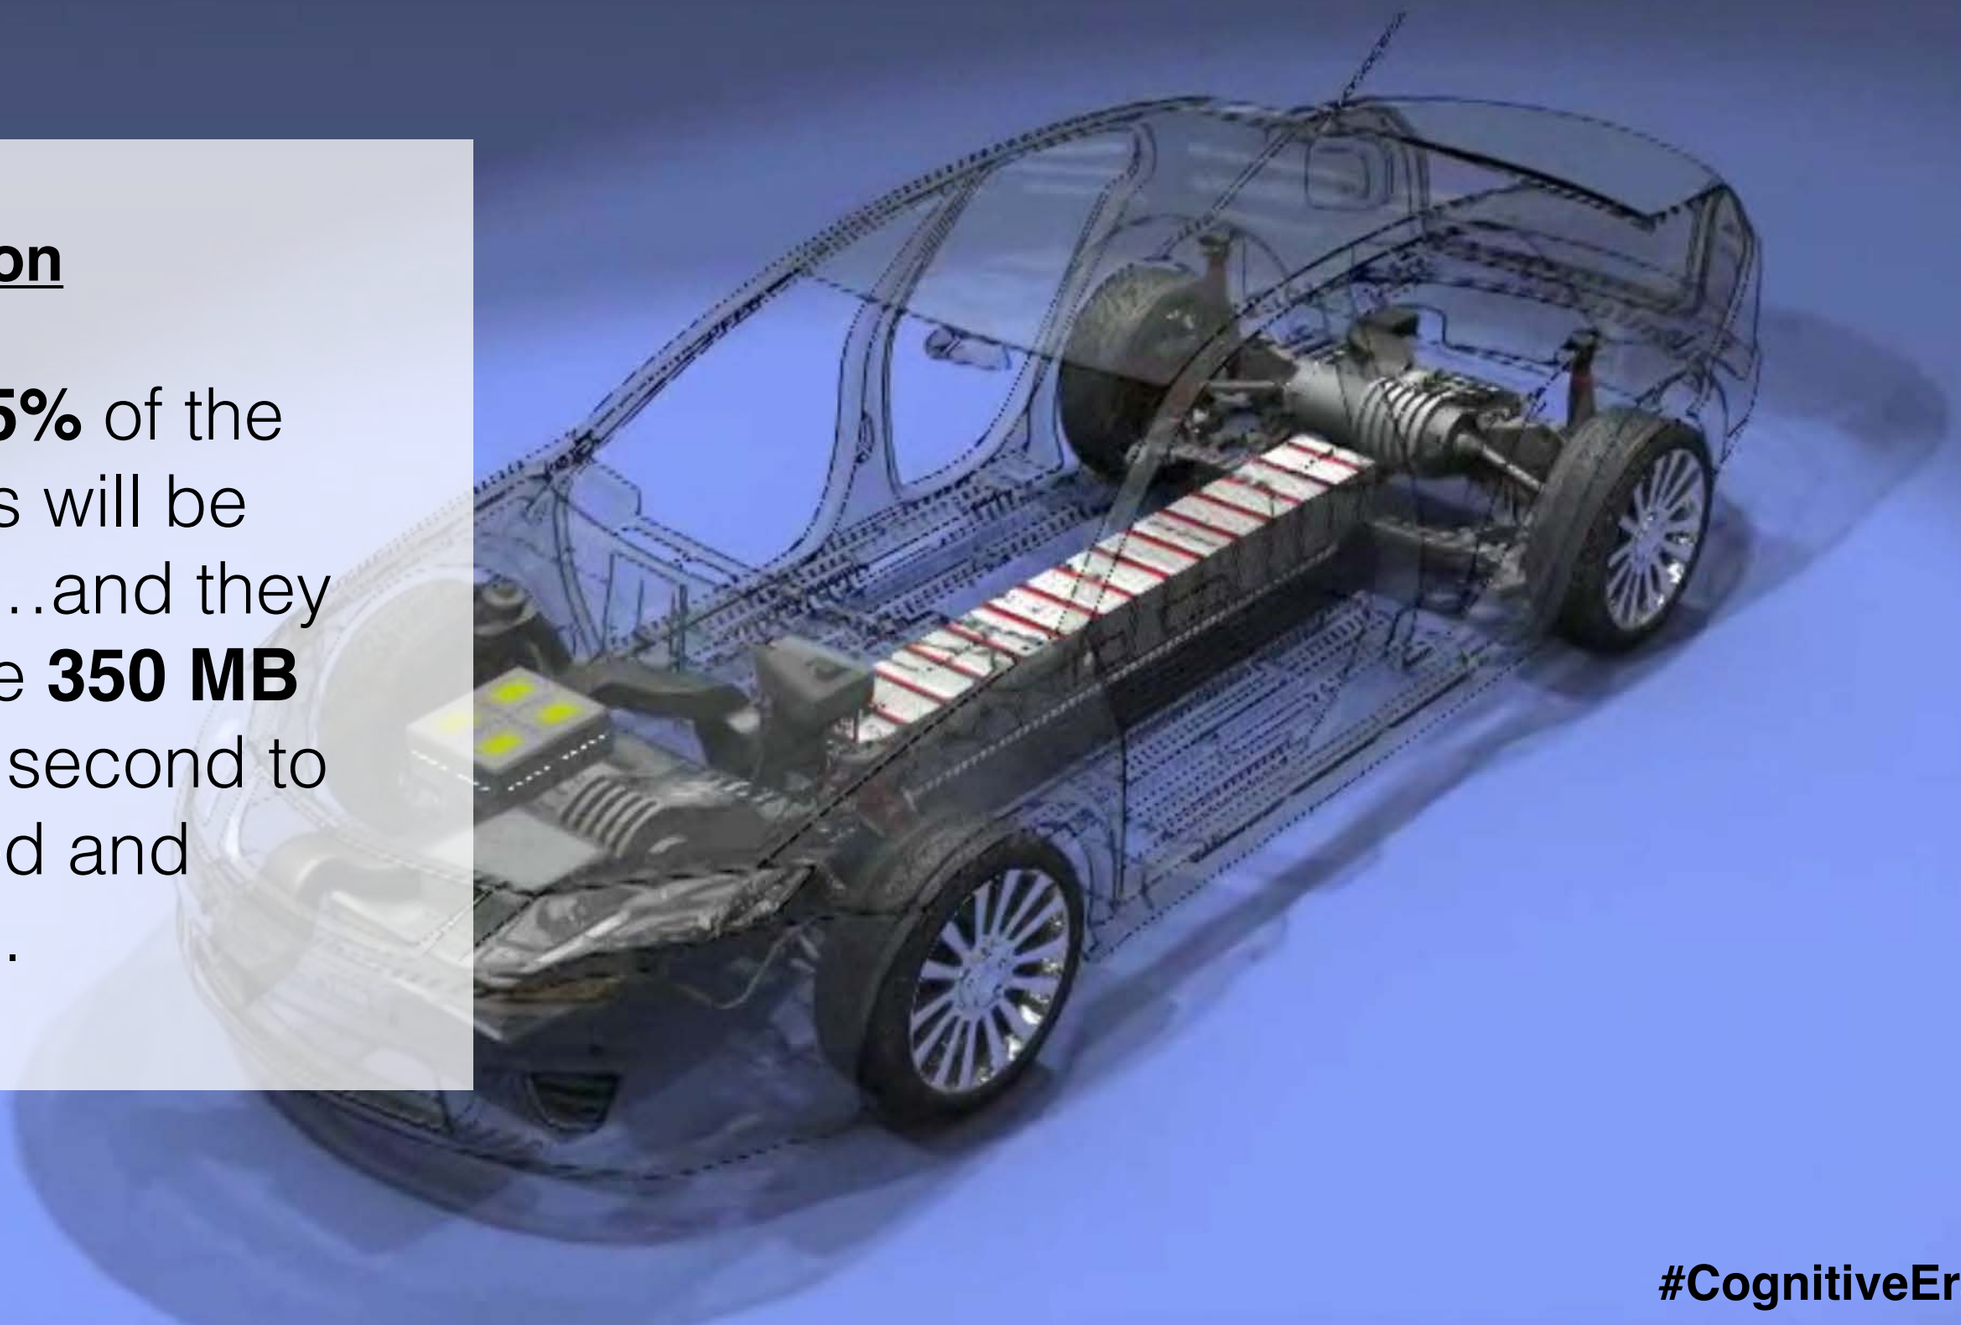

#CognitiveEra

Tabulating Systems Era

1900 - 1940s

Programmable Systems Era

1950s - Present

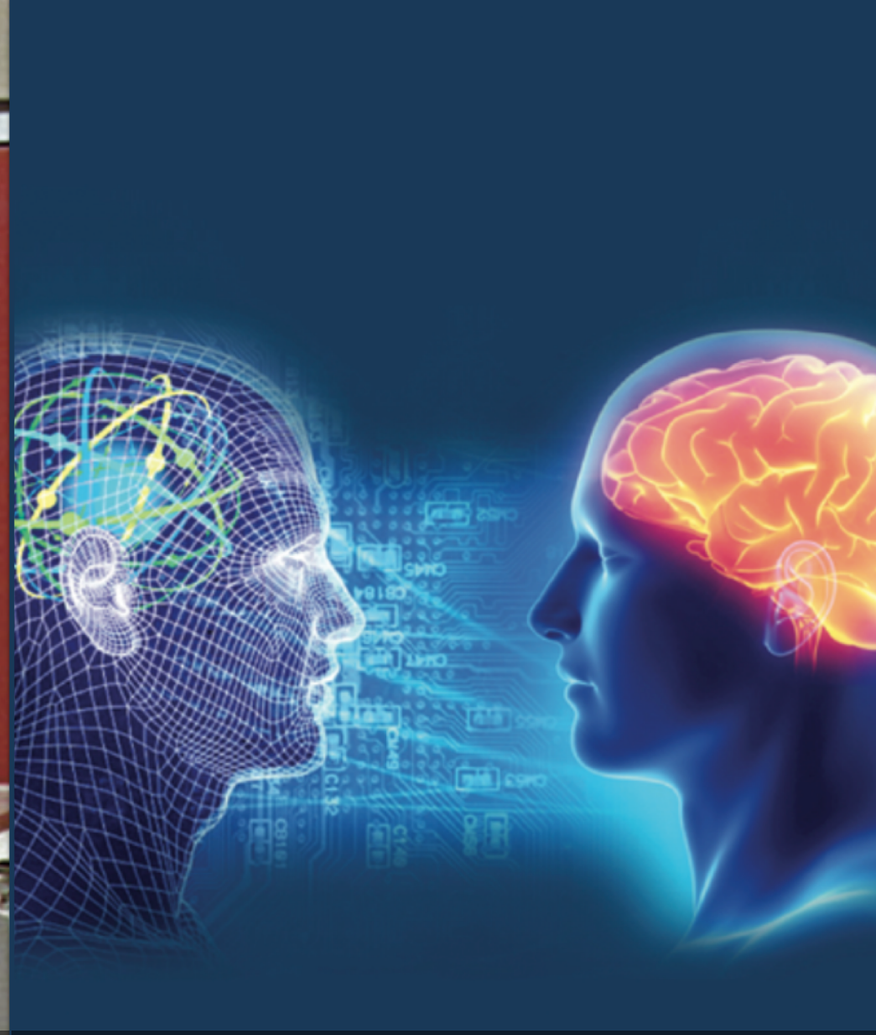

Cognitive Systems Era

2011 -

#CognitiveEra

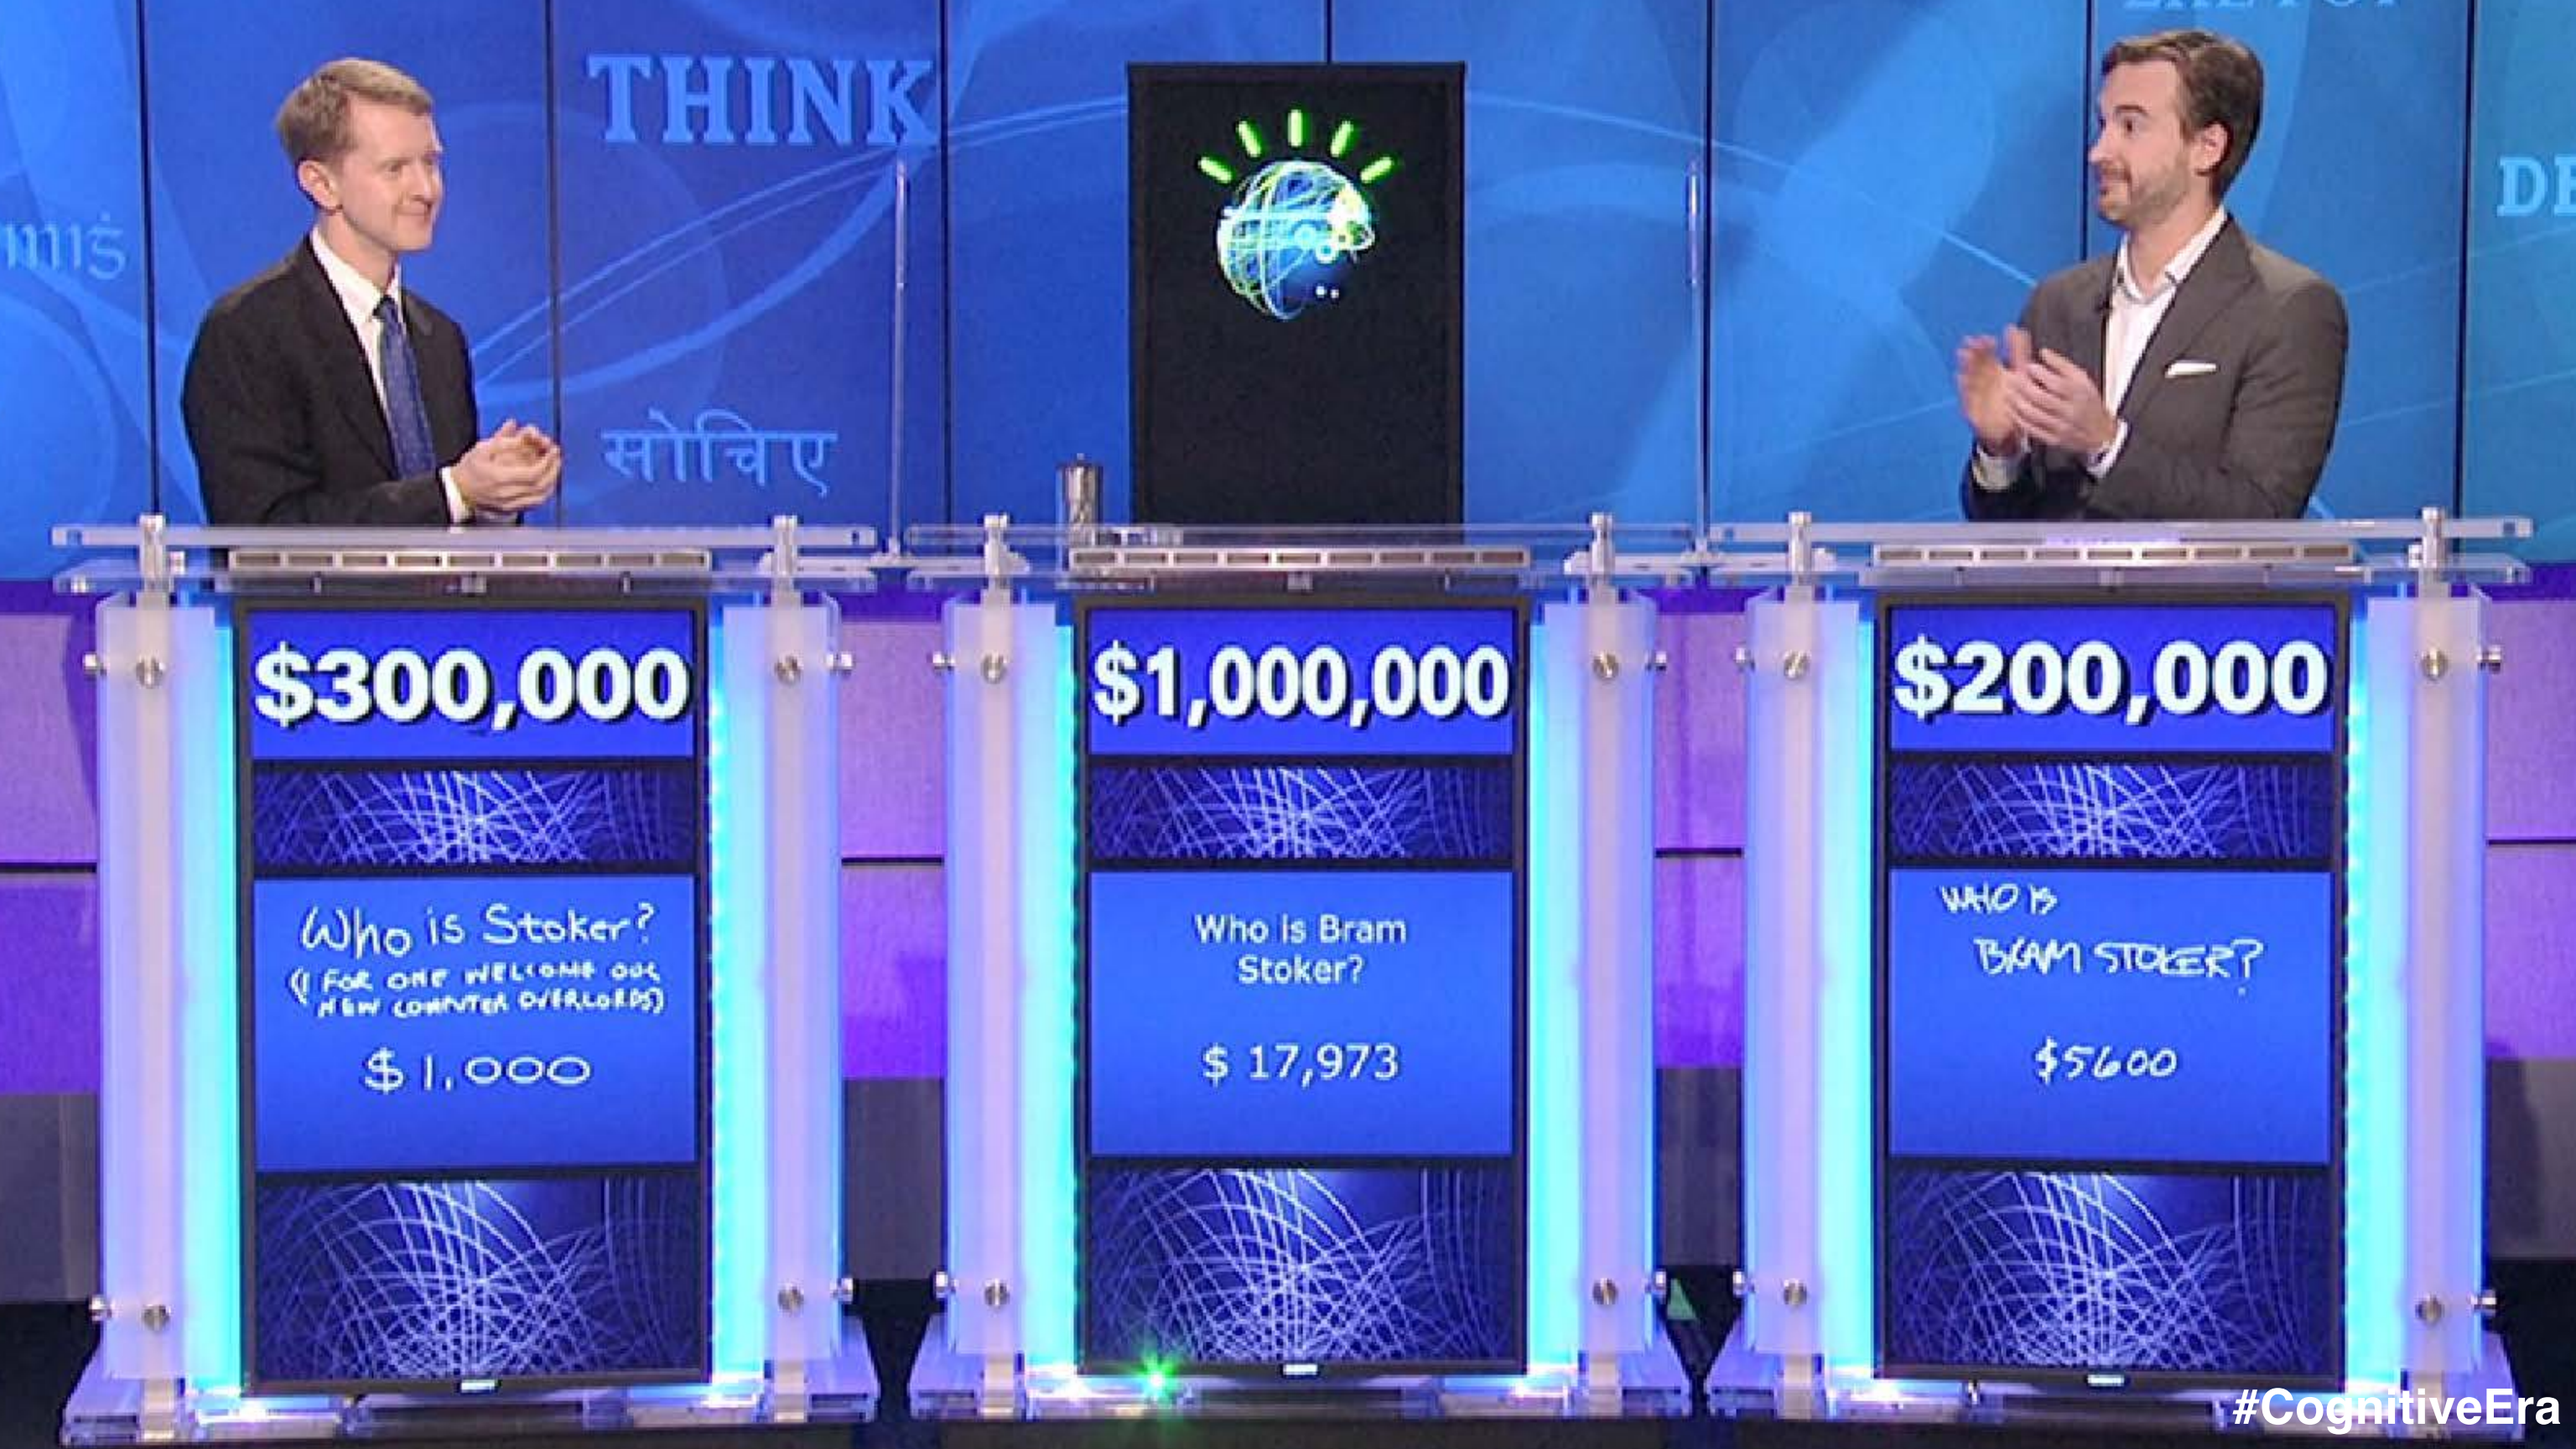

#CognitiveEra

Intuition

Design

Value judgements

Common sense

Deep Learning

Discovery

Large scale math

Fact checking

#CognitiveEra

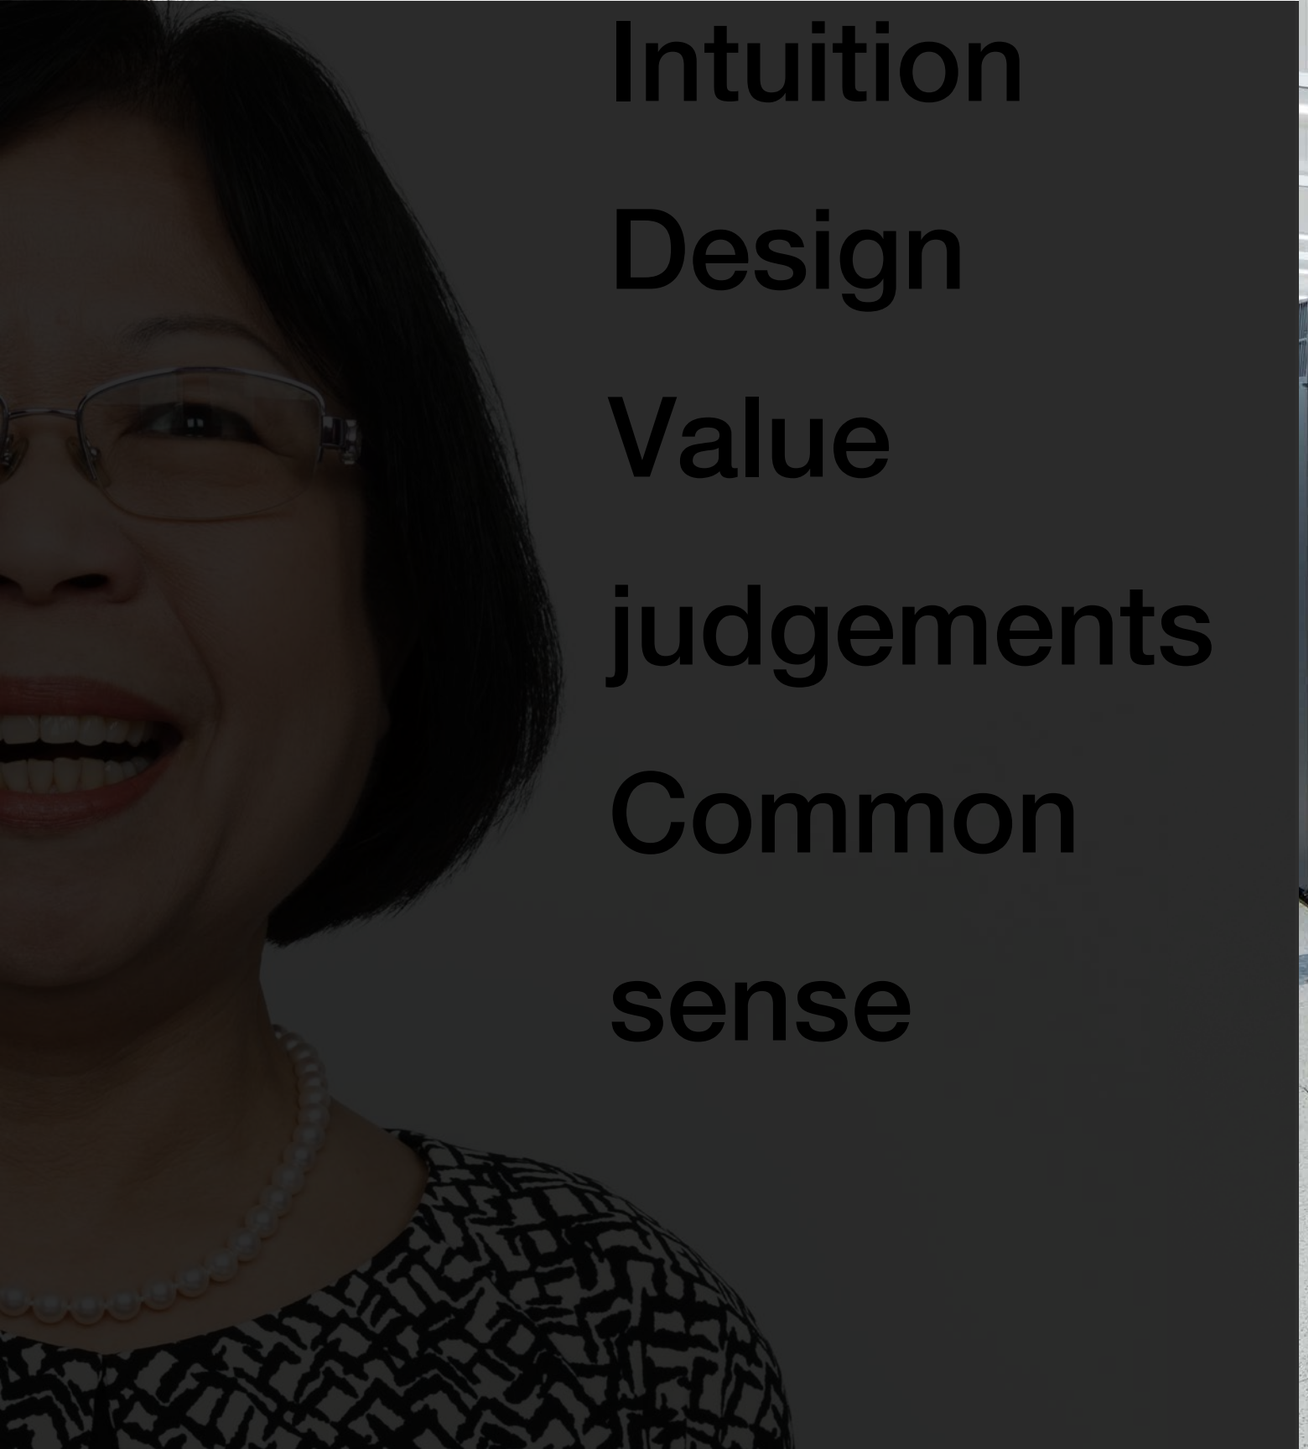

Intuition

Design

Value judgements

Common sense

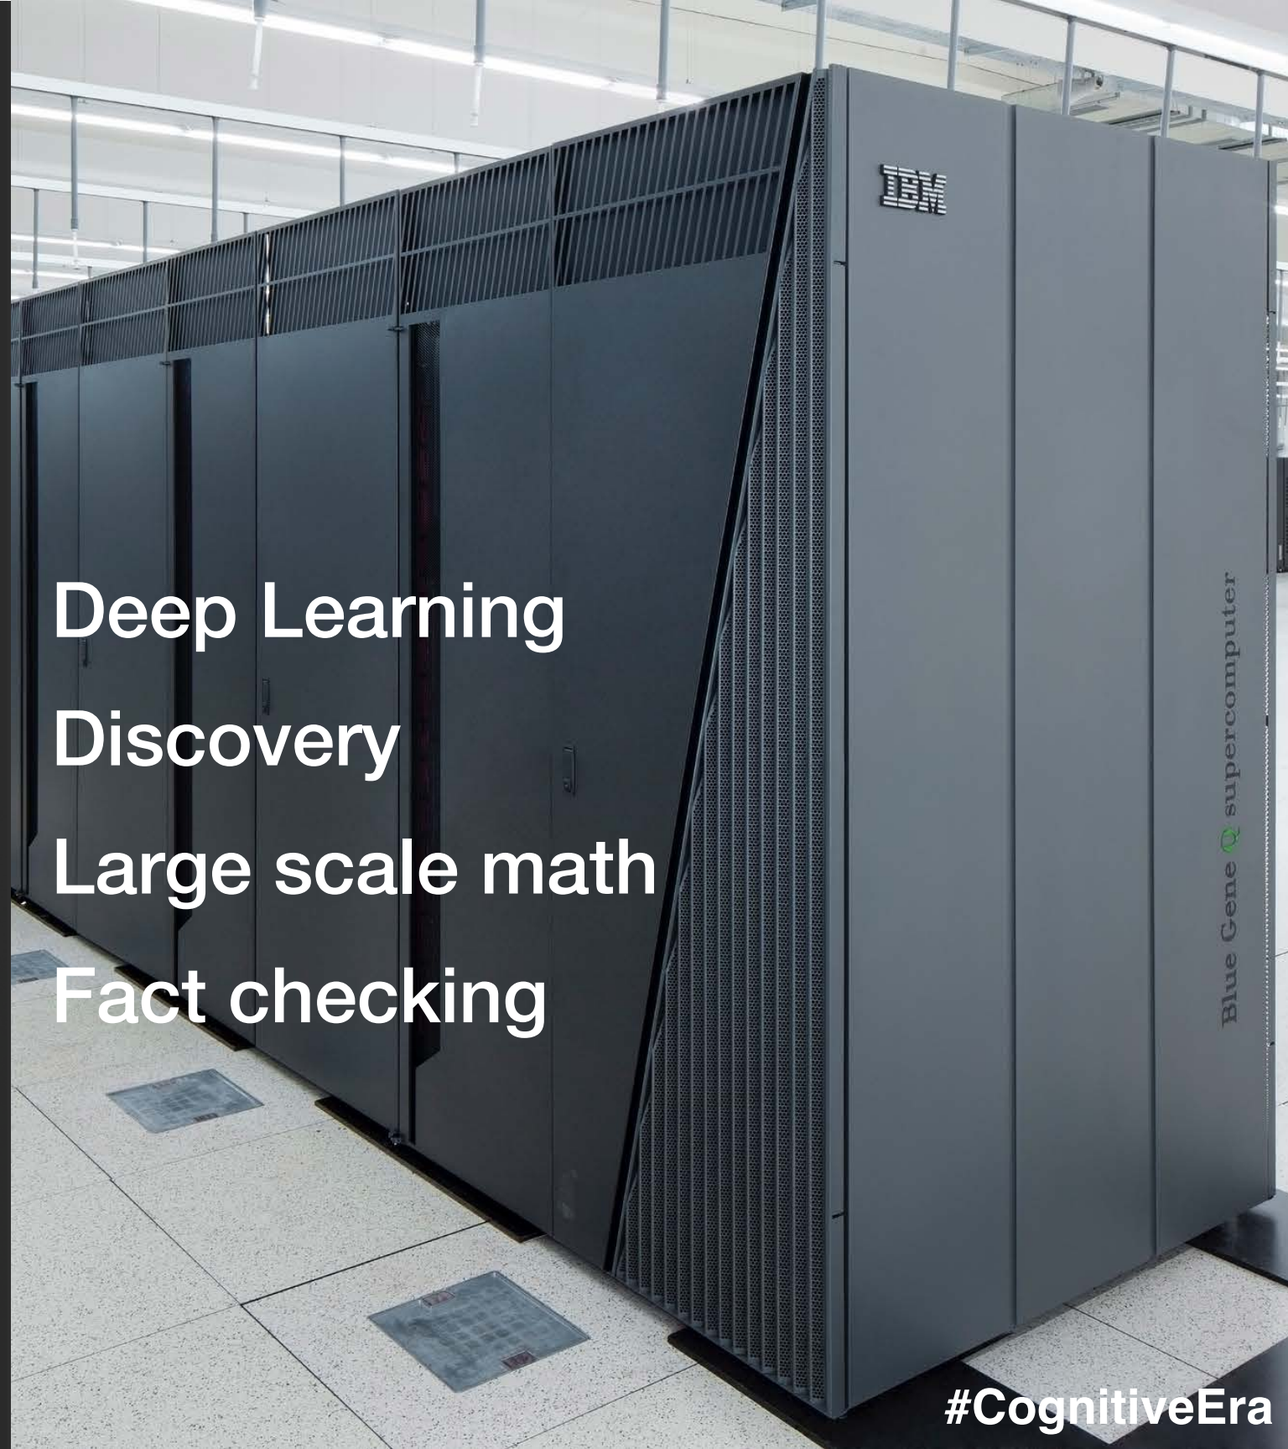

Deep Learning

Discovery

Large scale math

Fact checking

#CognitiveEra

Human + Machine

#CognitiveEra

The Watson that competed on *Jeopardy!* in **2011** comprised what is now a single API—**Q&A**—built on **five underlying technologies**.

Natural Language Processing

