STORAGE RESEARCH @ HP Labs

Jay Wylie, Senior Research Scientist
Storage and Information Management Platforms Lab (SIMPL),
HP Labs

Acknowledgements

- Stan Williams & *Information and Quantum Systems Labs*
- Partha Ranganathan, Niraj Tolia & *Exascale Computing Lab*
- *Storage and Information Platforms Lab*

Caveat:
HP Labs research projects *not* HP products!

Outline

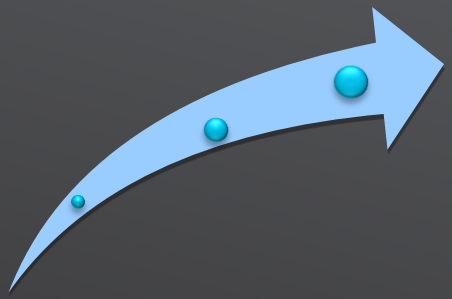

Information-aware

Global-scale storage

NVMem & data centers

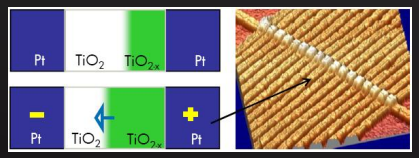

Memristors

3 fundamental passive linear circuit elements

Resistor – 1827
Georg Ohm

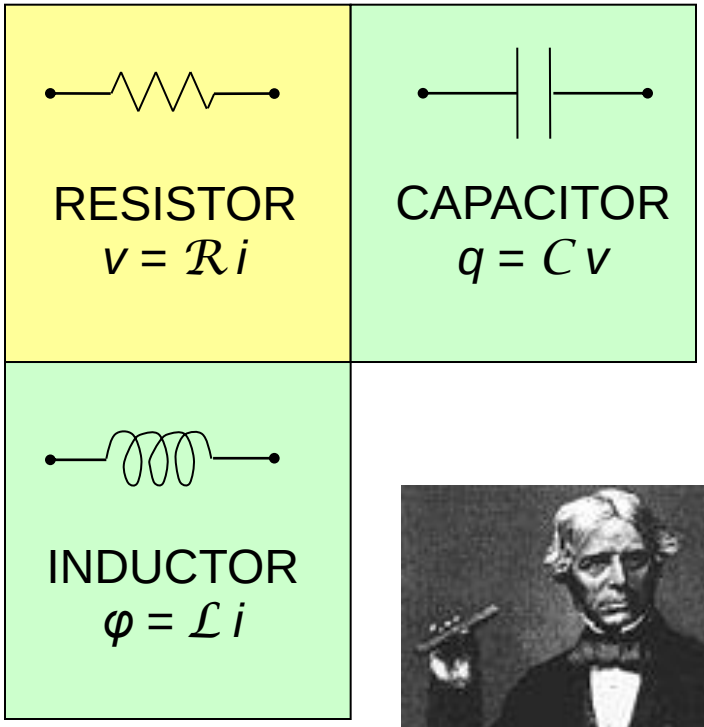

Capacitor - 1745
Volta / von Kleist & van Musschenbroek
Benjamin Franklin