Machine Learning

Question Analysis

Feature Engineering

Ontology Analysis

Since then, Watson has grown to a family of **28 APIs**.

By the end of 2016, there will be nearly **50 Watson APIs**—with more added every year.

The Watson that competed on *Jeopardy!* in **2011** comprised what is now a single API—**Q&A**—built on **five underlying technologies**.

Since then, Watson has grown to a family of **28 APIs**.

By the end of 2016, there will be nearly **50 Watson APIs**—with more added every year.

Natural Language Processing
Machine Learning
Question Analysis
Feature Engineering
Ontology Analysis

#CognitiveEra

Cognitive systems must learn at scale, reason with purpose, and interact with humans naturally.

#CognitiveEra

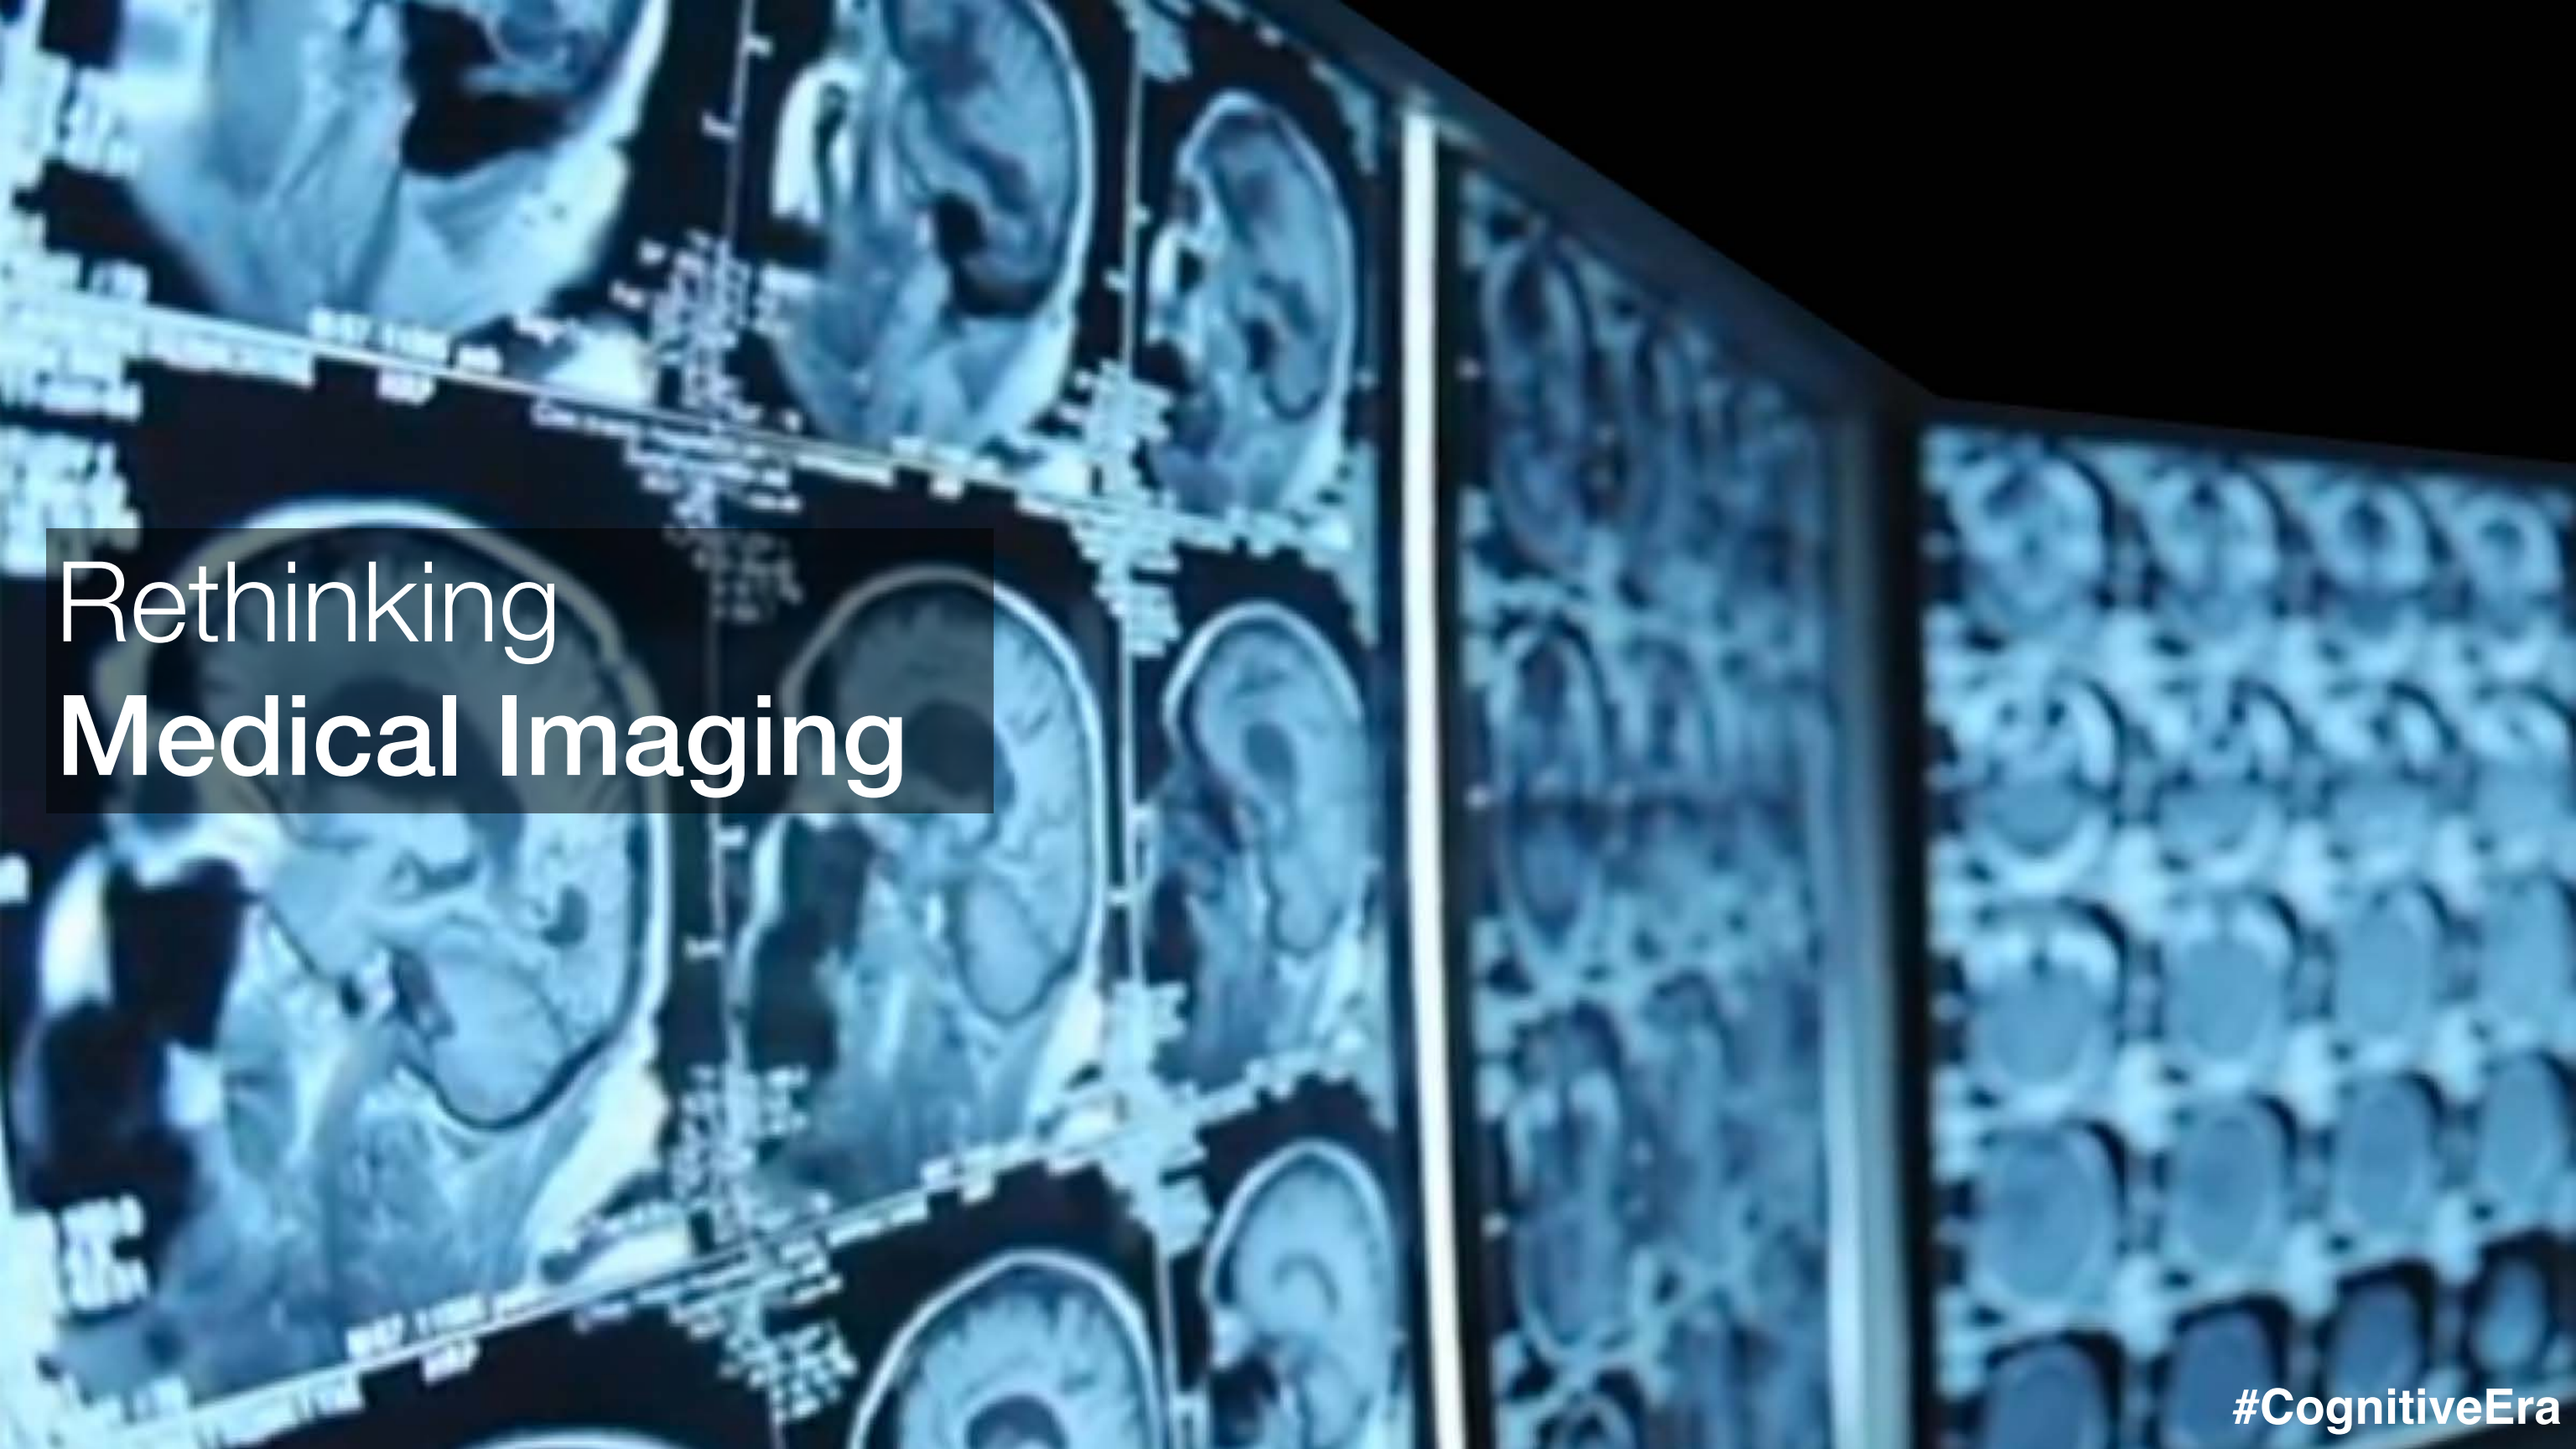
Rethinking
Medical Imaging
#CognitiveEra


Rethinking
Seismological
Exploration
#CognitiveEra

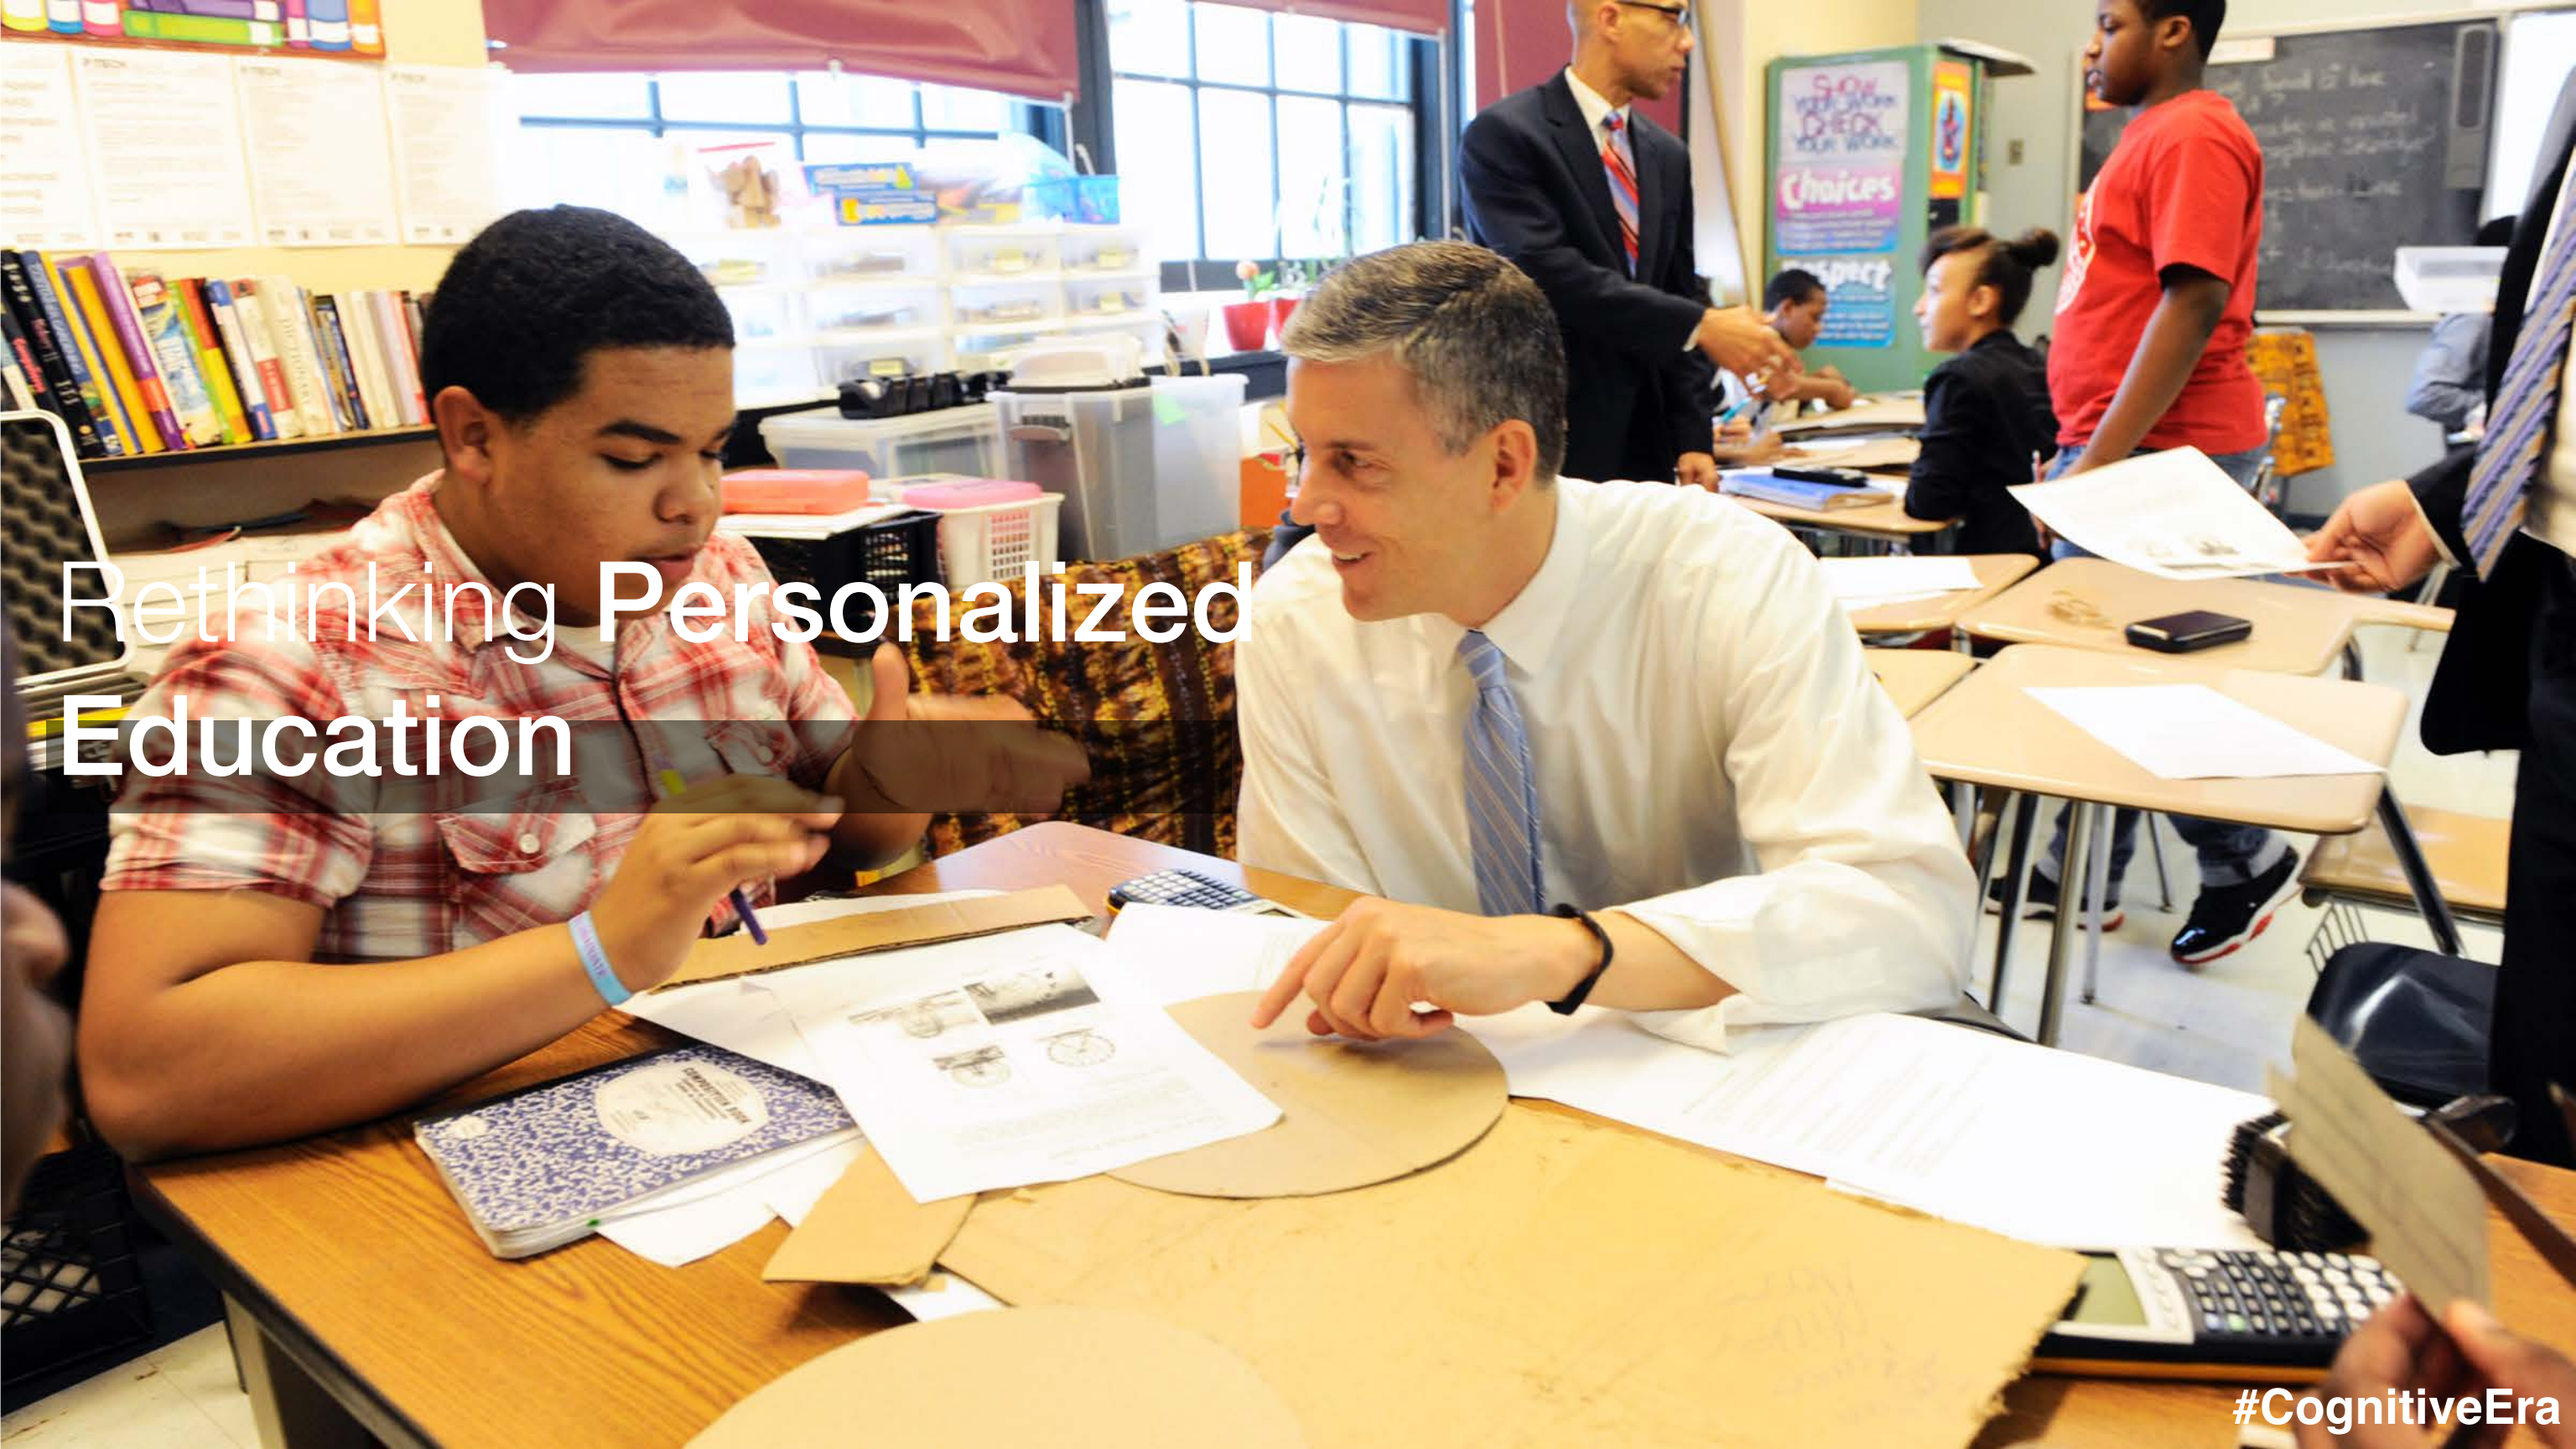

Rethinking **Personalized Education**

#CognitiveEra

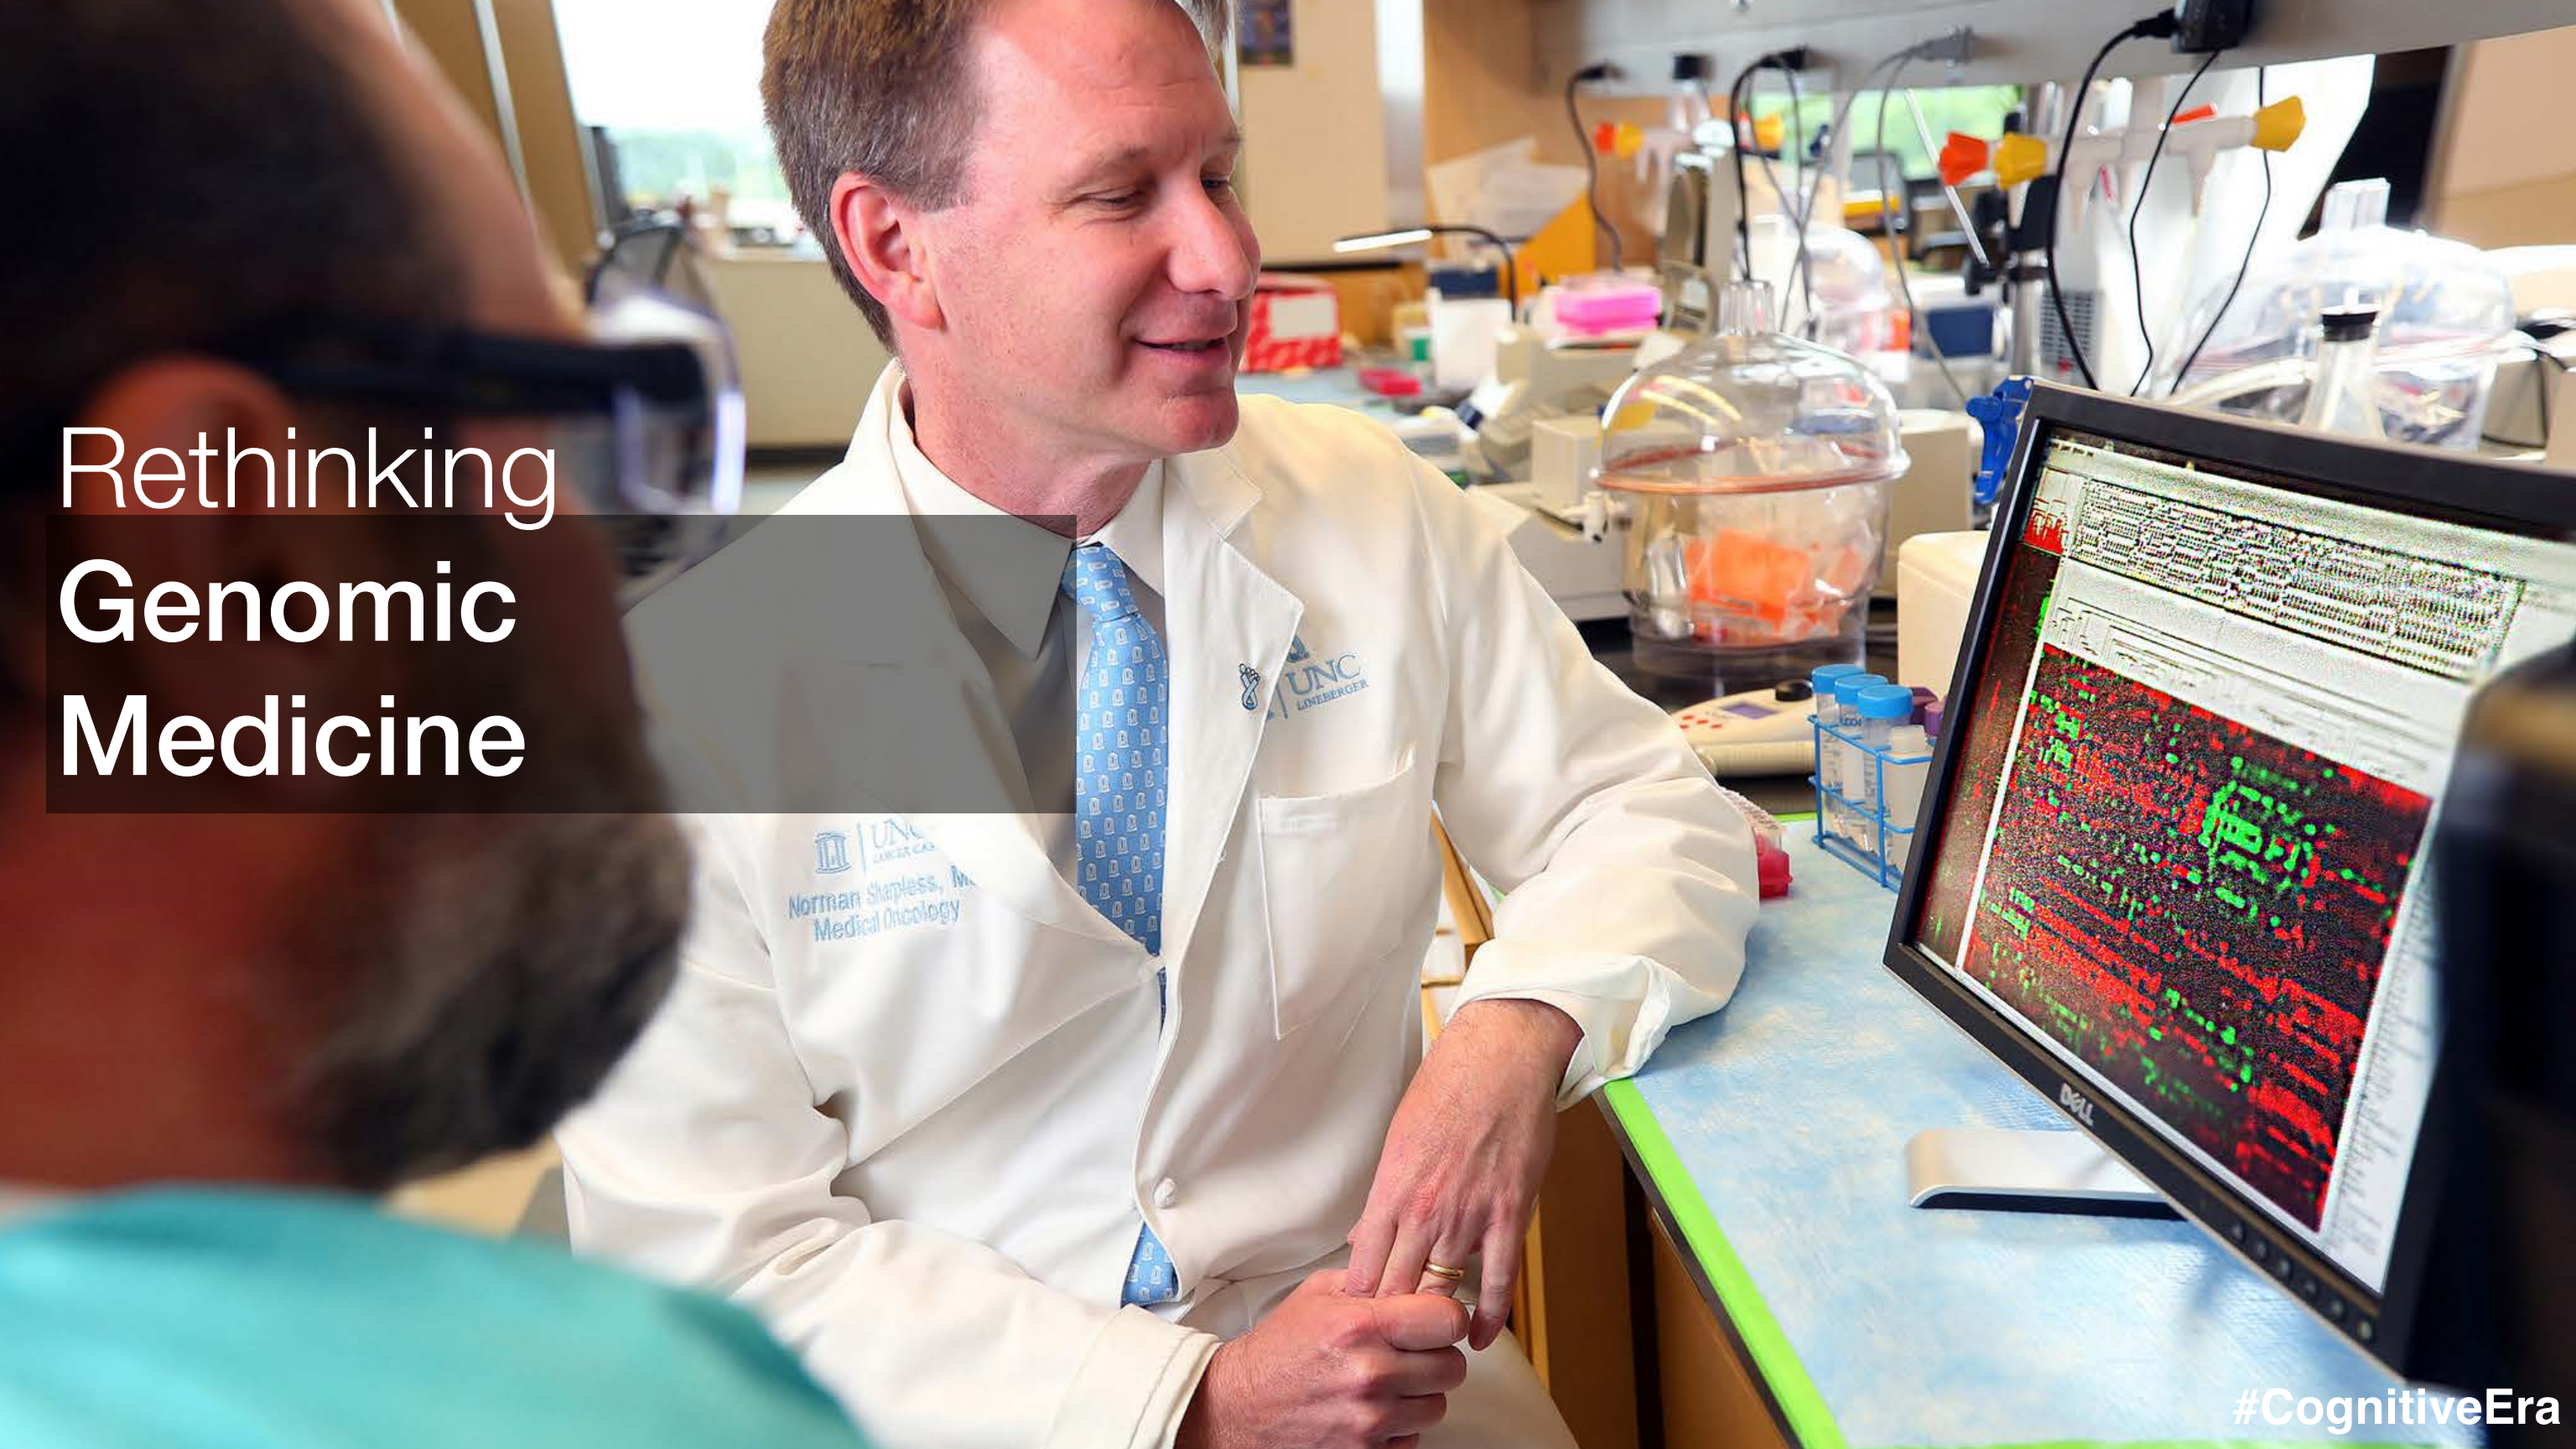
Rethinking
Genomic
Medicine
UNC
LINEBERGER
Norman Sharpless, M
Medical Oncology
DELL
#CognitiveEra

Trends, Shifts and Facts

500 million DVDs worth of data is generated daily

80% of the world's data is unstructured

1 trillion connected objects and devices by 2025

Data is becoming the world's new natural resource

85% of new software is being built for cloud

72% of developers say cloud-based services or APIs are central to the applications they are designing

25% of the world's applications will be available in the cloud by 2016

IT and business processes are being transformed into cloud services

80% of individuals are willing to trade their information for a personalized offering

84% of millennials say social and user-generated content influence what they buy

5 minutes response time users expect once they have contacted a company via social media

Social, mobile and access to data have changed how individuals are understood and engaged