Inductor – 1831
Michael Faraday
Joseph Henry

1960's – Leon Chua generalizes circuit theory to nonlinear systems

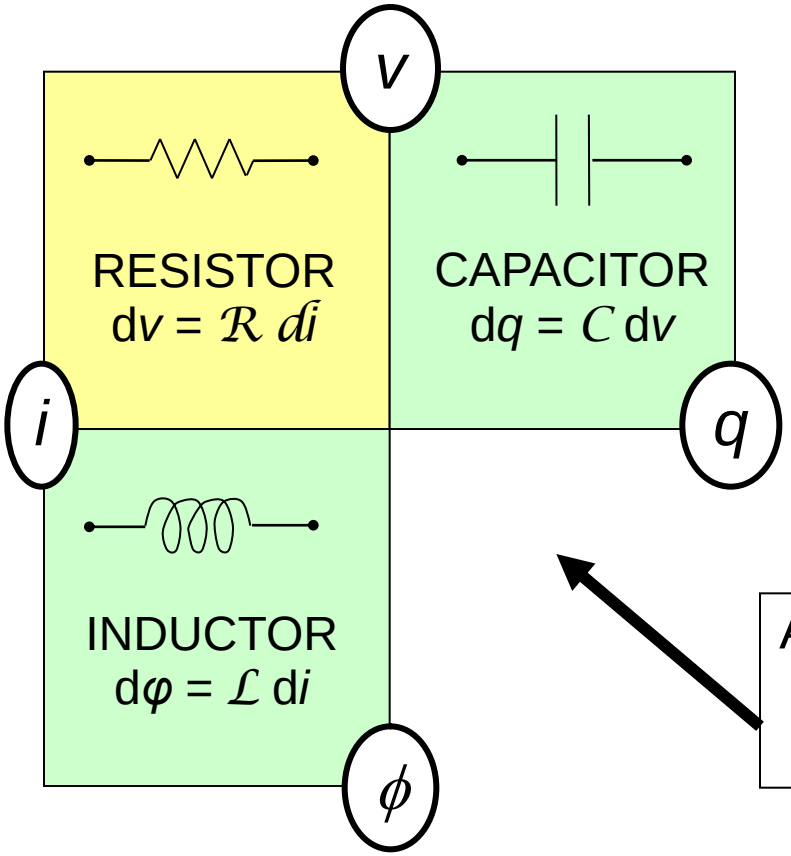

And sees that there is a hole where an obvious relation seems to be missing

Chua defined the memristor, but stated that there was no known example

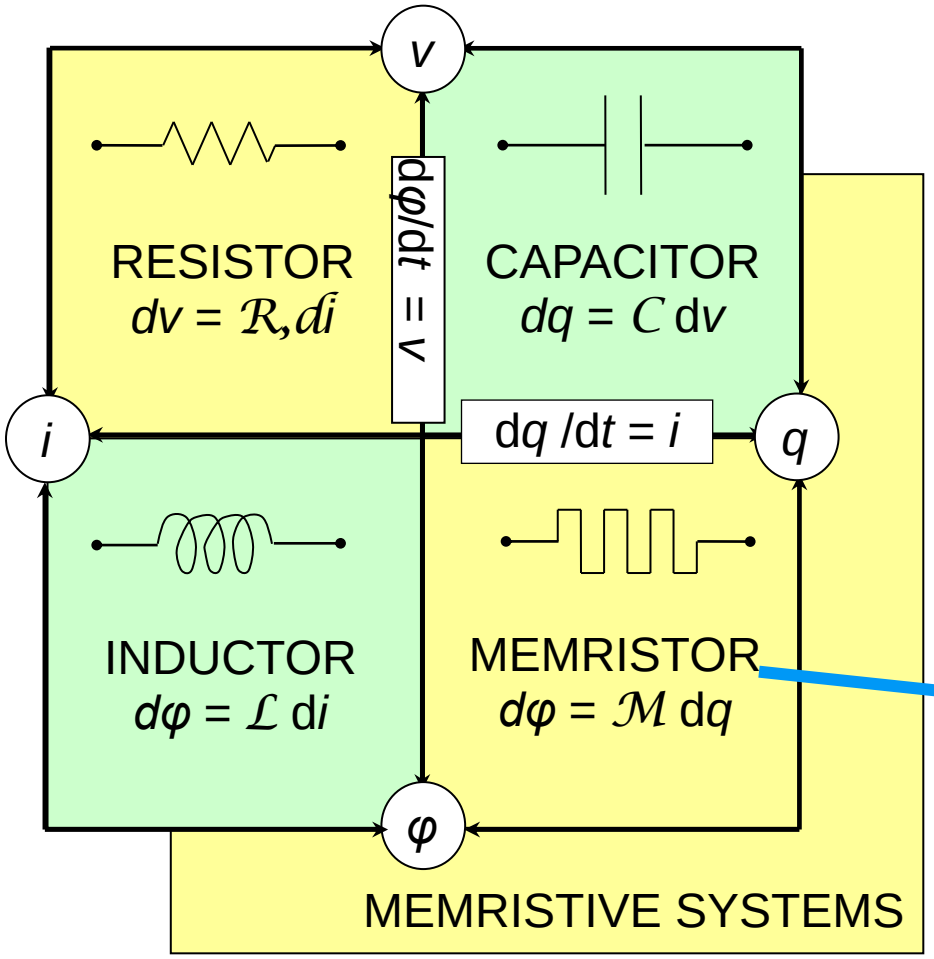

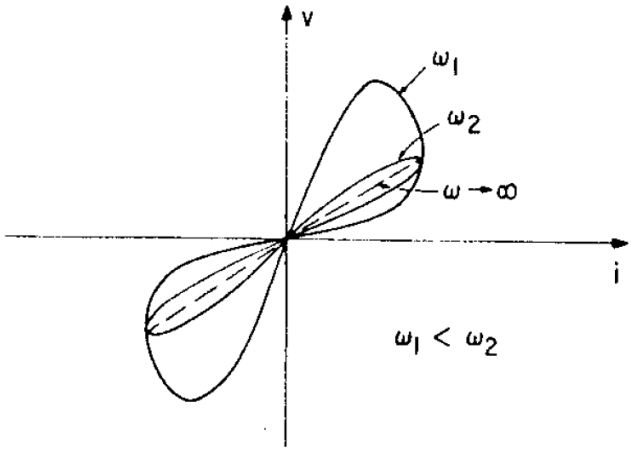

A Memristor can act like a resistor whose value varies with voltage. And, it remembers the value after current stops.

L. O. Chua, “Memristor - the missing circuit element,” IEEE Trans. Circuit Theory 18, 507–519 (1971).
L. O. Chua and S. M. Kang, "Memristive devices and systems," *Proc. IEEE,* 64 (2), 209-23 (1976).