Big Data: Analytics Multiplier and Gravity Effects

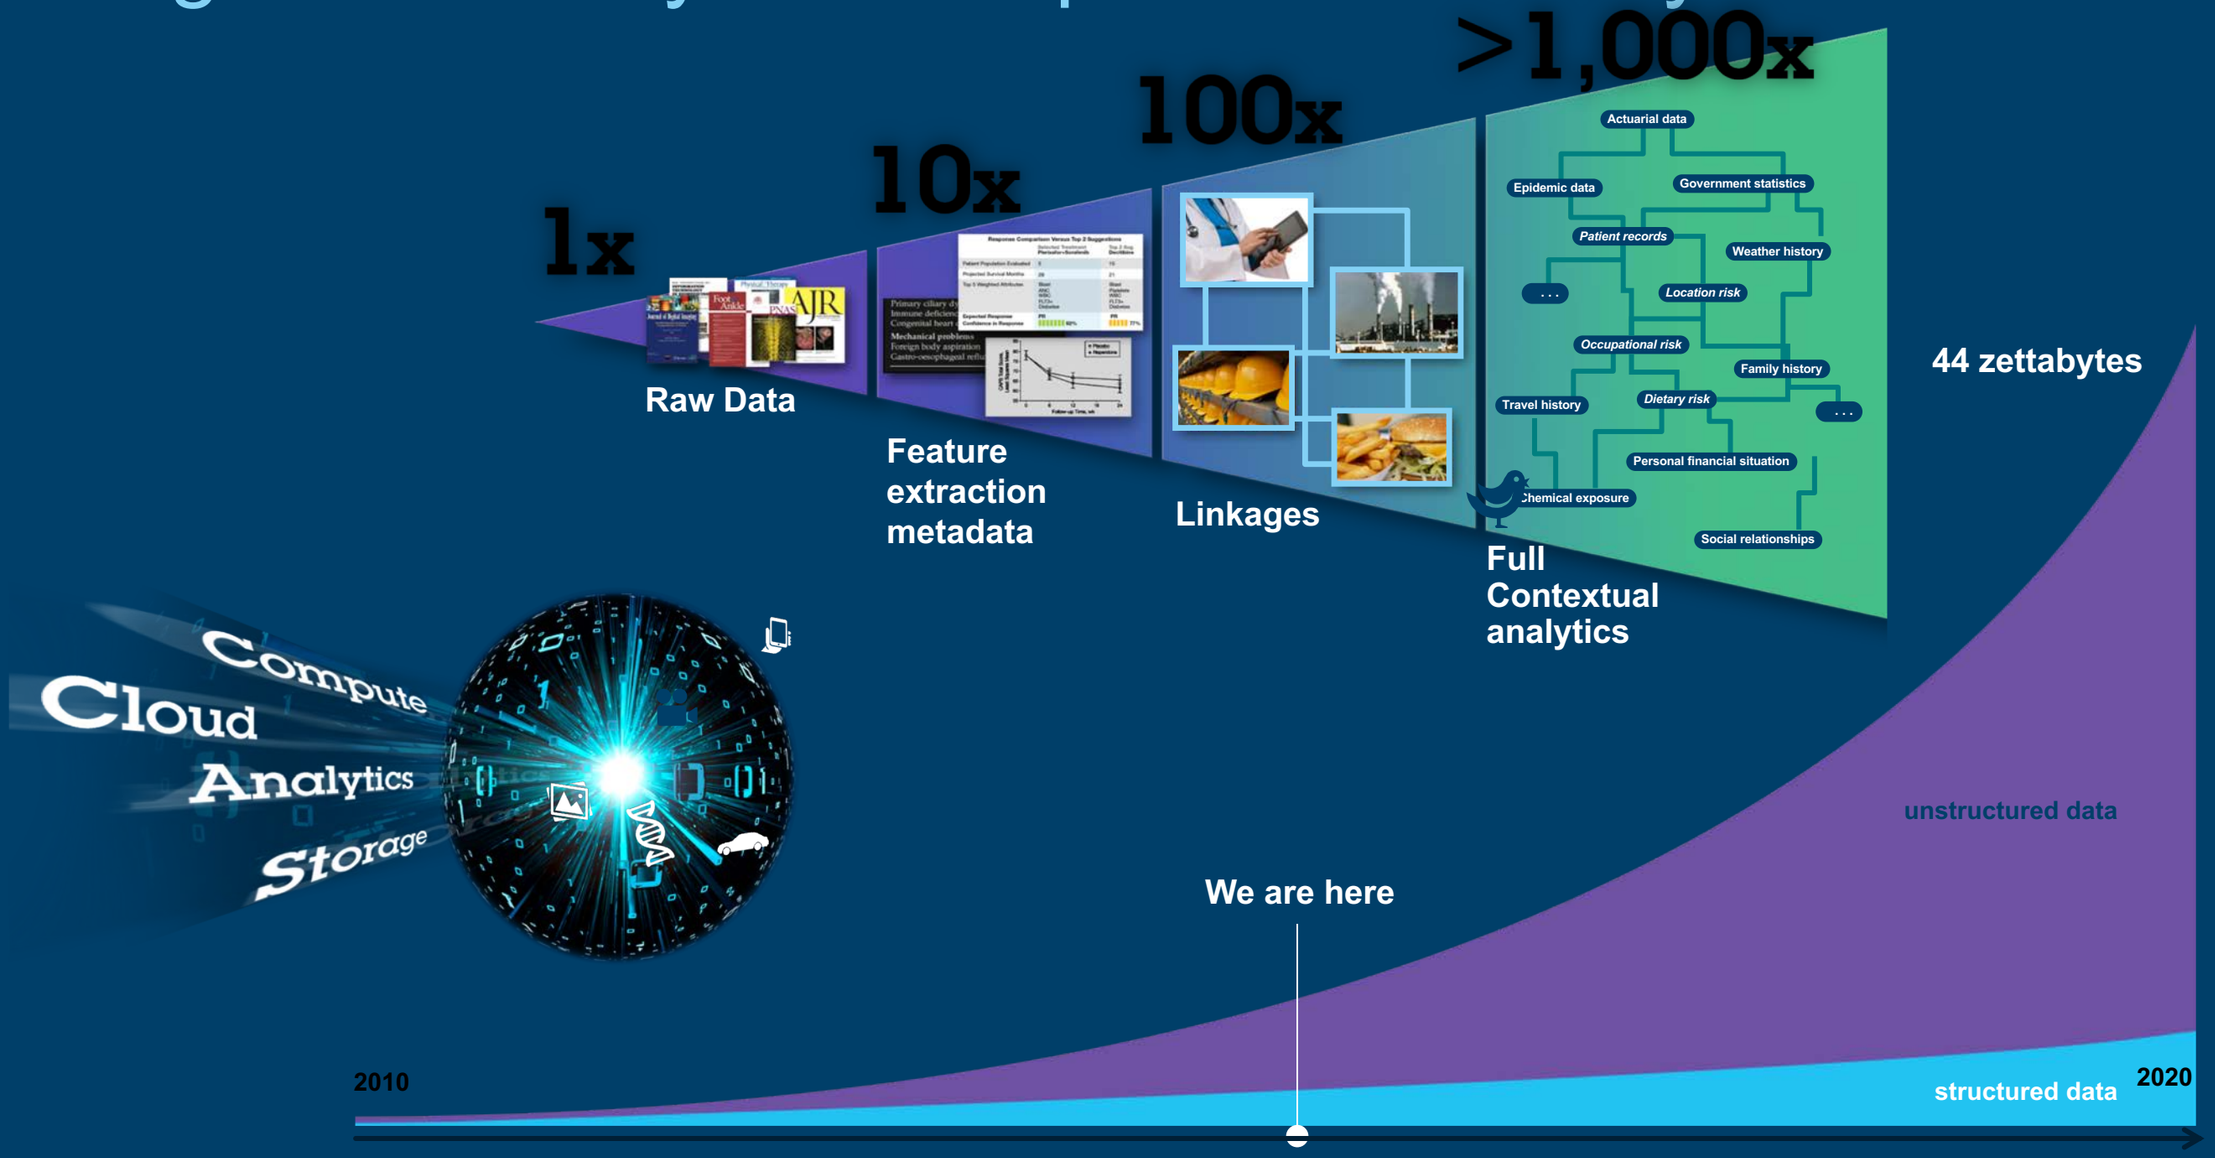

The Current Storage Model is being Disrupted by the Explosion of Data and Need for Speed

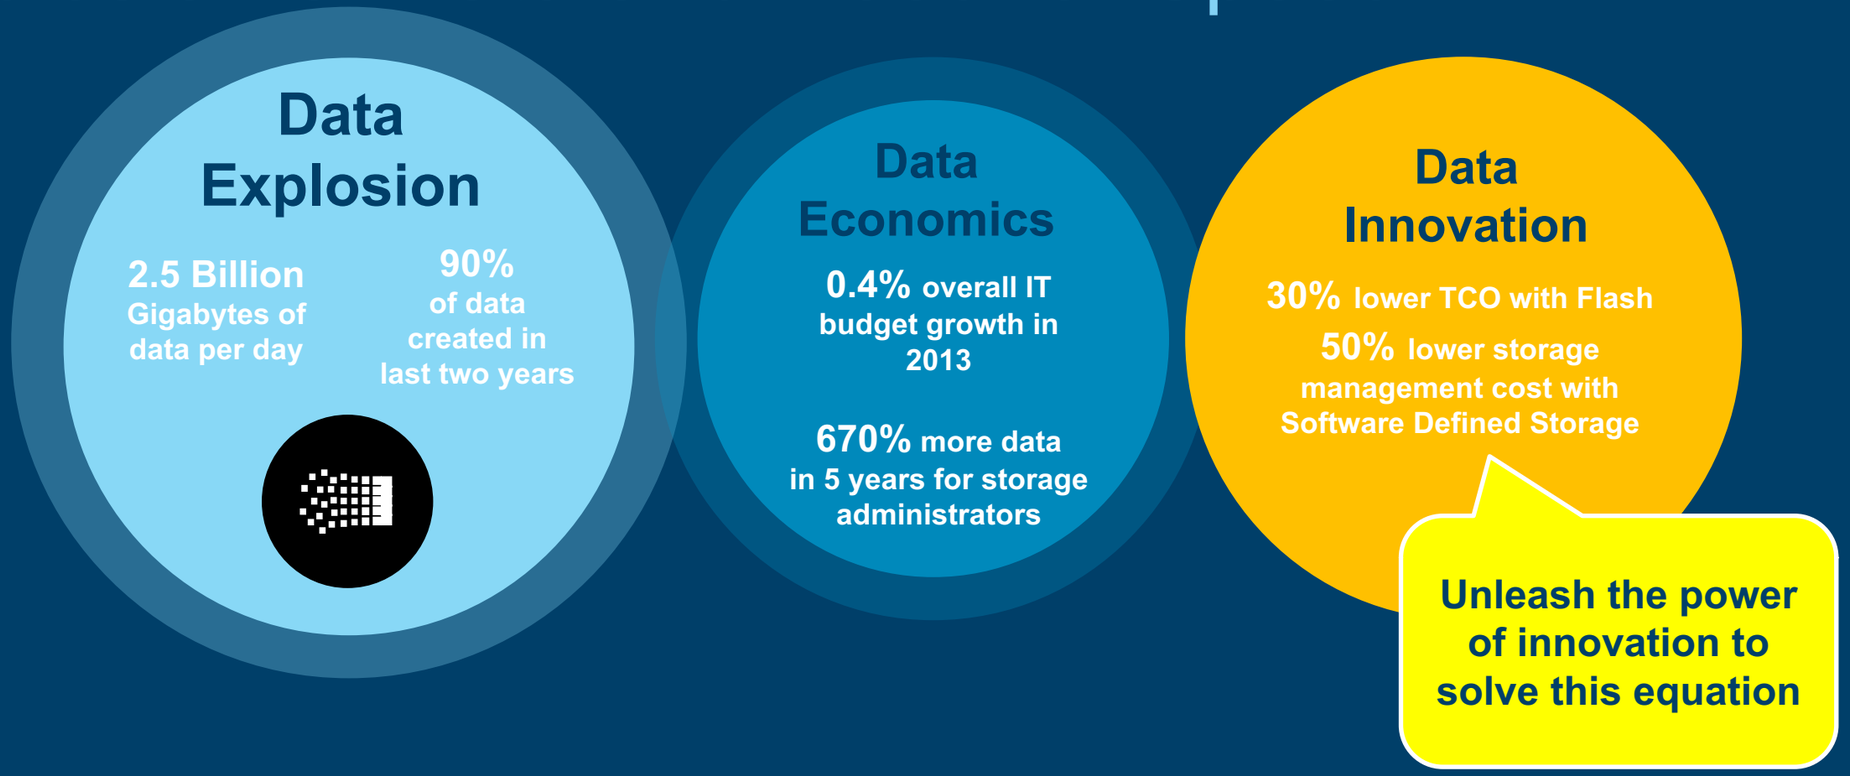

The top two challenges organizations face with IT infrastructure are storage related – Data Management and Cost Efficiency

SOURCE: *2014 IBM Institute for Business Value Study on Infrastructure Matters; Gartner IT Metrics

Systems Evolution to Meet New Workload Requirements

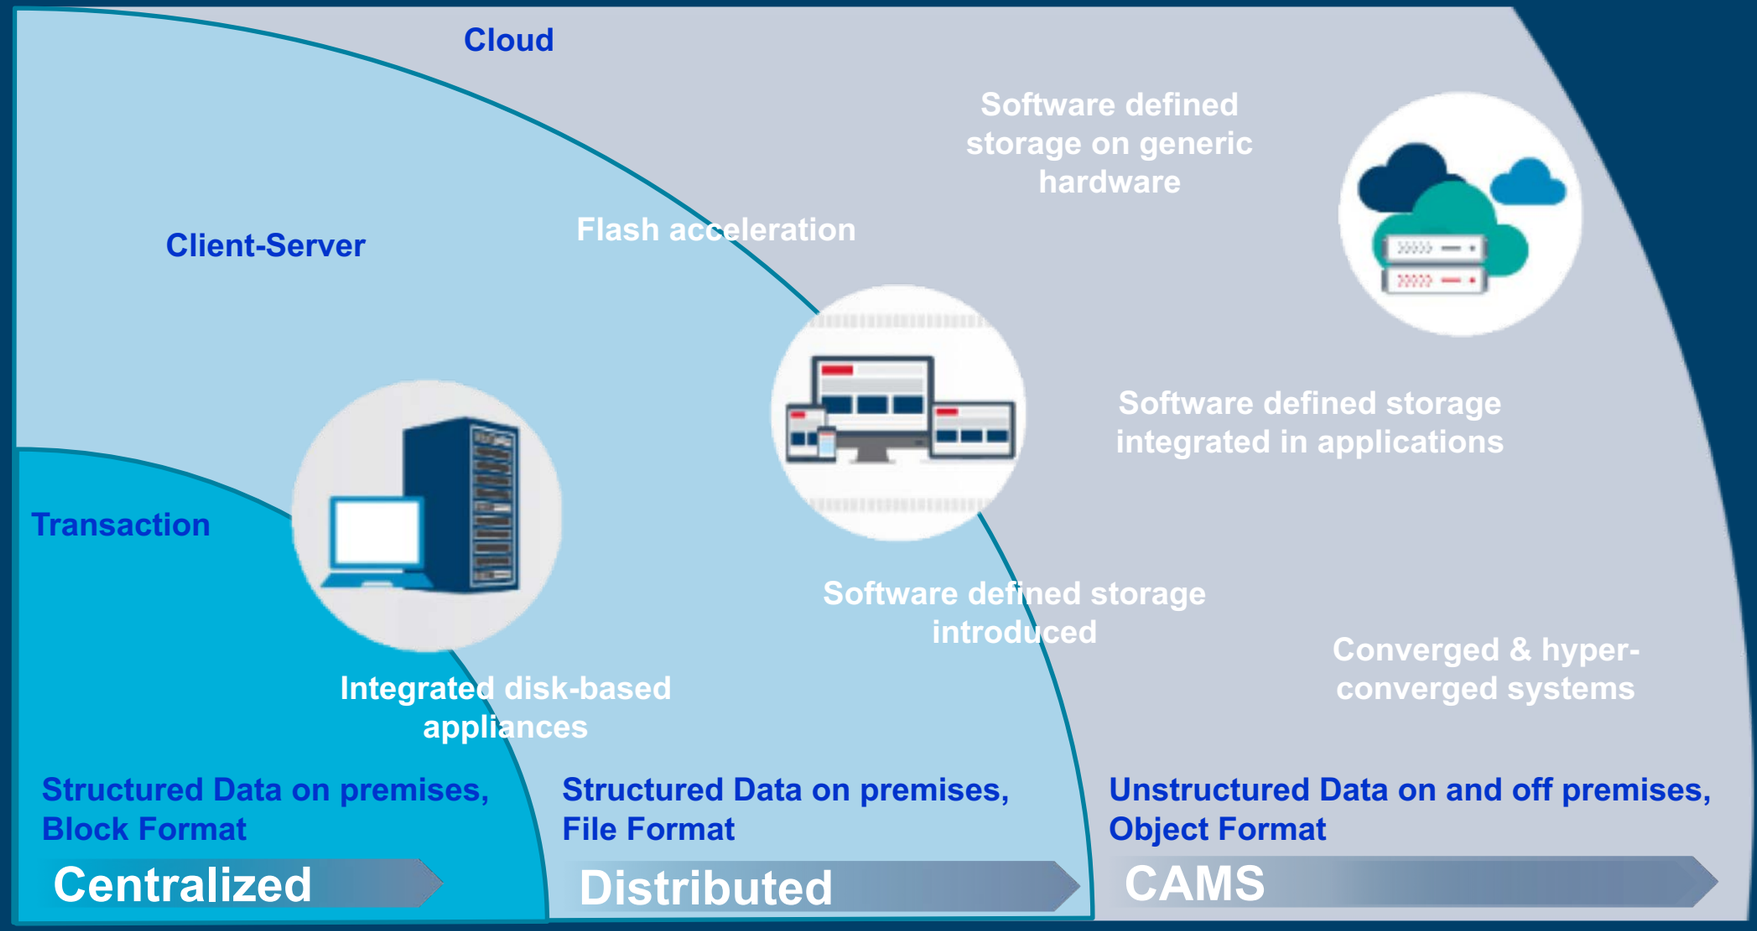

SDS means applications drive their storage, and storage turns into software on generic hardware

Expected Benefits of SDS – Client Survey

SDS is seen as the architecture model for storage innovation

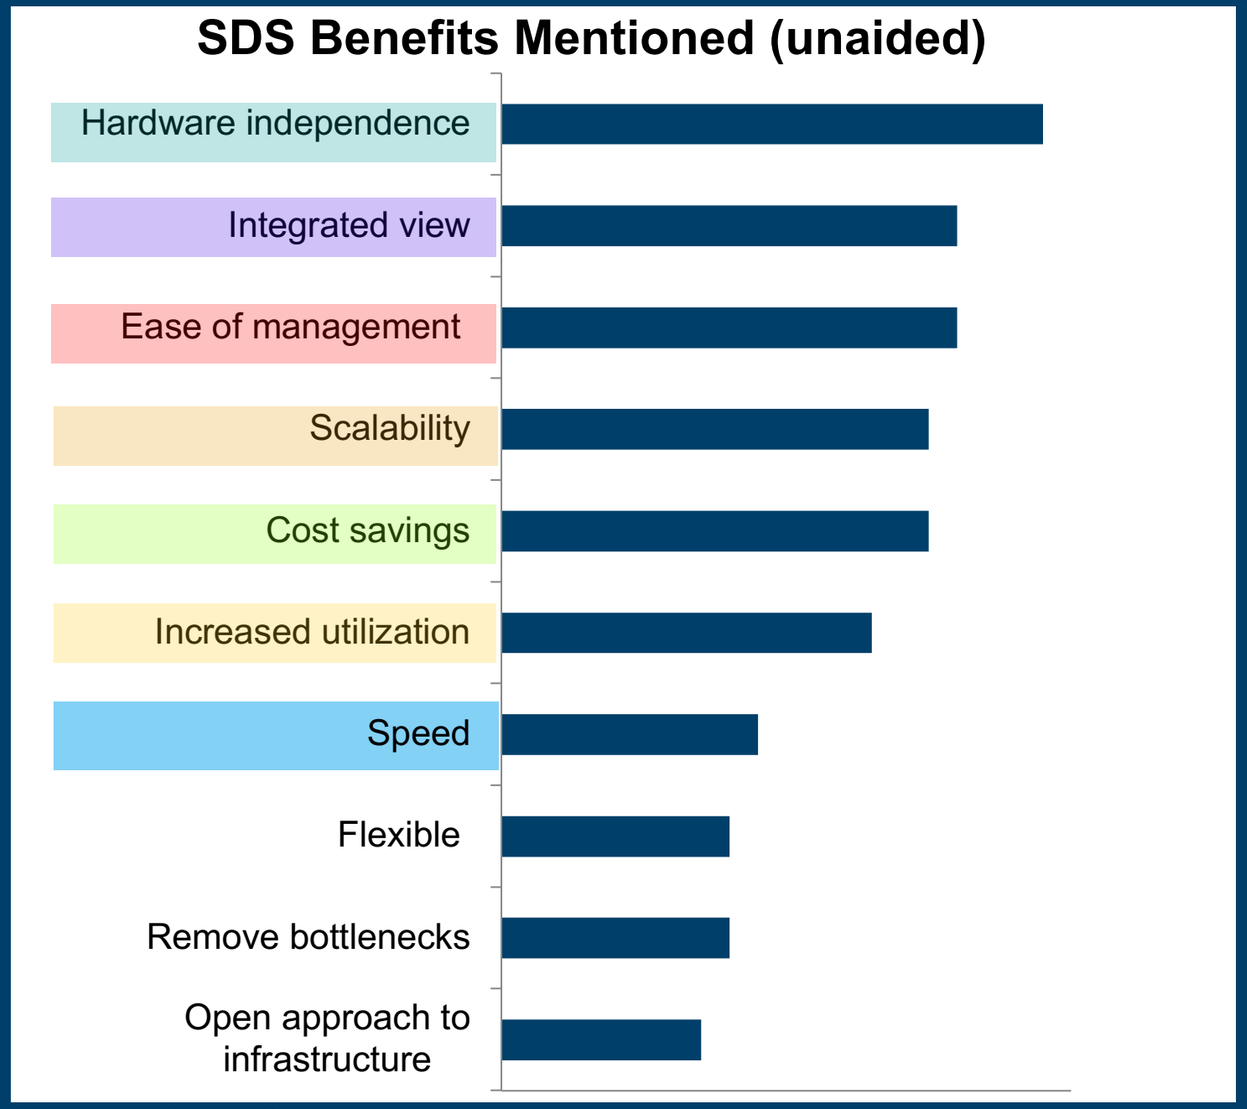

Software Defined Storage - Client Survey Findings

35

Software-defined Storage Business Value

- Delivered as an appliance, as software, or as a cloud service – with seamless paths between them to enable hybrid clouds

Simplified storage management	Scalability with data anywhere	Improved data economics	Openness
Unified and streamlined storage management and data protection across all applications and data types, wherever that data is stored	Elastic scalability with high performance for new analytics, big data, social and mobile applications. This includes unifying silos to deliver data without borders, securely, with built-in hybrid cloud support	Leverages commodity hardware and intelligently moves data to the right location at the right time, from Flash for fast access to tape and cloud for the lowest cost tier	Supports industry standards such as OpenStack and Hadoop to ensure the ability to complement vendor innovations with ones from other providers and communities

Exploring IBM Software Defined Storage Capabilities
IBM Spectrum Storage Portfolio

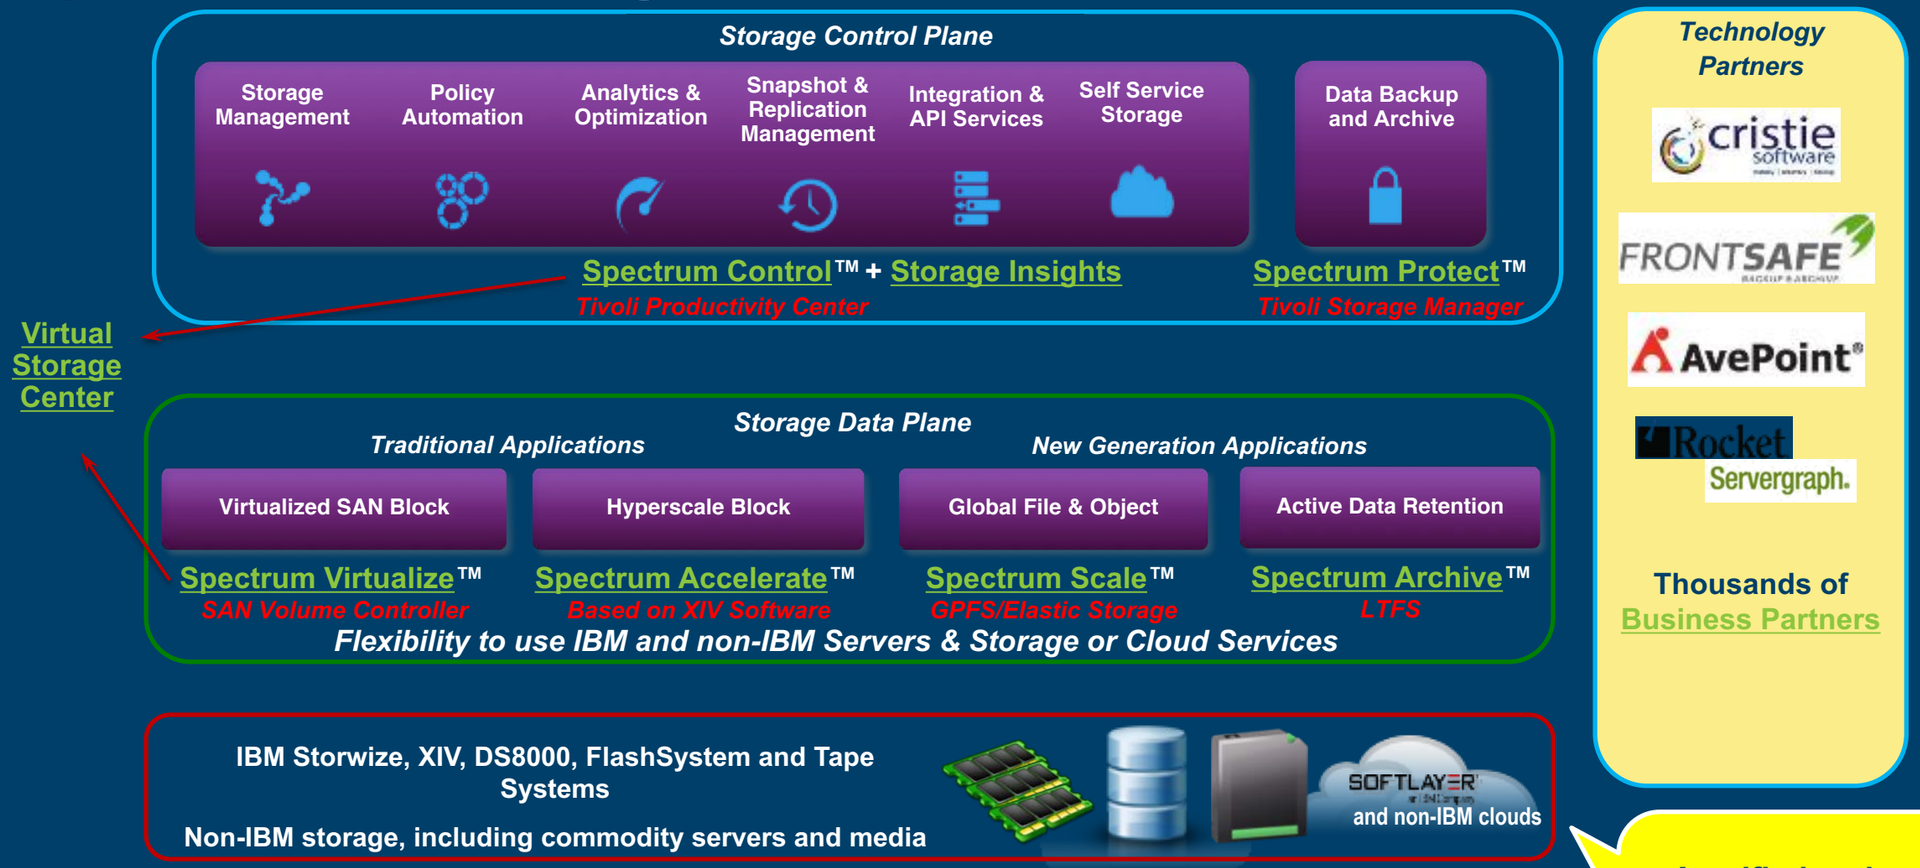

Technology Shifts among Storage Tiers

Primary data

Flash Arrays & Solid-State Drives (SSD)

Secondary & Backup data

Disk and VTL

Active Archive Nearline Objects

Slow/Cheap Disk

Optical Libraries

Large Objects Offline Archive Cold Data Store

High Capacity Tapes + Linear Tape File System (LTFS)

Storage Innovation Themes

- **Flash**-based solid-state storage
  - Flash density increasing, $/GB dropping, but performance & endurance dropping as well
  - Enterprise all-Flash arrays address the low-latency, high-IOPS segment efficiently
  - What about the rest of the datacenter and cloud?
  - Can software treat the limitations of low-cost Flash?
- **Tape** has addressed, so far, the low TCO segment more efficiently than disk
  - Can tape maintain its substantial cost advantage over disk?
  - How can tape be integrated into the cloud for object store?
- **Traditional** storage systems are still mission-critical in the enterprise
  - How can traditional storage be cloud-enabled?
  - How can the cloud provide a reliable and secure backend for traditional systems?

Enterprise Flash

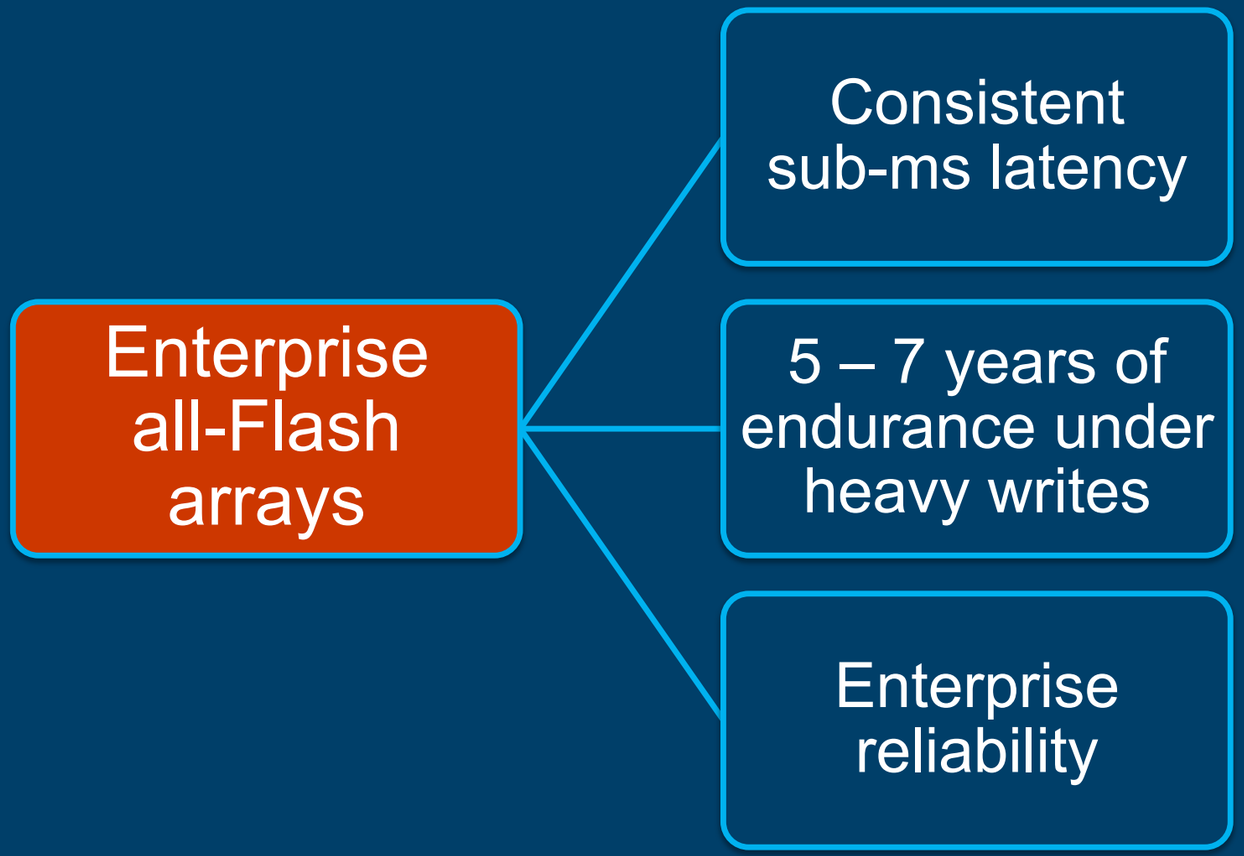

How?

- Use of MLC Flash
- Purpose-built hardware (FPGAs)
- Hardware-only data path
- Purpose-built ECC schemes
- Advanced Signal Processing
- State-of-the-art Garbage Collection
- Fine-grained Health Management

Endurance Improvement

Gains due to innovations in hardware
(ECC, Signal Processing, etc.)

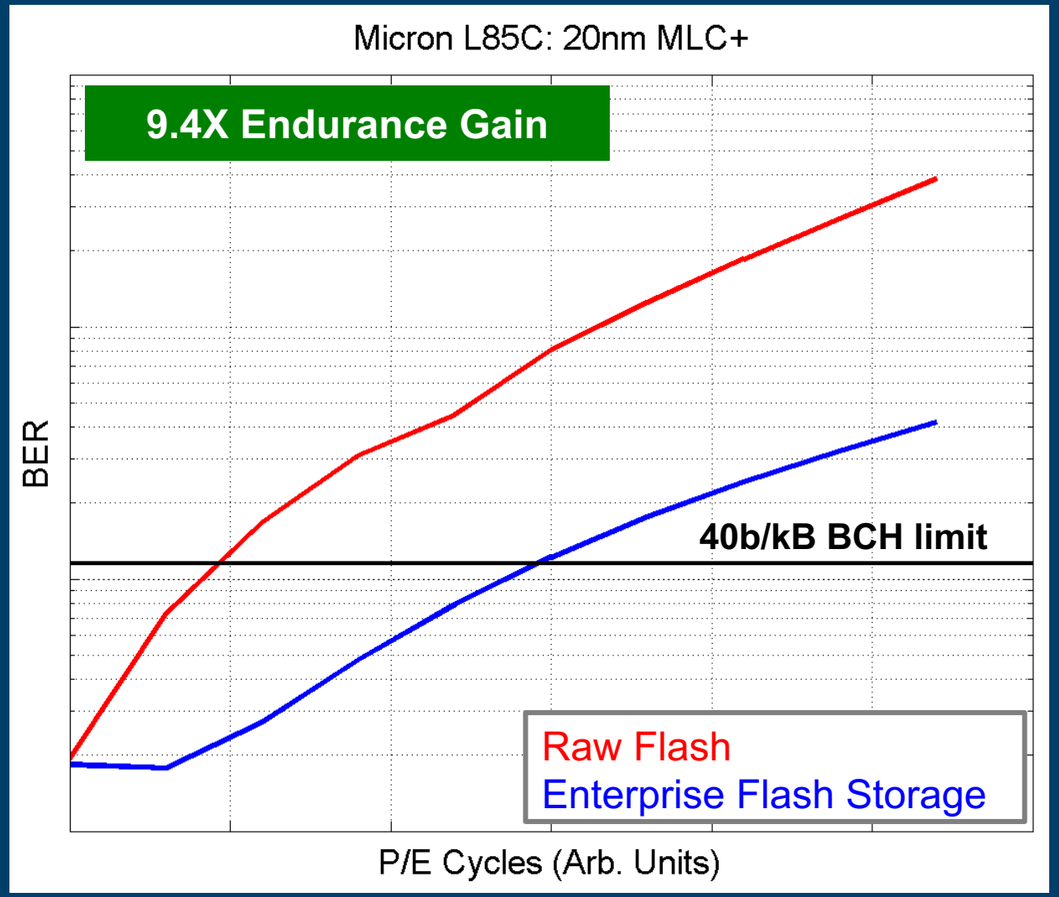

Gains due to innovations in FTL
(GC, Wear-leveling, etc.)

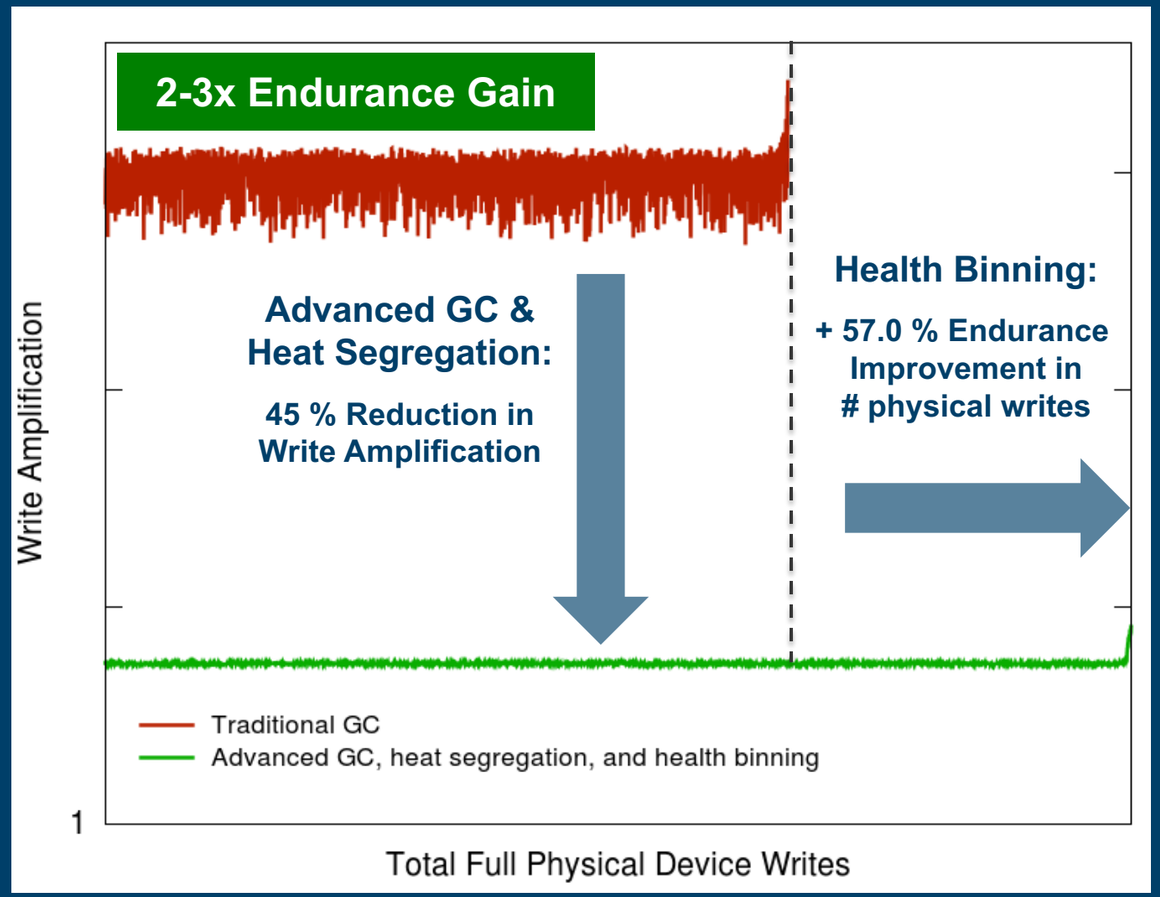

Dramatic endurance improvement, millions of IOPS, 100-200usec consistent latency!

New Category of **Big Data Flash**

- Many workloads **do not really need the write performance** and **endurance** of "good" Flash
  - In certain environments data actually is immutable
- What matters is **high density, low cost,** and **good read performance**
  - Current Flash architectures are not a good fit

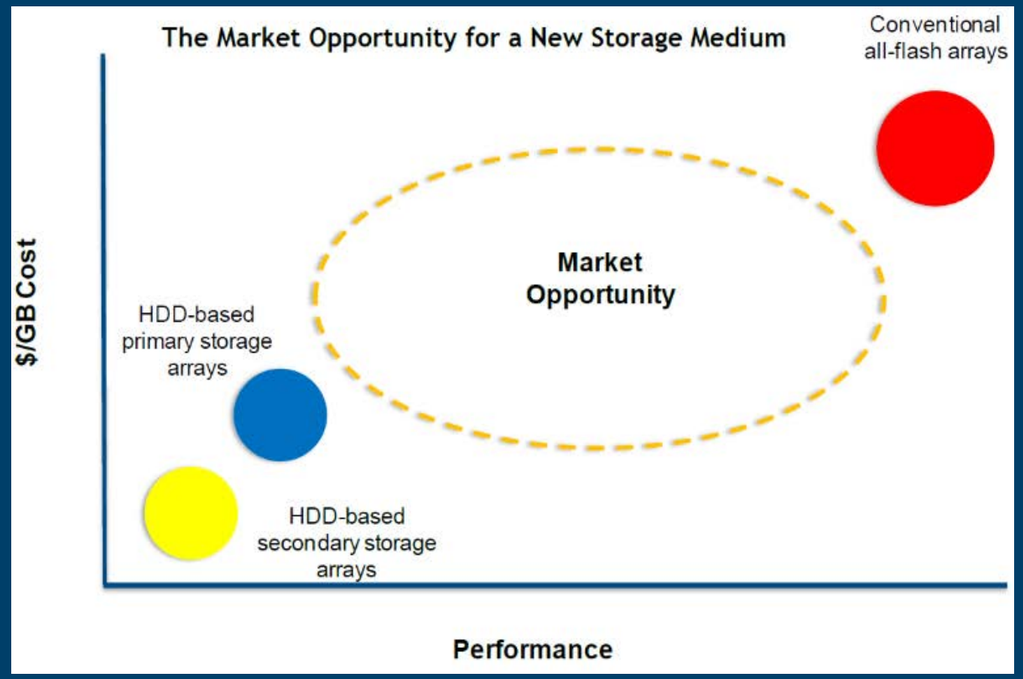

eBay: "*We could live with 1/3rd the number of writes that normal flash supports as long as we could get it for 1/4th the price.*"

- IDC just introduced a new market category **of Big Data Flash** (March 2015)
- Content repositories, media and streaming services, Big Data and analytics, NoSQL, Object storage, Web infrastructure

At < 1 $/GB for Flash, total acquisition cost becomes the same as an HDD-based solution, with much lower TCO.
- IDC

Low-cost Flash technology (c-MLC, TLC)

Can't we just use low-cost SSDs?

- Low-cost Flash suffers from high write latency, low endurance
  - For example, TLC, 3D-NAND, c-MLC
- Low-cost SSDs have limited resources, simple controllers to keep the cost as low as possible (~ $0.4/GB!)
- Therefore, they only use simple Flash management
  - Sufficiently good read performance
  - But: **limited write endurance, terrible write performance**

Raw low-cost SSDs are **practically unusable** in a real datacenter

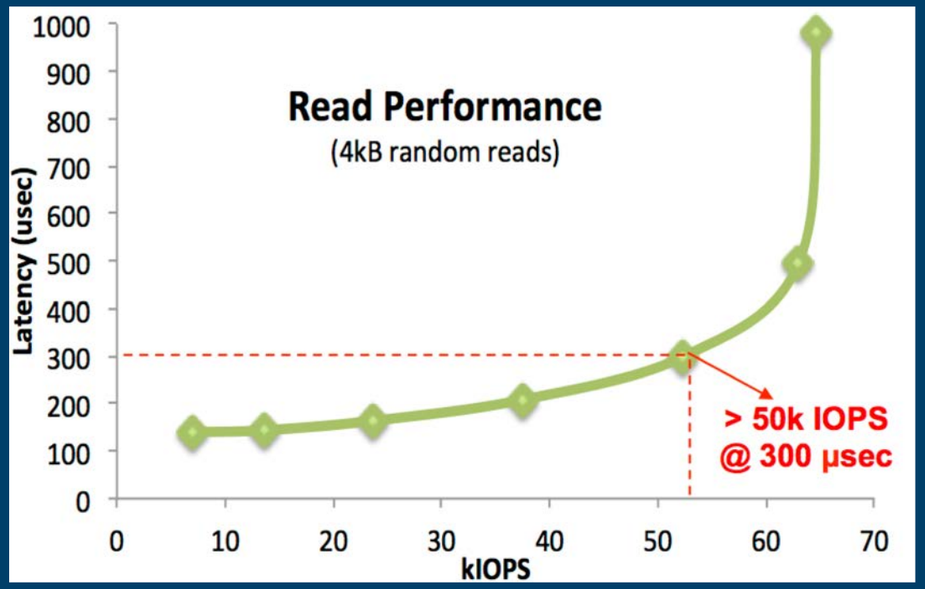

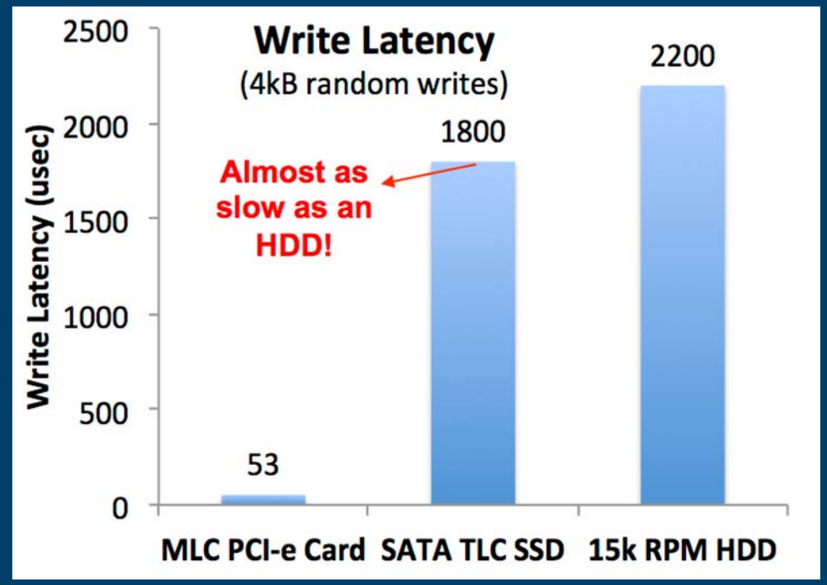

SoftwAre Log-Structured Array

What?
A **Flash-optimized** I/O stack that elevates the performance and endurance of consumer-level SSDs to enterprise standards.

Why?
Offer **cost-effective all-Flash** storage in public and private clouds, mainly for read-dominated workloads, complementing our high-end FlashSystem offerings.

How?
1. Use high-density, low-cost, off-the-shelf Flash SSDs
2. Move complexity from hardware to software to reduce cost
3. Optimize **end-to-end** for low Write Amplification
4. Employ aggressive Data Reduction
5. Natively support Object Storage

Squeeze the most capacity out of Flash

SALSA

- ✓ Implements the state-of-the-art Flash Management **in software**
- ✓ Exposes **standard interfaces**
  - File-systems and applications run unmodified on top of SALSA
- ✓ Is ideal for cost-optimized scale-out storage systems like **CEPH**
  - SALSA enables CEPH on low-cost SSDs, offering high performance and endurance

Experiments – Block Storage

- Using SALSA in a commodity Linux server to create an array out of 5 SSDs
  - With RAID5-equivalent parity protection
- Comparing against RAID0, RAID5 on the same SSDs

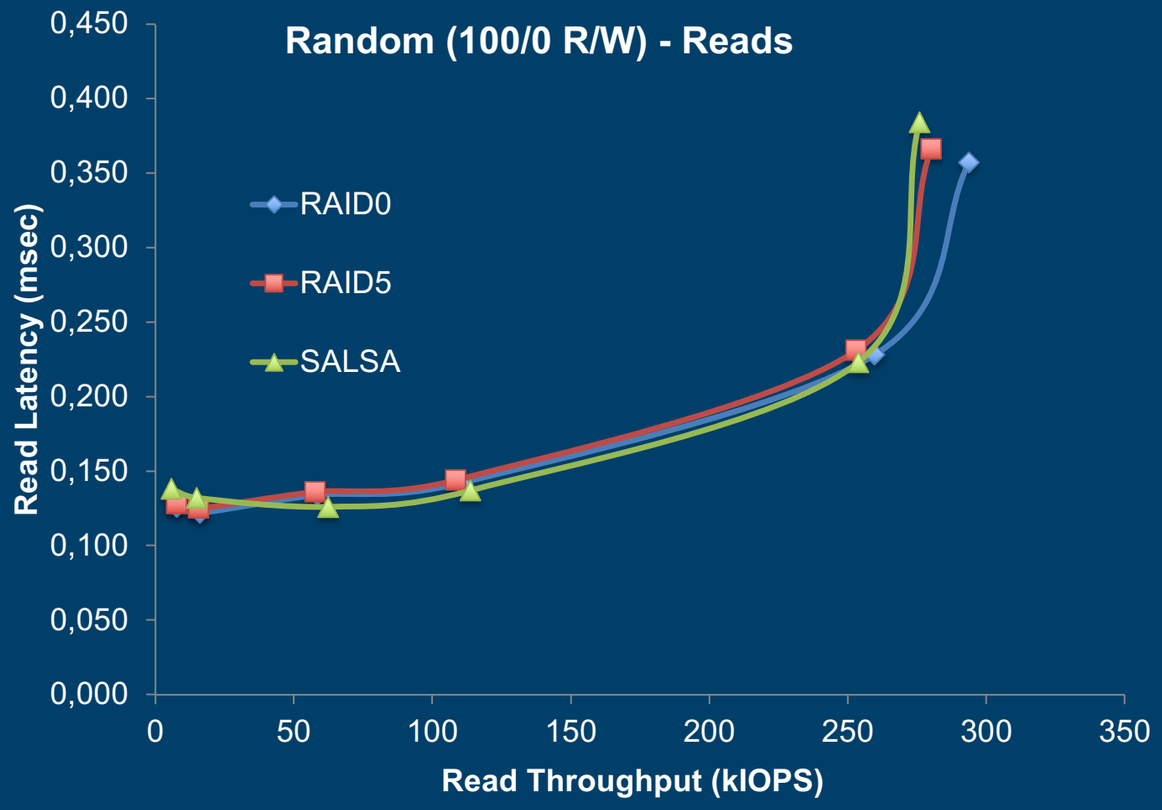

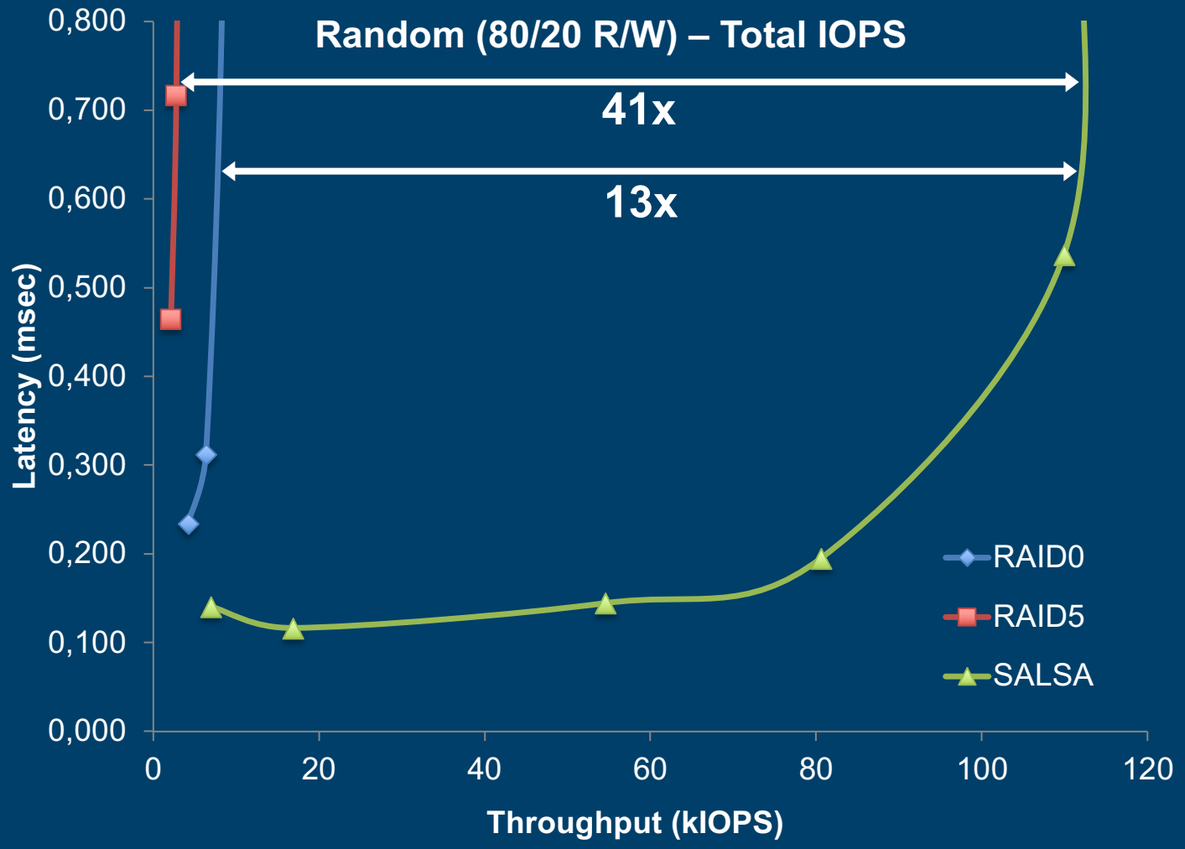

SALSA dramatically improves performance in the presence of writes

CEPH on SALSA

- 3-node x86 cluster
- 10 Gbit Ethernet network
- 2 x 1 TB TLC SSDs per node
- Replication factor of 3
- Mixed read/write random I/O