HP Labs: Memristor from theory to practice

The missing Memristor found,
D. Strukov *et al., Nature* 453, 80 (2008)
http://www.nature.com/nature/journal/v453/n7191/full/nature06932.html

The Mysterious Memristor,
IEEE Spectrum, May 2008
http://spectrum.ieee.org/semiconductors/design/the-mysterious-memristor

Memristive switches enable stateful logic Operations via material implication,
J. L. Borghetti et al., *Nature* 464, 873 (2010)
http://www.nature.com/nature/journal/v464/n7290/full/nature08940.html

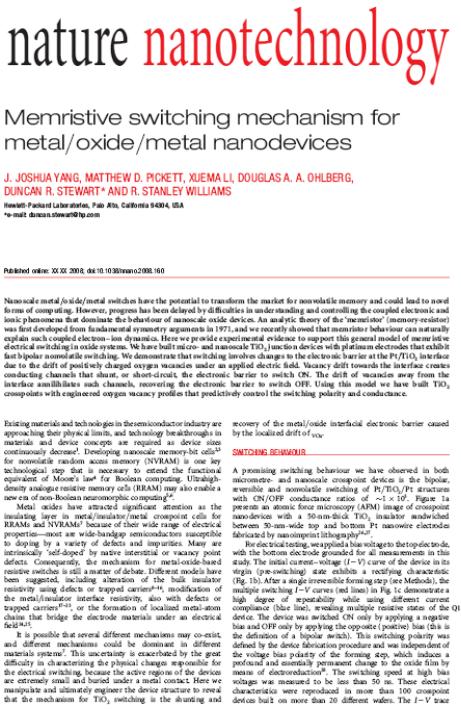

nature nanotechnology

Memristive switching mechanism for metal/oxide/metal nanodevices

J. JOSHUA YANG, MATTHEW D. PICKETT, XUEMA LI, DOUGLAS A. A. OHLBERG, DUNCAN R. STEWART* AND R. STANLEY WILLIAMS
Hewlett-Packard Laboratories, Palo Alto, California 94304, USA
*e-mail: duncan.stewart@hp.com

Published online: XX XX 2008; doi:10.1038/nnano.2008.160

Nano scale metal/oxide/metal switches have the potential to transform the market for nonvolatile memory and could lead to novel forms of computing. However, progress has been delayed by difficulties in understanding and controlling the coupled electronic and ionic phenomena that dominate the behaviour of nanoscale oxide devices. An analytic theory of the 'memristor' (memory-resistor) was first developed from fundamental symmetry arguments in 1971, and we recently showed that memristor behaviour can naturally explain such coupled electron–ion dynamics. Here we provide experimental evidence to support this general model of memristive electrical switching in oxide systems. We have built micro- and nanoscale TiO_2 junction devices with platinum electrodes that exhibit fast bipolar nonvolatile switching. We demonstrate that switching involves changes to the electronic barrier at the Pt/TiO_2 interface due to the drift of positively charged oxygen vacancies under an applied electric field. Vacancy drift towards the interface creates conducting channels that shunt, or short-circuit, the electronic barrier to switch ON. The drift of vacancies away from the interface annihilates such channels, recovering the electronic barrier to switch OFF. Using this model we have built TiO_2 crosspoints with engineered oxygen vacancy profiles that predictively control the switching polarity and conductance.

Existing materials and technologies in the semiconductor industry are approaching their physical limits, and technology breakthroughs in materials and device concepts are required as device sizes continuously decrease[1]. Developing nanoscale memory-bit cells[2,3] for nonvolatile random access memory (NVRAM) is one key technological step that is necessary to extend the functional equivalent of Moore's law[4] for Boolean computing. Ultrahigh-density analogue resistive memory cells (RRAM) may also enable a new era of non-Boolean neuromorphic computing[5,6].

Metal oxides have attracted significant attention as the insulating layer in metal/insulator/metal crosspoint cells for RRAMs and NVRAMs[7] because of their wide range of electrical properties—most are wide-bandgap semiconductors susceptible to doping by a variety of defects and impurities. Many are intrinsically 'self-doped' by native interstitial or vacancy point defects. Consequently, the mechanism for metal-oxide-based resistive switches is still a matter of debate. Different models have been suggested, including alteration of the bulk insulator resistivity using defects or trapped carriers[8–14], modification of the metal/insulator interface resistivity, also with defects or trapped carriers[15–17], or the formation of localized metal-atom chains that bridge the electrode materials under an electrical field[18,19].

It is possible that several different mechanisms may co-exist, and different mechanisms could be dominant in different materials systems[7]. This uncertainty is exacerbated by the great difficulty in characterizing the physical changes responsible for the electrical switching, because the active regions of the devices are extremely small and buried under a metal contact. Here we manipulate and ultimately engineer the device structure to reveal that the mechanism for TiO_2 switching is the shunting and recovery of the metal/oxide interfacial electronic barrier caused by the localized drift of $v_{O}^{\cdot\cdot}$.