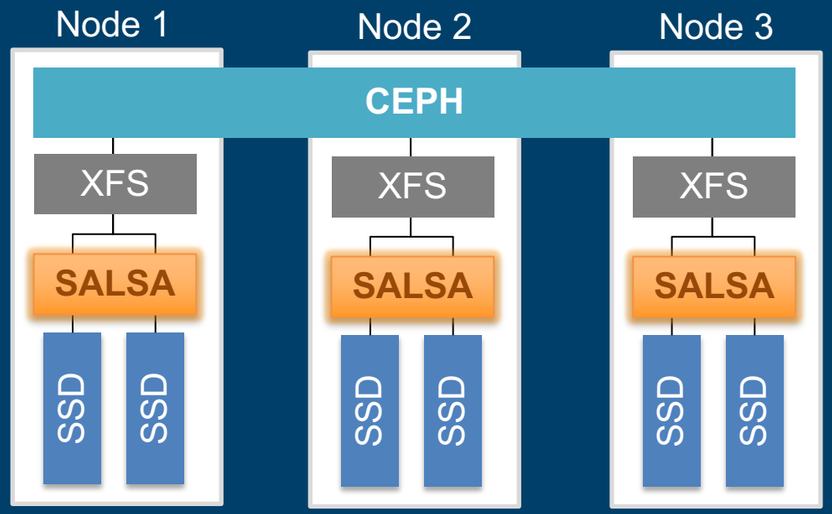

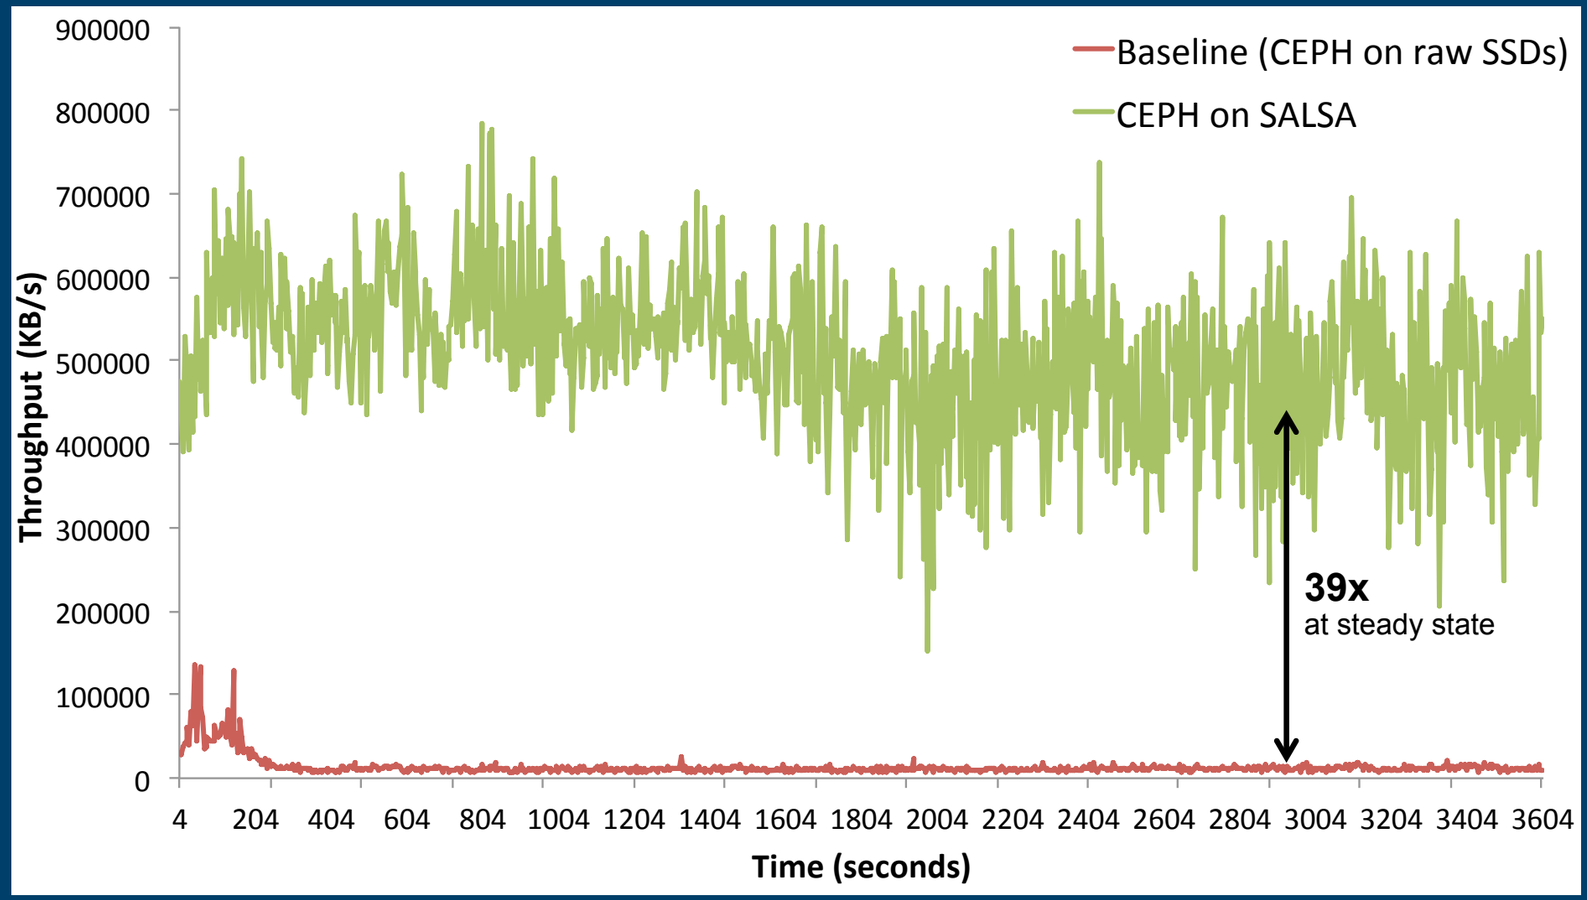

SALSA can enable CEPH on Flash with high performance at a low cost!

SDS using SALSA

using low-cost SSDs

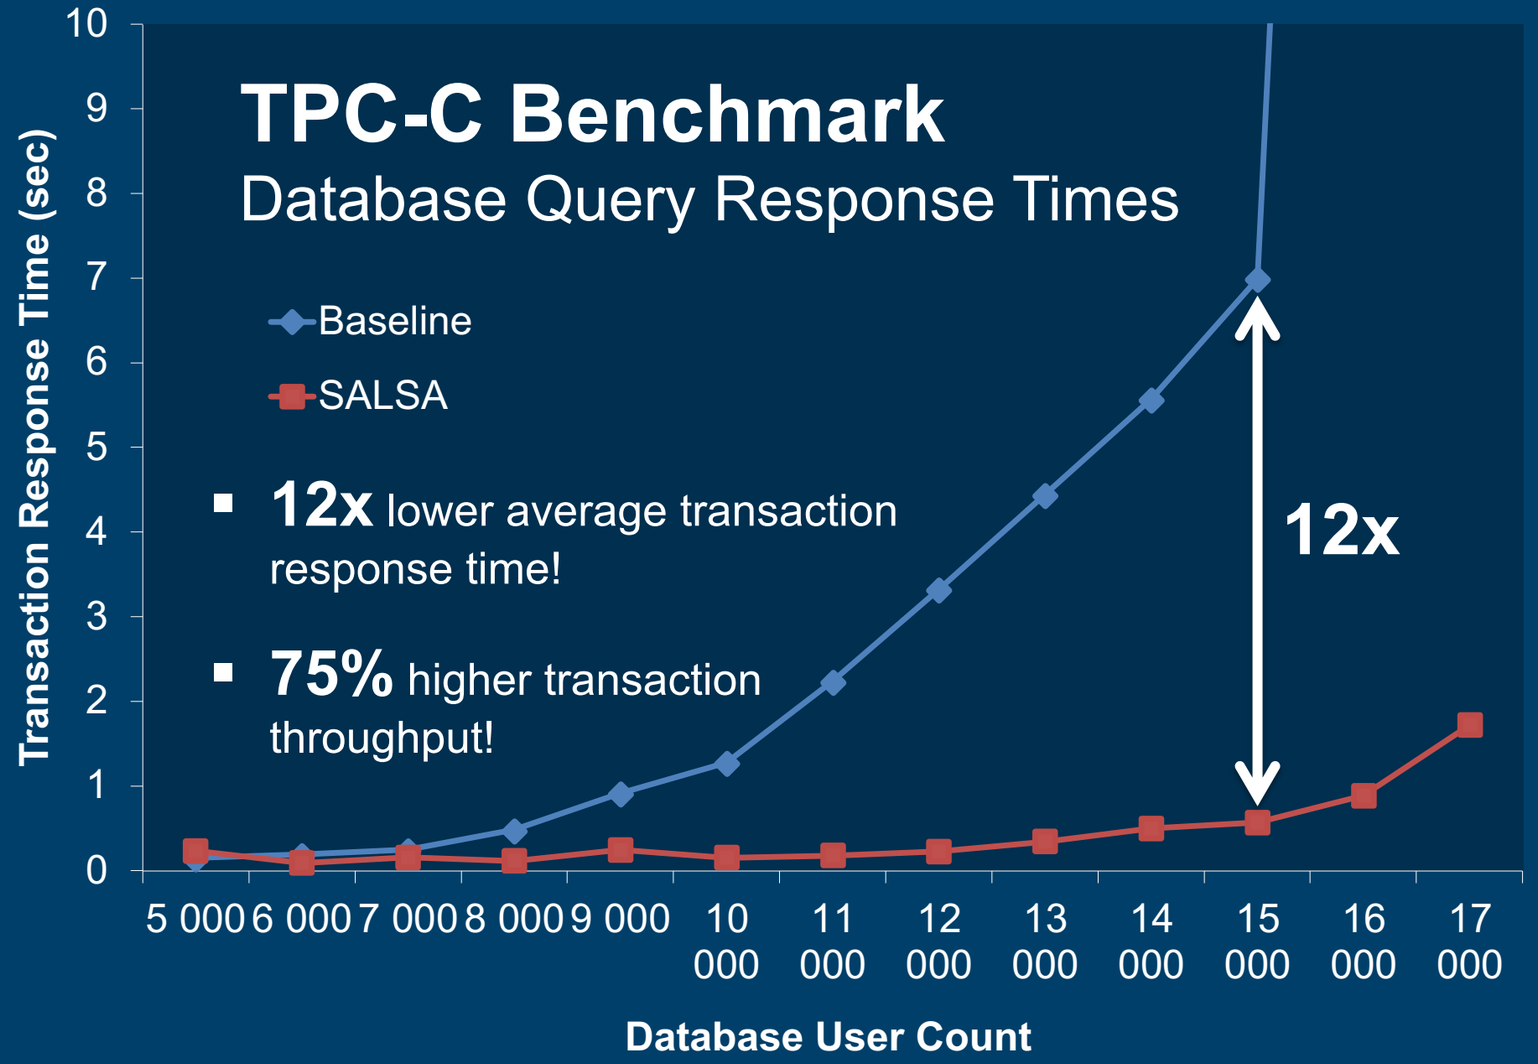

SALSA value-add:

1. **Extend SDS to Flash:** SALSA extends the benefits of SDS to take advantage of lowest-cost commodity Flash storage.
2. **Address read-dominated workloads:** SALSA enables the use of lowest-cost commodity SSDs for to specifically address the growing set of read-dominated workloads, by working around the physical limitations using software intelligence.
3. **Competitive advantage:** SALSA is a unique competitive advantage not available elsewhere, whether hardware Flash arrays, or SDS such as ScaleIO, vSAN or Ceph.
4. **Competitive against disk:** SALSA makes Flash cost competitive vs. a disk-based solution for read-dominated workloads.

The Data Deluge

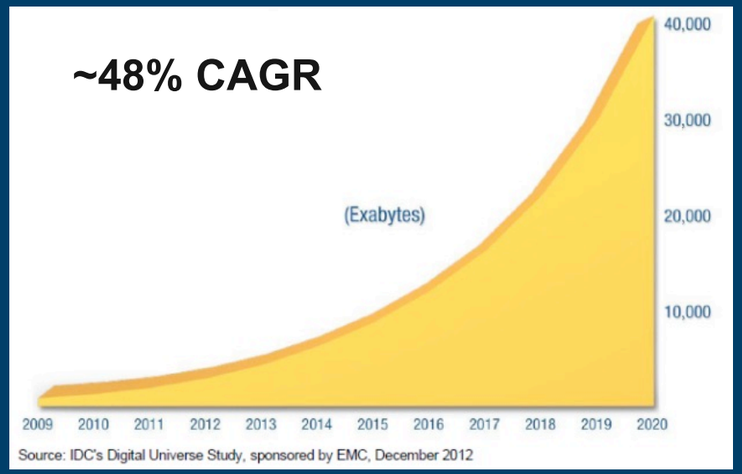

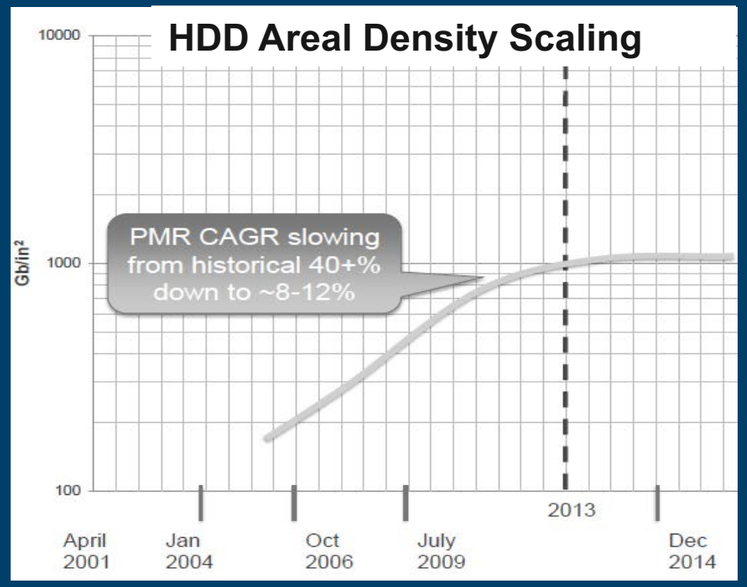

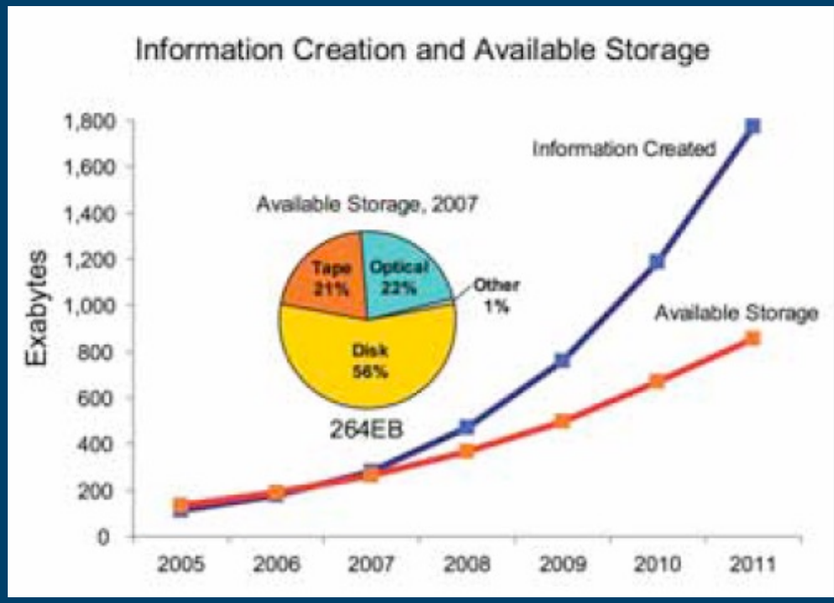

80% of all files created are inactive – no access in at least 3 months!

Source: D. Anderson, 2013 IEEE Conf. on Massive Data Storage

Advantages of Tape

- **Capacity Density**
  - Commercially available today: 10 TB per cartridge
- **Energy efficiency**
  - No power needed after data has been recorded
- **Security**
  - Drive-level encryption
  - Data inaccessible when drive is not loaded
- **Long media lifetime**
  - 30+ years
- **Reliability**
  - Orders of magnitude better error protection than disk
  - Typically, no data loss in case of drive failure

The main net advantage of tape is **low cost**

Areal Density Trends

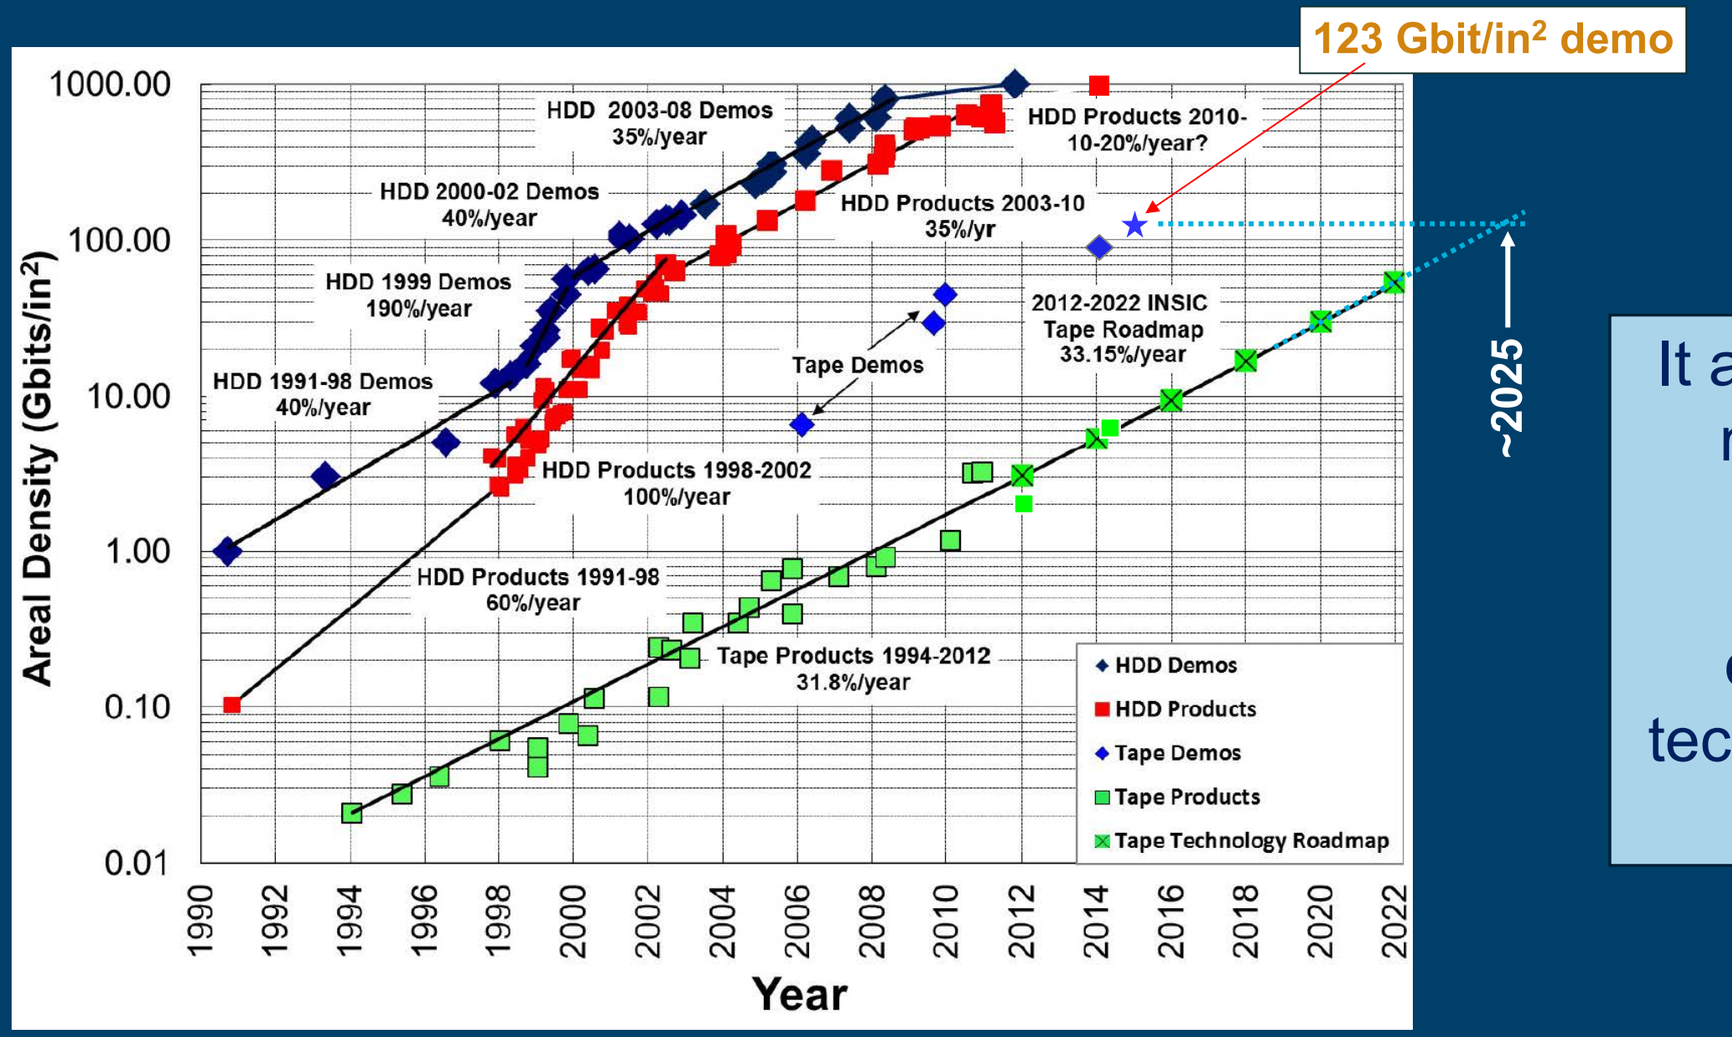

It appears that tape will maintain its almost 10x or more cost advantage over other storage technologies for at least another decade.

Why is Tape Interesting Again?