SWITCHING BEHAVIOUR

A promising switching behaviour we have observed in both micrometre- and nanoscale crosspoint devices is the bipolar, reversible and nonvolatile switching of Pt/TiO_2/Pt structures with ON/OFF conductance ratios of $\sim 1 \times 10^3$. Figure 1a presents an atomic force microscopy (AFM) image of crosspoint nano devices with a 50-nm-thick TiO_2 insulator sandwiched between 50-nm-wide top and bottom Pt nanowire electrodes fabricated by nanoimprint lithography[20,21].

For electrical testing, we applied a bias voltage to the top electrode, with the bottom electrode grounded for all measurements in this study. The initial current–voltage (I–V) curve of the device in its virgin (pre-switching) state exhibits a rectifying characteristic (Fig. 1b). After a single irreversible forming step (see Methods), the multiple switching I–V curves (red lines) in Fig. 1c demonstrate a high degree of repeatability while using different current compliance (blue line), revealing multiple resistive states of the QI device. The device was switched ON only by applying a negative bias and OFF only by applying the opposite (positive) bias (this is the definition of a bipolar switch). This switching polarity was defined by the device fabrication procedure and was independent of the voltage bias polarity of the forming step, which induces a profound and essentially permanent change to the oxide film by means of electroreduction[18]. The switching speed at high bias voltages was measured to be less than 50 ns. These electrical characteristics were reproduced in more than 100 crosspoint devices built on more than 20 different wafers. The I–V trace

Memristors

– Advantages
- Fast, cheap, dense, & low energy
- Can be fabbed with CMOS logic & 3D layers on single die

– Challenges
- Understand wear out
- Moving from the lab to the fab

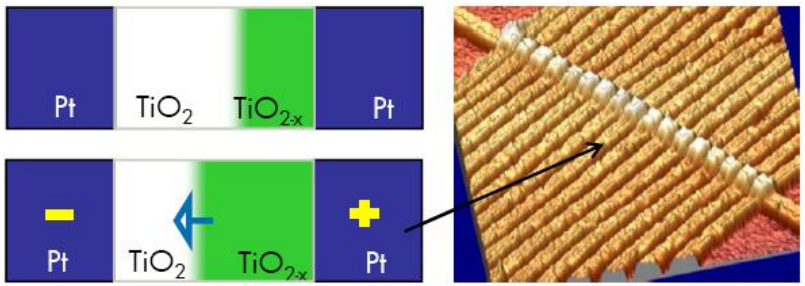

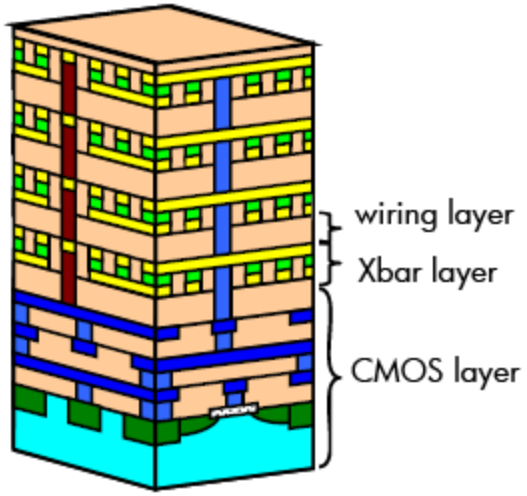

Figures from Stan Williams

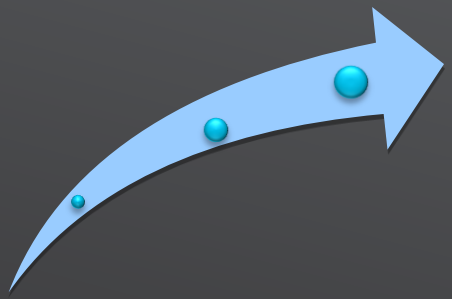

Information-aware

Global-scale storage

NVMem & data centers

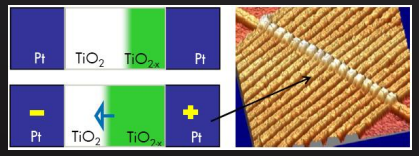

Memristors

Emerging NVMem Technologies

Technology	Density (μm^2/bit)	Bandwidth (GB/s)	Read latency (ns)	Write latency (ns)	Read energy per bit (pJ)	Write energy per bit (pJ)
Hard disk	n/a	0.5	3,000,000	3,000,000	2500	2500
Flash SSD	0.0021	1.0	25,000	200,000	250	250
DRAM	*0.0038*	*51.2*	55	55	*24*	*24*
PCRAM (22nm)	0.0058	variable	48	150	2	19
Memristor (22nm)	0.0048	variable	100	100	1-3	1-3

Table from Partha Ranganathan & Jichuan Chang

Compare-and-Swap NVMem

– Great fit for Sinfonia (Aguilera et al. SOSP 2007)
  - Simplified scalable, distributed data structures
– Leverage CAS NVMem to build mini-transactions

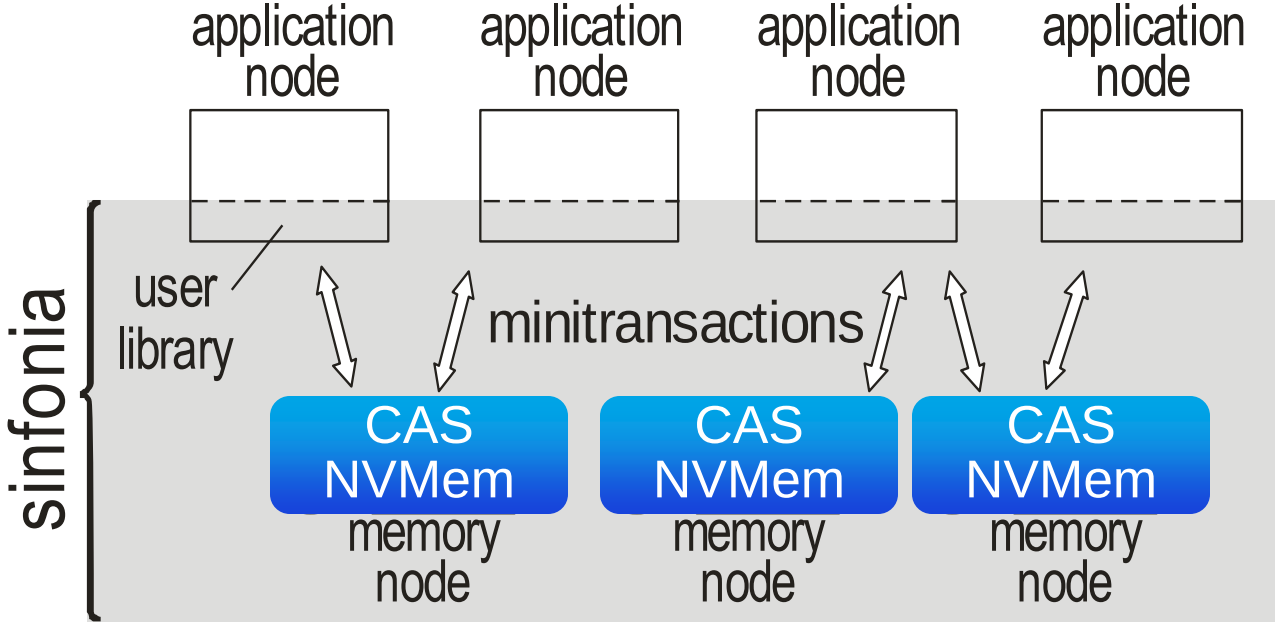

Figure from Mehul Shah

Compute-centric vs. data-centric architecture

Organize architecture around storage. Colocate compute with storage. Flatten memory hierarchy; introduce compute hierarchy. Convert memory bottleneck to network. Codesign software.

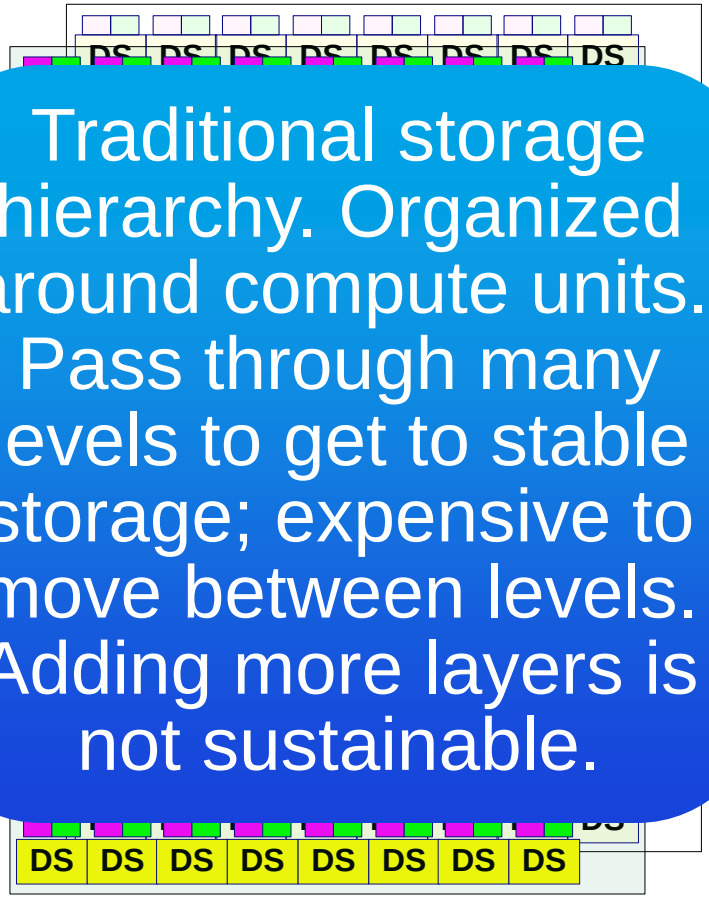

Traditional storage hierarchy. Organized around compute units. Pass through many levels to get to stable storage; expensive to move between levels. Adding more layers is not sustainable.