1

Tape areal density trends

- April 2015: IBM and Fujifilm demonstrated **123 Gbit/inch^2 on conventional BaFe media** (220 TB in a single cartridge!)
- Tape can grow capacity at least 40% per year, continuously increasing its $/GB advantage over disk
- The INSIC roadmap indicates that data rates of **2 GB/s** appear feasible by 2022

2

Linear Tape File System – LTFS

- An open format and interface to tape
- For the first time, tape can be used with open software

IBM demonstration shows that tape can sustain the roadmap for at least **another decade**

The Linear Tape File System (LTFS)

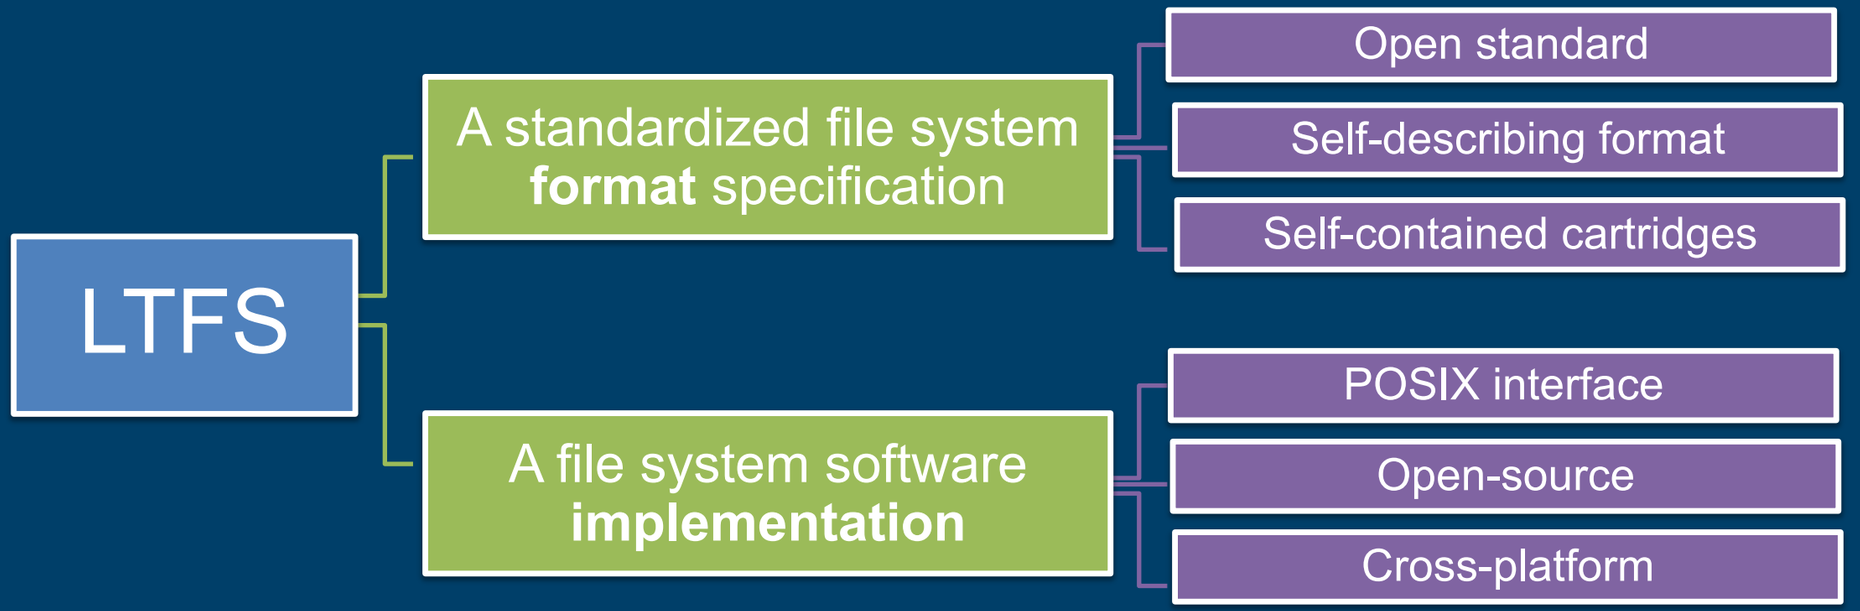

IBM LTFS

- **LTFS Single Drive Edition (LTFS SE)**
  - Support for standalone drives only (no robotics)
  - **Open-source**
- **LTFS Library Edition (LTFS LE)**
  - Support for tape libraries and library automation
- **LTFS Enterprise Edition (LTFS EE)**
  - Fully automated cartridge management
  - Disk-based cache
  - Seamless integration with GPFS

IceTier: Integrating Tape with OpenStack Swift

Goal

- Augment cloud object storage with a low-cost, cold storage tier for archive use cases
- Reduced availability (on the order of minutes)
- Reduced cost (significantly lower than disk)

Main Idea

- Integrate LTFS with a standard disk-based OpenStack Swift installation

Facts about tape

- Tape is at least 6x cheaper than disk
- Tape density scaling and cost are projected to be advantageous over disk for the next 10 years

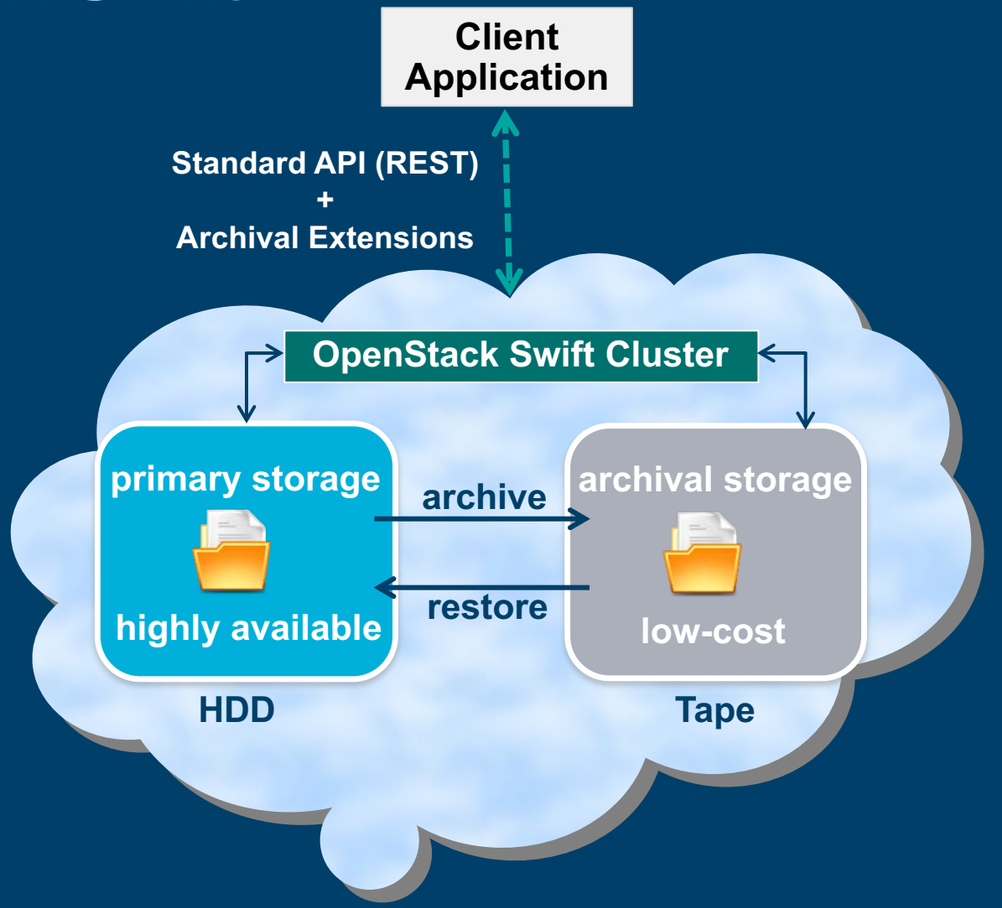

Offerings today

Amazon Glacier (3h) :	0.01 $/GB/month
Google Nearline (3s) :	0.01 $/GB/month

Integrating Tape with OpenStack Swift

Approach: Introduce LTFS-based Tape Data Ring

- Introduce Data Ring consisting of multiple LTFS instances (arranged in availability zones and geographic regions)
  - Reuse Swift's replication function for data reliability across tape libraries and geographies
  - Reuse Swift's scalability by adding more LTFS instances as needed

- Minor modifications in Swift required
  - Extensions to the standard REST API (e.g. special extended attributes to control tape tiering, timeouts)
  - Modified auditing and data collocation

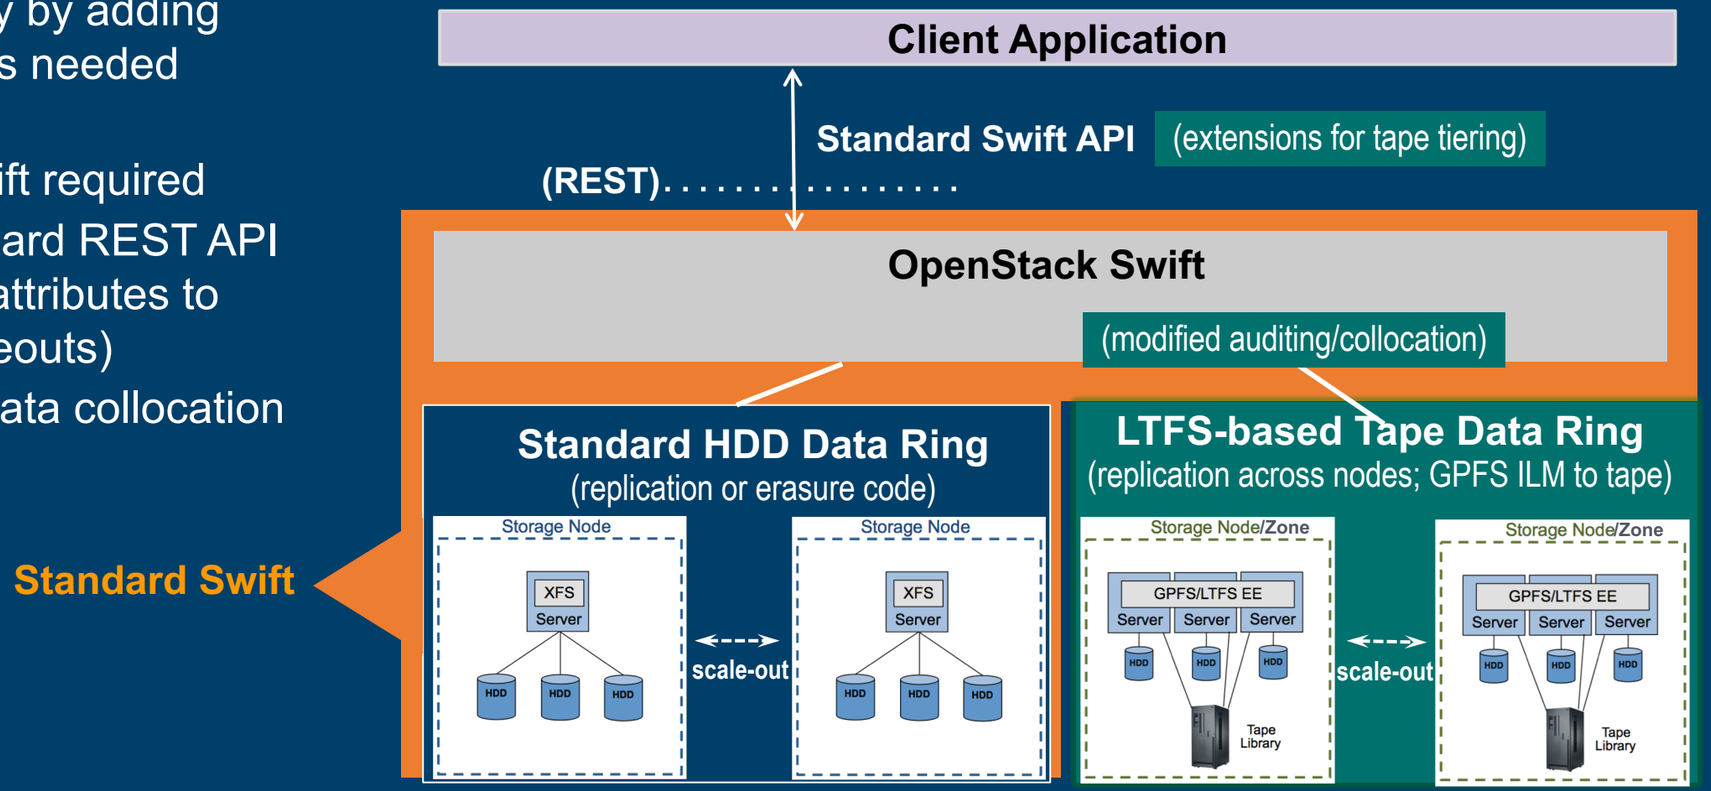

Swift Community Effort to Support Archival Storage

Tape Tier Object Retrieval Time – Simulation Results

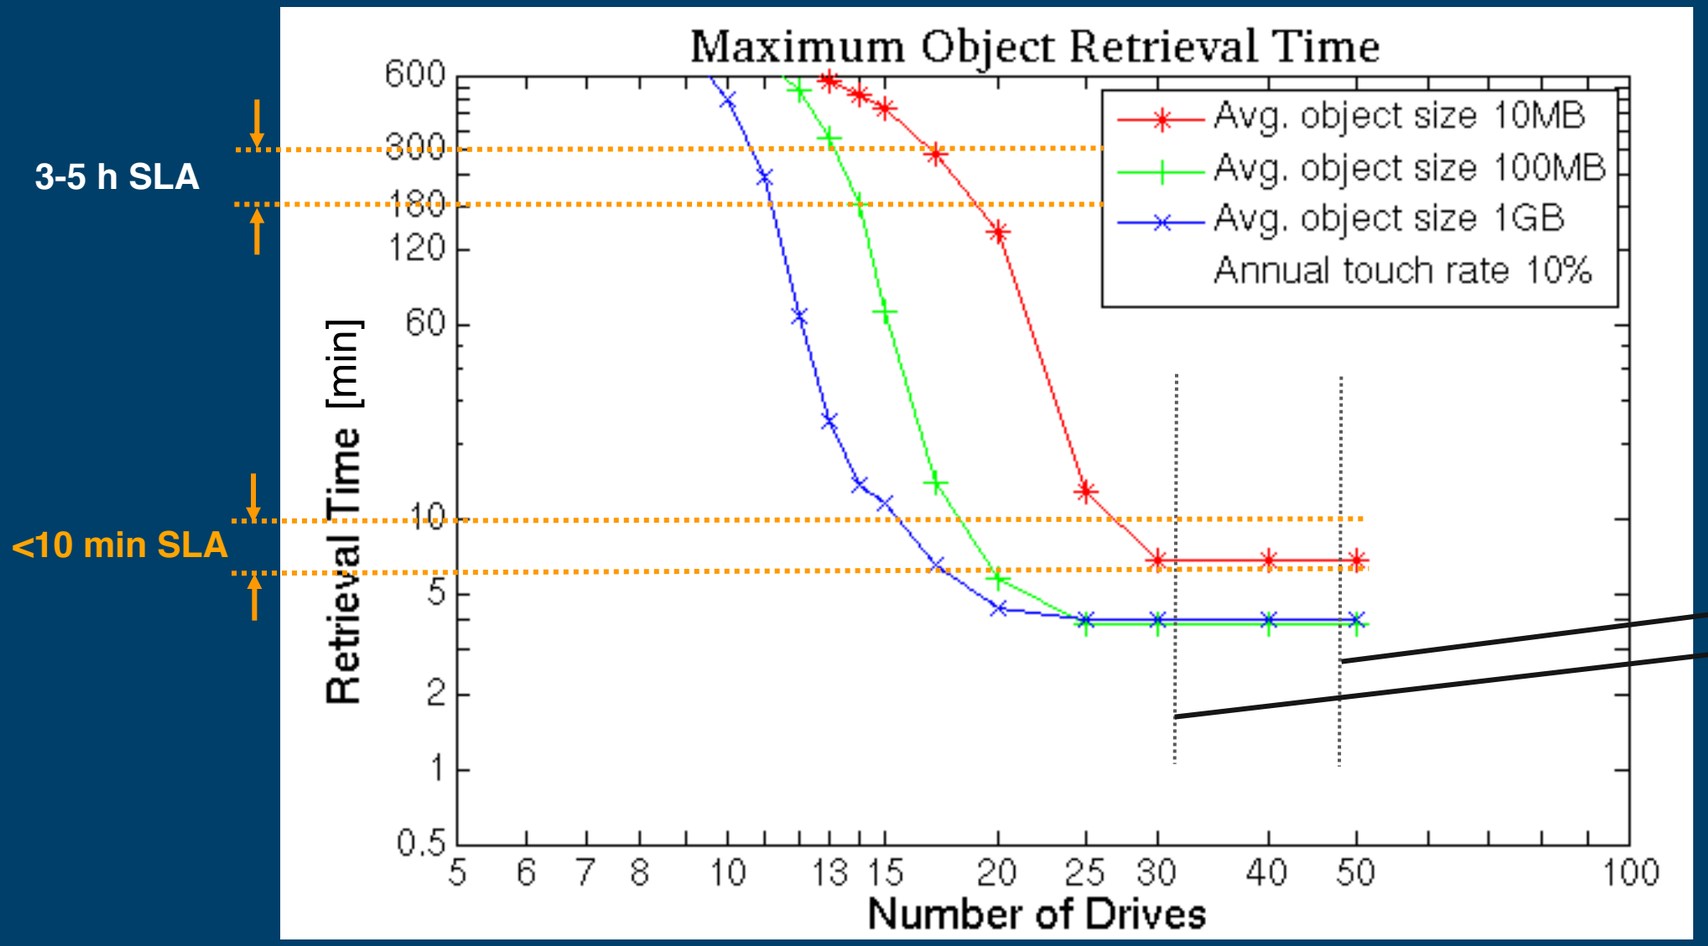

Retrieval time of less than 10 min achievable with 30 drives / 20 PB!

20 PB (4 racks):
12 drives/rack (typical)
7.5 drives/rack (required)

Workload assumptions:
- Poisson arrivals of individual or grouped requests, corresponding to 10% annual touch rate, exponentially distributed object size
- Some collocation of requests is assumed: on average, not less than 1GB of data is read per tape mount

From Cloud to InterCloud

Cloud-based object storage systems is a success story
- Simple APIs (Put/Get Key Value)
- Amazon S3 exceeded 2 trillion objects in April 2013
- Prices and scale which can't be met with traditional architectures

but with several **limitations**
- Poor passive security of stored data
- Simplicity goes hand-in-hand with lack of enterprise features
- Cloud downtime causes data unavailability
- Each cloud provider has their own API → large switching effort & risk

Storage in the InterCloud
- Use multiple clouds (public and/or private)
  - Limits trust in a single provider
  - Heterogeneity ensures genuine failure independence
- Add Client-side intelligence
  - Offers enterprise-grade features
- Provider-agnostic API
  - One API exposed to the client, different APIs in the background

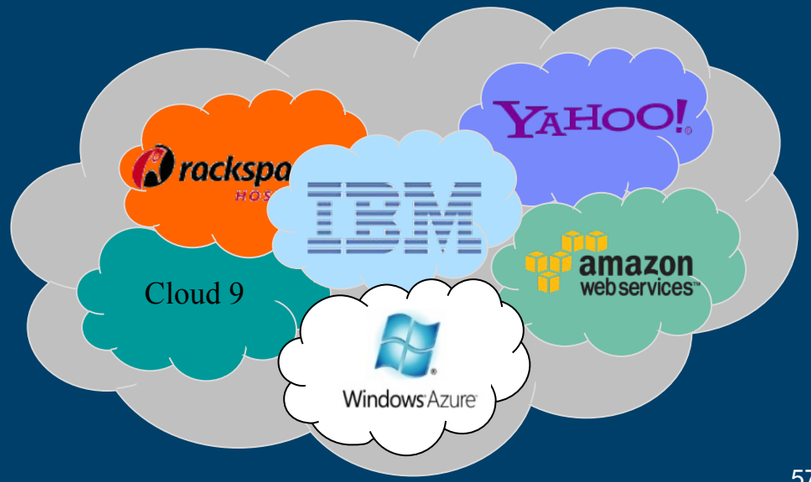

Multi-Cloud Storage Toolkit

What: A software-defined enterprise cloud storage **gateway**

Why: Address customer concerns regarding cloud **security**, **resilience**, and **vendor lock-in**

Goal: Enable existing storage products to natively support public/private cloud storage

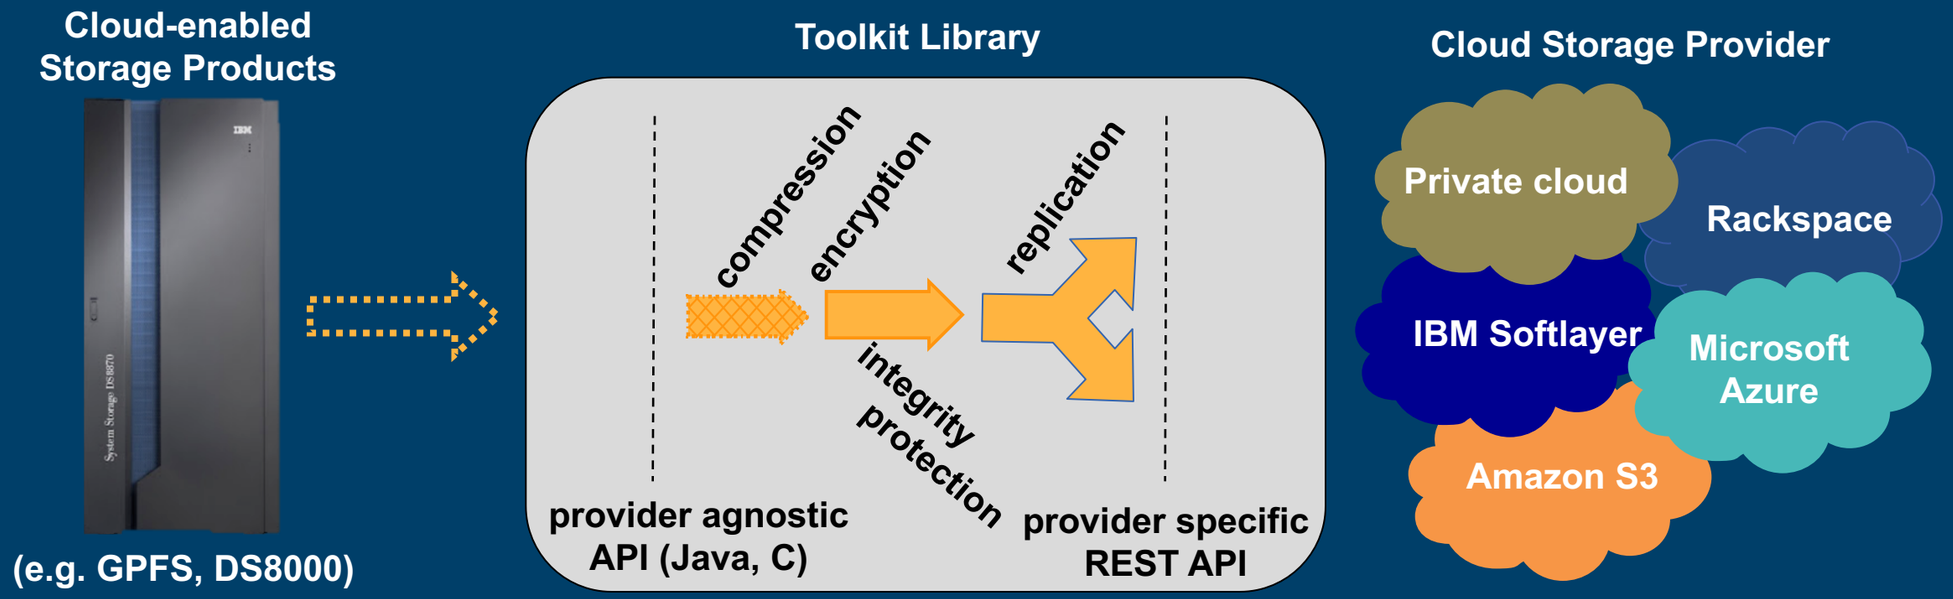

A Software-Defined Enterprise Cloud Storage Gateway for:

- transparent data migration
- backup and disaster recovery
- security and high availability

Use Case 1: Cloud ILM for File Storage (GPFS)

Goal: Enable a **secure**, **reliable**, **transparent** cloud storage tier in GPFS (a.k.a. Spectrum Scale)

Motivation: **Manage data growth** by placing file data

- in the right tier at the right time according to its value
- while being available under one common name space at any time
- leveraging the economy of scale of the cloud

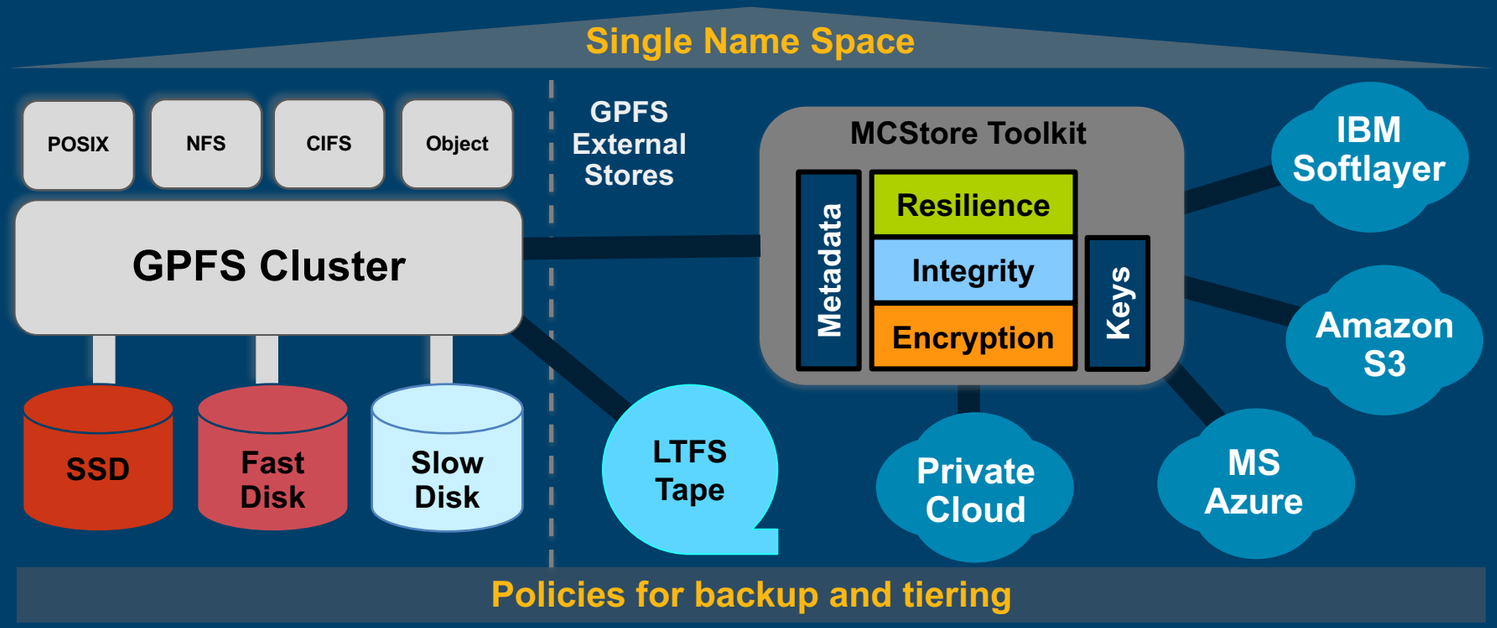

Value

- Seamless file migration between local disk and cloud
- File system backup for DR
- Efficient data sharing between remote clusters
- Multiple cloud storage tiers (using multiple cloud providers)
- Run workloads locally or in the cloud

Use Case 2: Cloud Backup for Block Storage

A cloud-based Time Machine for enterprise block storage

Built-in easy-to-use features to support various use cases:

- Backup
- Disaster recovery
- Data sharing
- Migration/archiving
- Compliance/auditing

- The toolkit applies encryption and integrity protection
- The toolkit stores full and/or incremental snapshots of volumes in the cloud.
- Snapshots can be restored from the cloud to the original or to a new volume.