Figures from Partha Ranganathan & Niraj Tolia

NVMem and data centers

- Novel API's
  - Transactional interface (e.g., Compare-and-Swap NVMem)
  - Data-structure oriented NVMem (e.g., B-tree NVMem)
- Novel architectures
  - Data-centric architecture
  - NVMem + photonics
- Other challenges
  - Tolerating and masking (transient) faults
  - Navigate density-latency tradeoff
  - Balance processing, storage, and networking bandwidths

Information-aware

Global-scale storage

NVMem & data centers

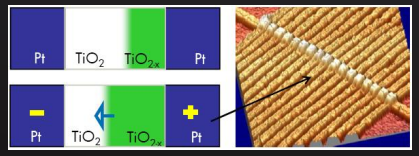

Memristors

Global-scale storage

- Global infrastructure
- Global clients/applications
- Research challenges
  - Scalability
  - Availability
  - Low cost
  - Flexibility

HP-Key Value Store (KVS)

- Research prototype
- Simple interface:
  - put(key, value)
  - value = get(key)
- Goals:
  - 1+ Exabyte, 1T objects
  - 10+ data centers

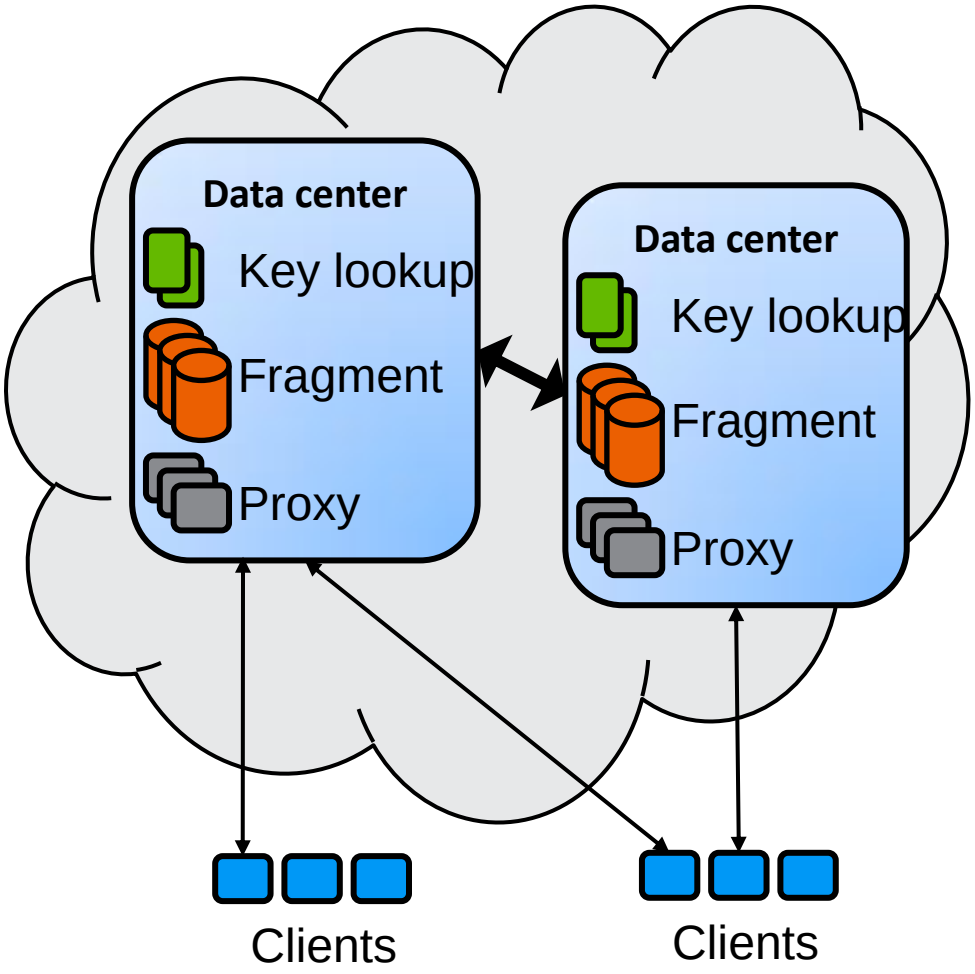

Storage interface & consistency

- Traditional interfaces based on strong consistency
  - Block-based accesses
  - File systems
  - DBMSes
- Simple put/get interface
  - Values `Put` into the system may be *eventually* consistent
  - `Get` operations may observe different values
  - Photo sharing, video serving, social networking, big data