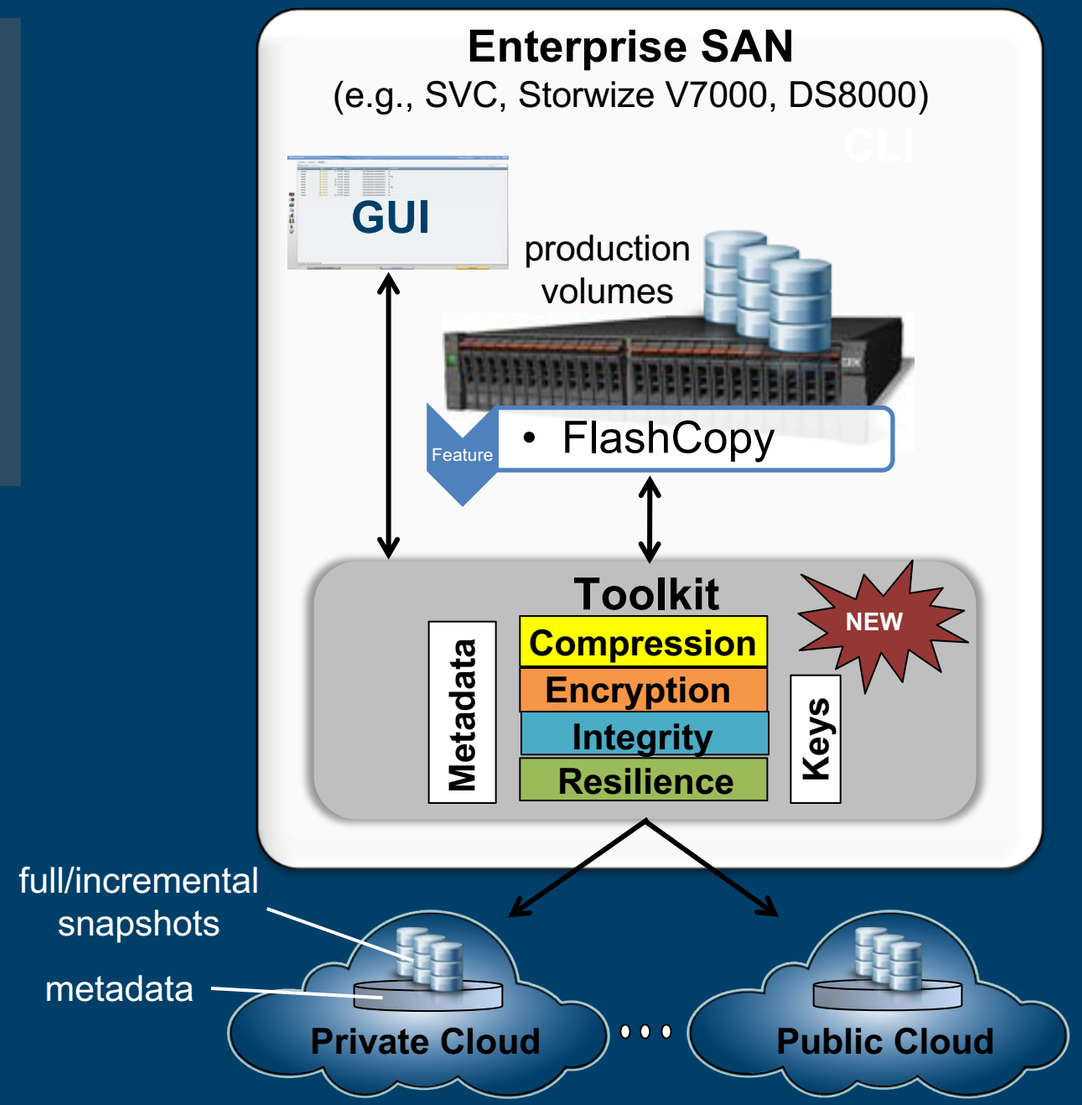

HyperScale Convergence: Platform Computing and Spectrum Storage

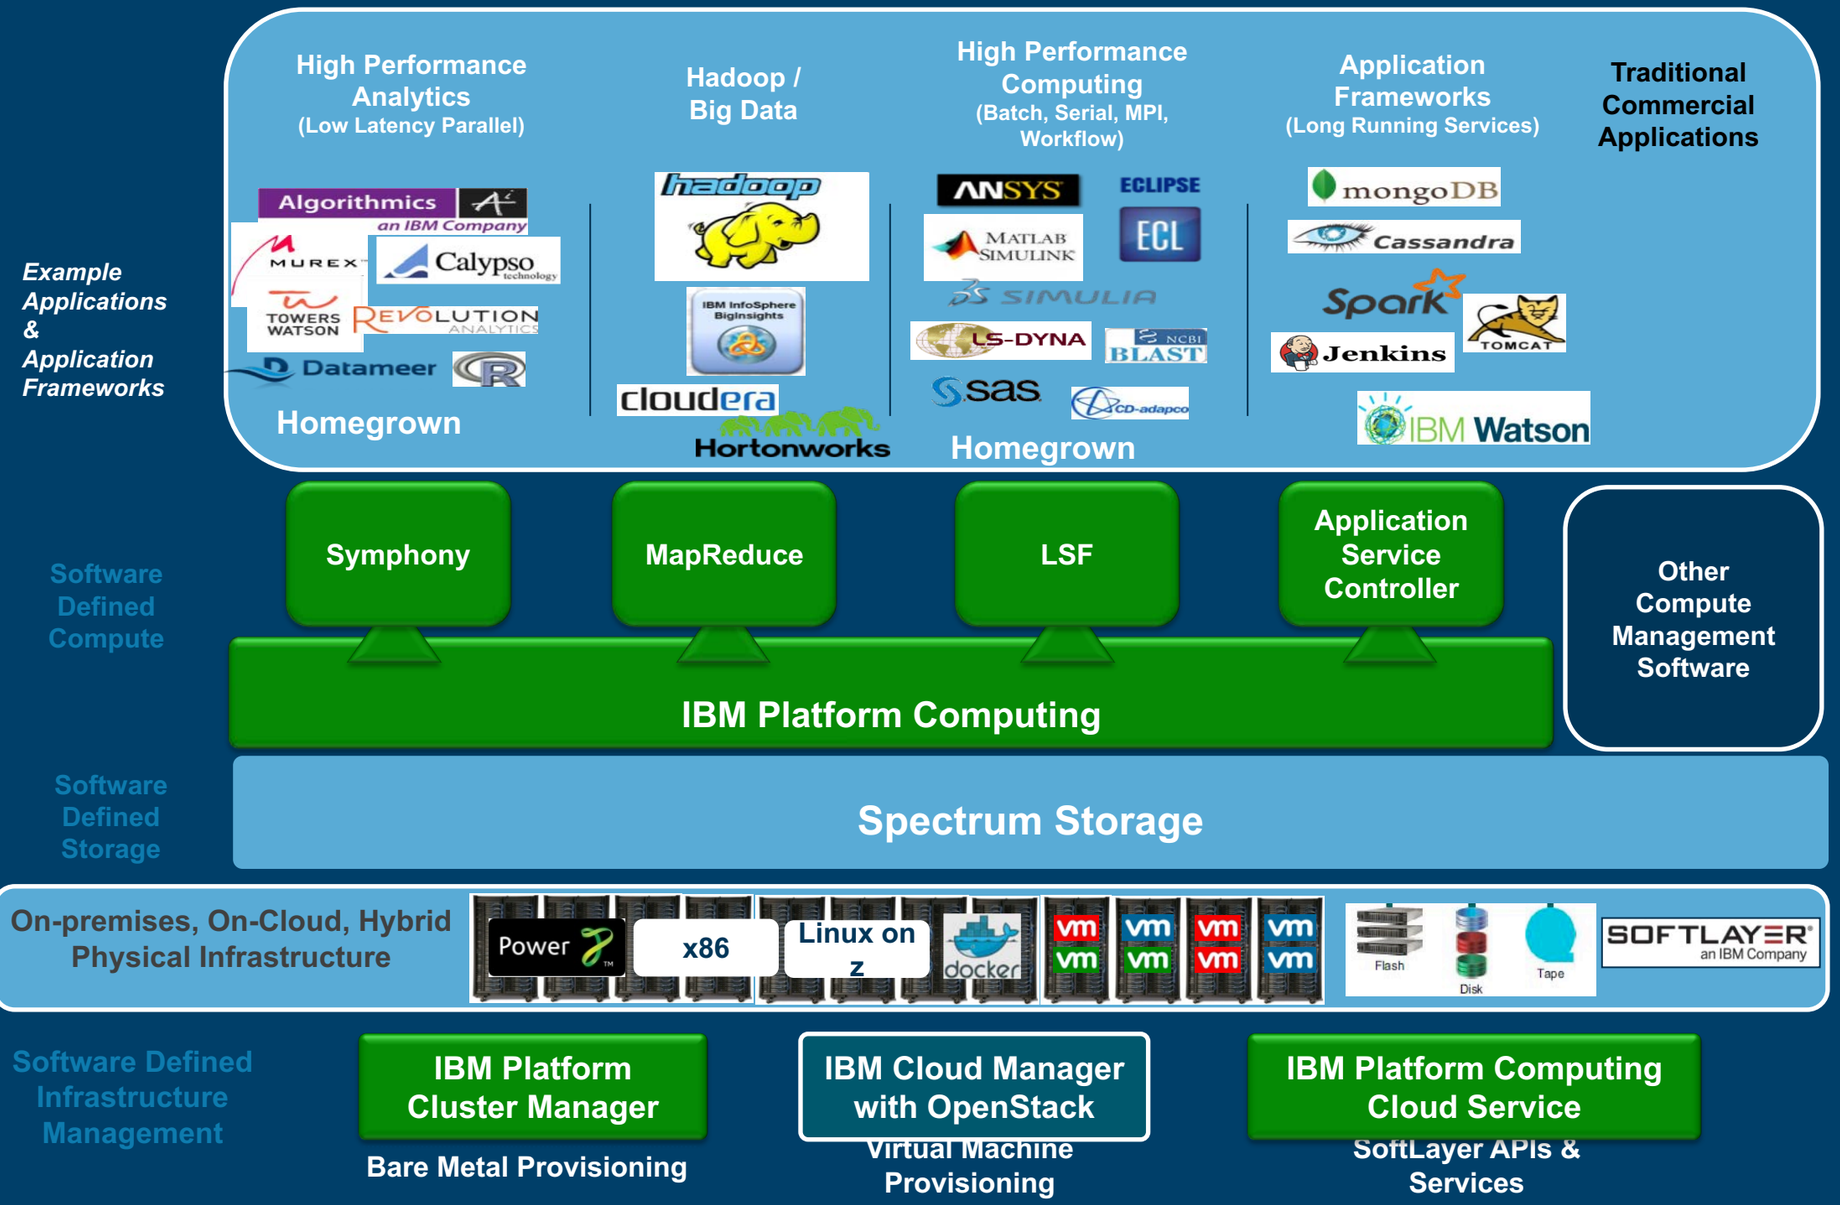

Hyper-scale Converged

Single Interface
installation, configuration, monitoring

Automation
automated provisioning,
SLA management, volume / file creation

Advanced Monitoring & Reporting
single-pane monitoring, remediation

REST APIs
application integration

Store
optimized data placement with global access

Analyze
workload and data-aware converged infrastructure

Protect
data lifecycle management and protection

A complete offering for modern data-centric applications

Eliminate silos

- Multi-tenant converged application and data fabric

Simplify administration

- Single pane management and monitoring

Improve data availability

- Global shared access to ensure data is available right when its needed

Minimize deployment time

- Automated deployment of physical, virtual and containerized resources

Future proof your data center

- Modular building blocks for multi-dimensional scalability

Conclusions

- Unleash the value of commodity Flash with **SALSA**
  - Transform user access patterns to be as Flash-friendly as possible
  - Elevate the performance and endurance of low-cost SSDs to meet datacenter requirements
- **Tape** will keep increasing its cost advantage over disk (for at least 10 years)
  - Use **LTFS** to build open systems based on tape
  - Integrate **Tape into Swift**, for open object archival based on tape
- **Cloud-enable** traditional storage transparently, securely, reliably
  - Leverage economies of scale for both File and Block storage
  - Enable new use cases using multiple cloud providers
- **Hyperscale Convergence** to integrate storage seamlessly in cloud-native data-intensive applications

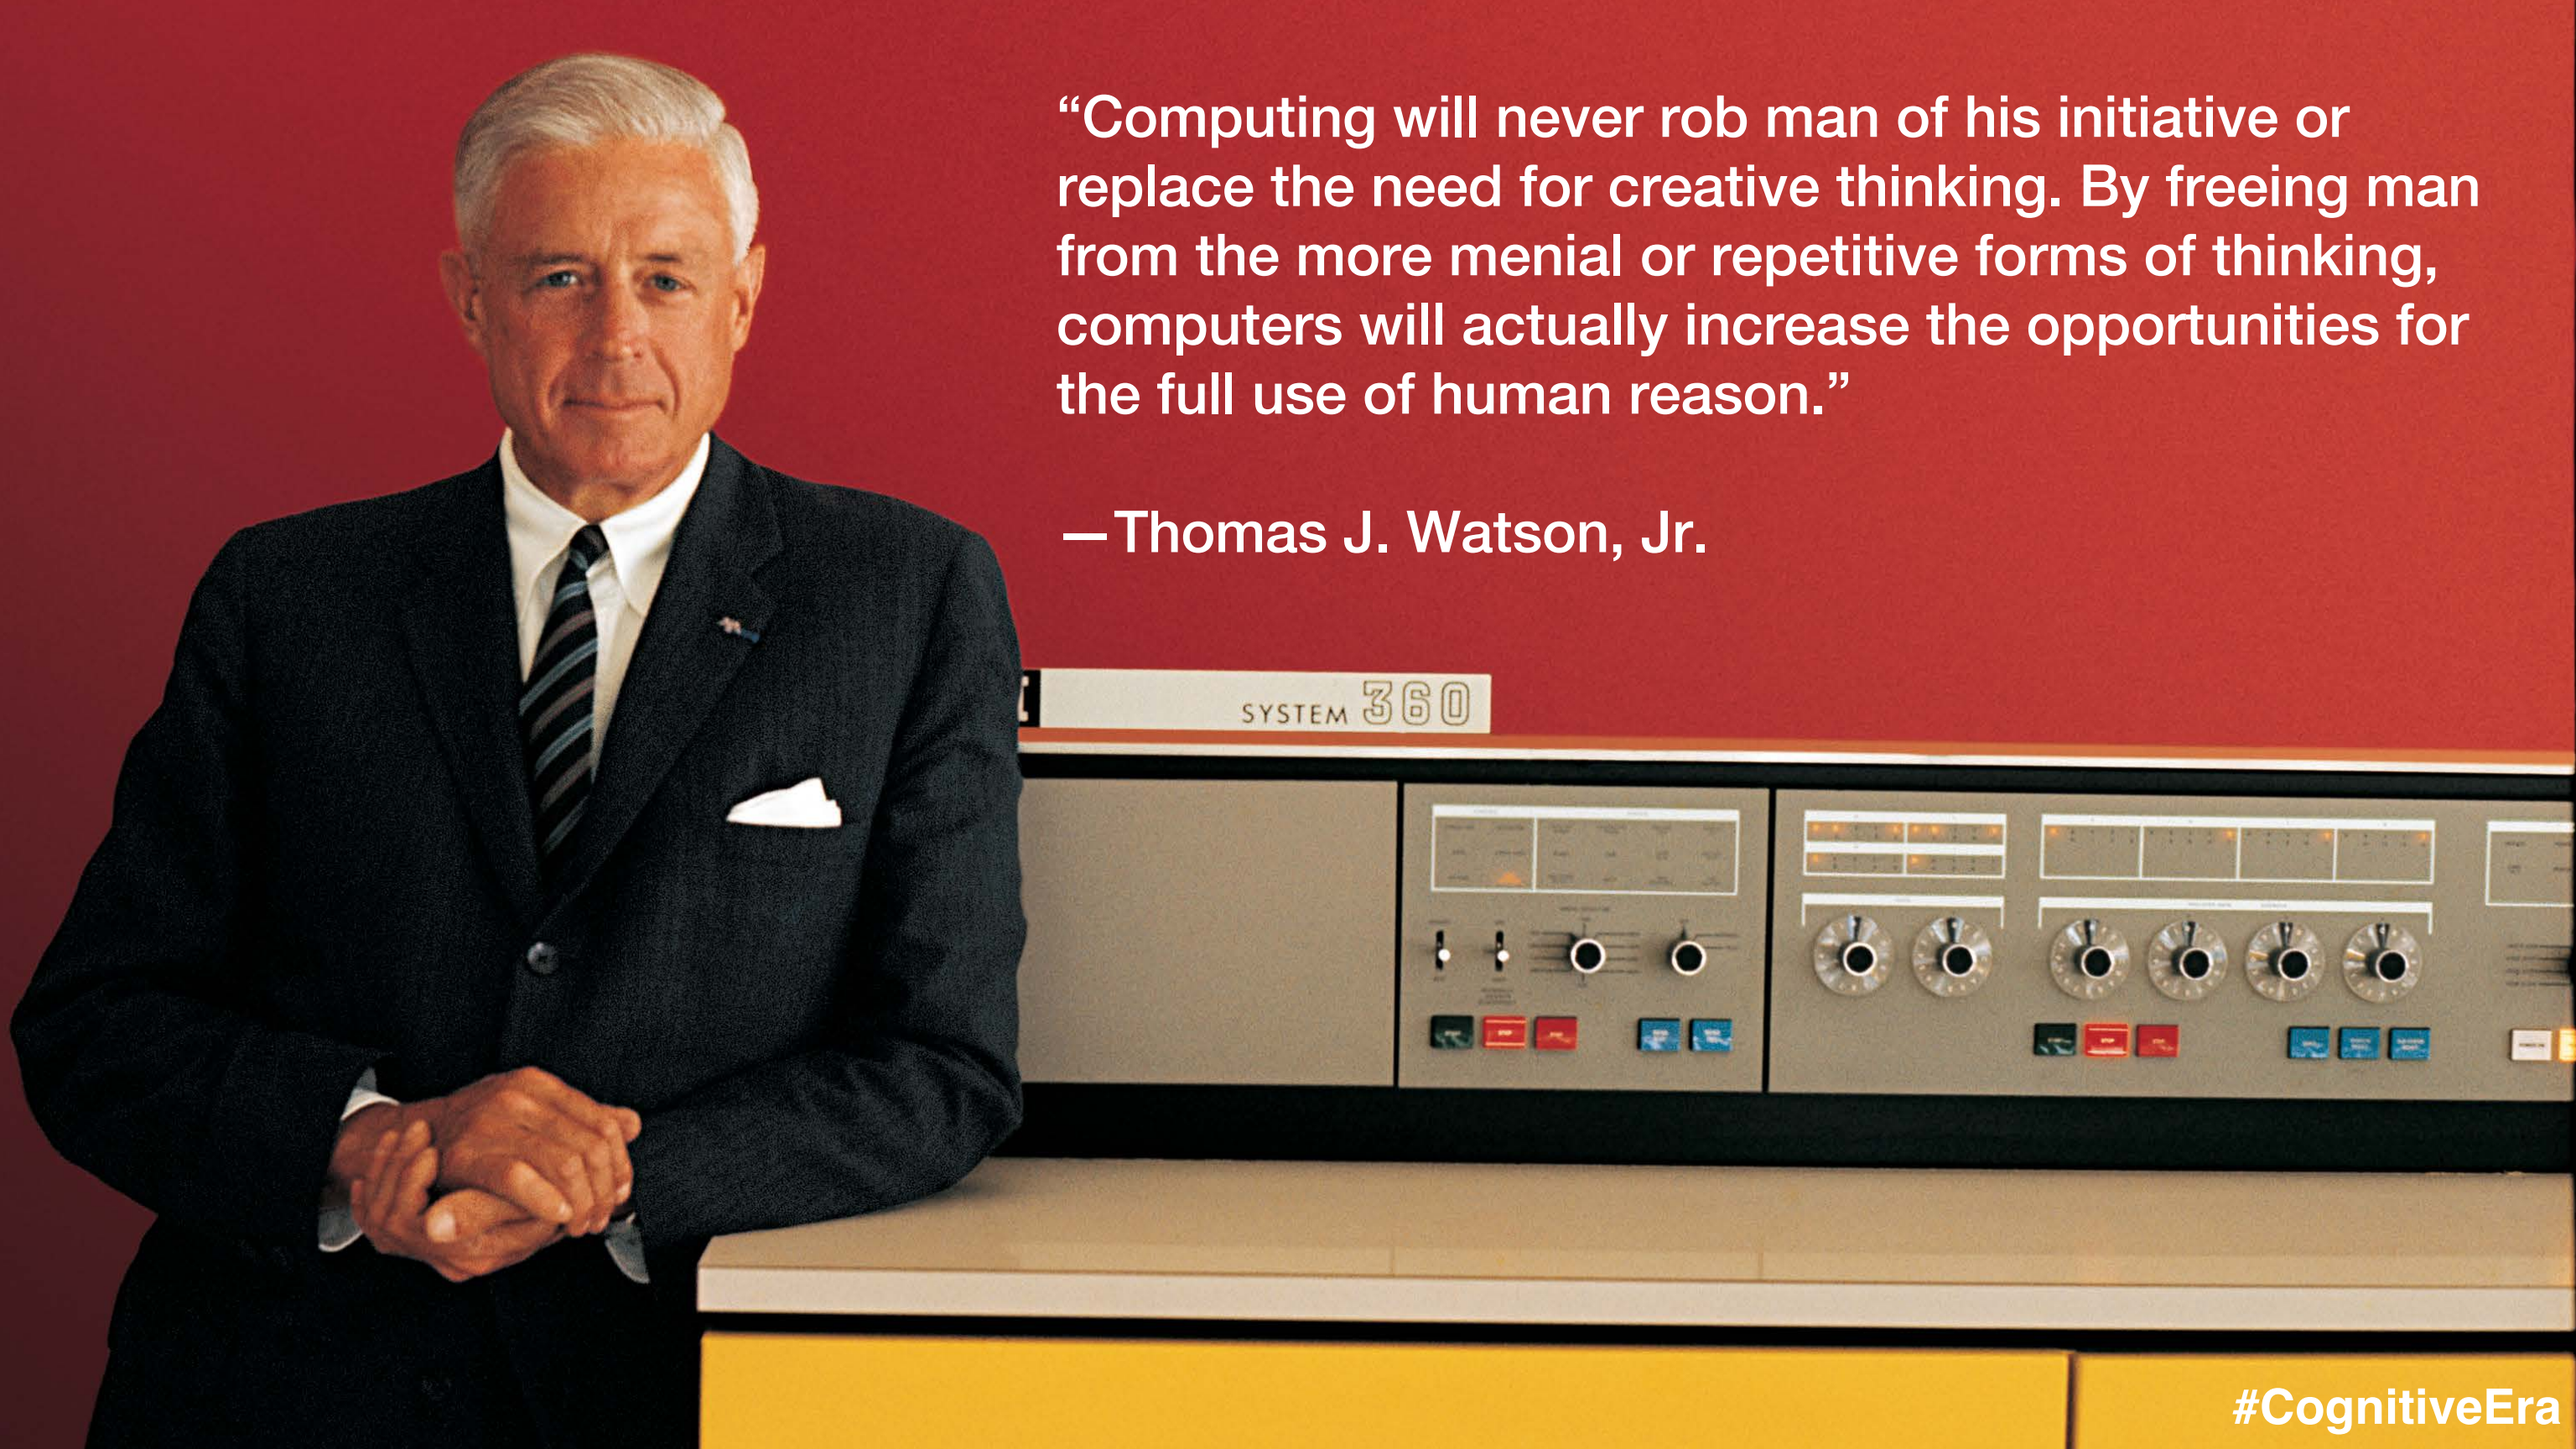
“Computing will never rob man of his initiative or replace the need for creative thinking. By freeing man from the more menial or repetitive forms of thinking, computers will actually increase the opportunities for the full use of human reason.”
—Thomas J. Watson, Jr.
SYSTEM 360
#CognitiveEra