Example: Eventual consistency

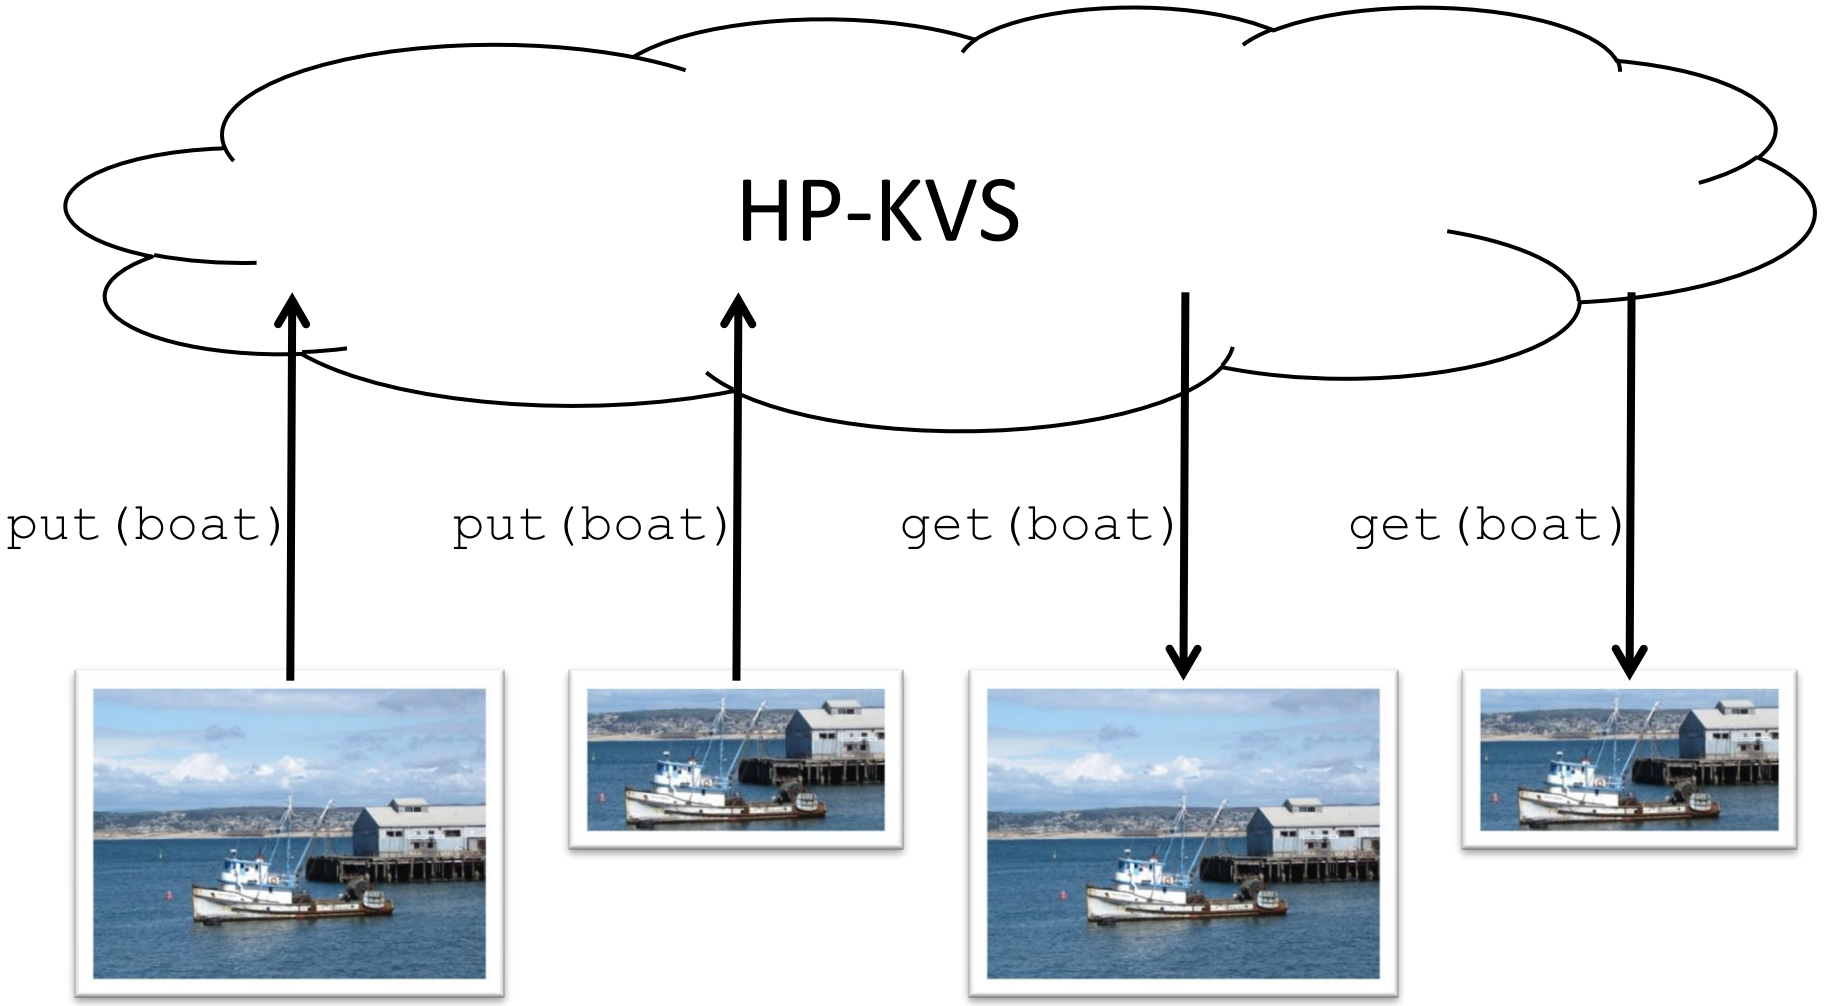

HP-KVS: Tunable and degradable consistency

– Consistency tunable per-object
  - Atomic (*strong*)
  - Regular
  - Eventual (*weak*)
  - Any
– Degradable consistency
  - Under failures may choose
  - Weaker consistency
  - *or* Unavailability

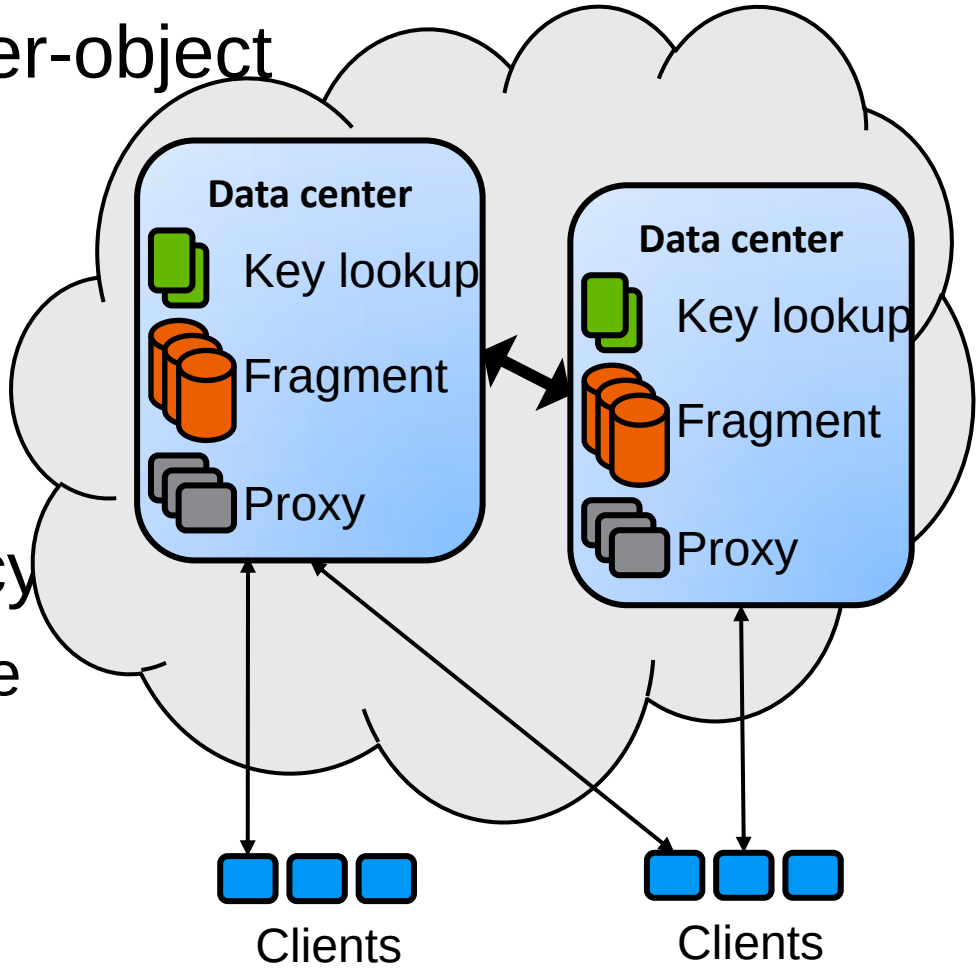

HP-KVS: Research

- Tunable and degradable consistency
  - APIs
  - SLAs
  - Measurement
- Tunable availability/reliability
  - Erasure codes (network RAID++)
  - Quorum systems (active:active, primary:failover, …)

Information-aware

Global-scale storage

NVMem & data centers

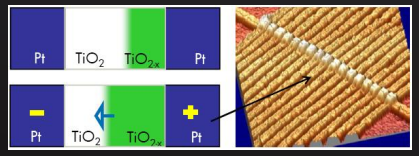

Memristors

Vision: From content to insight

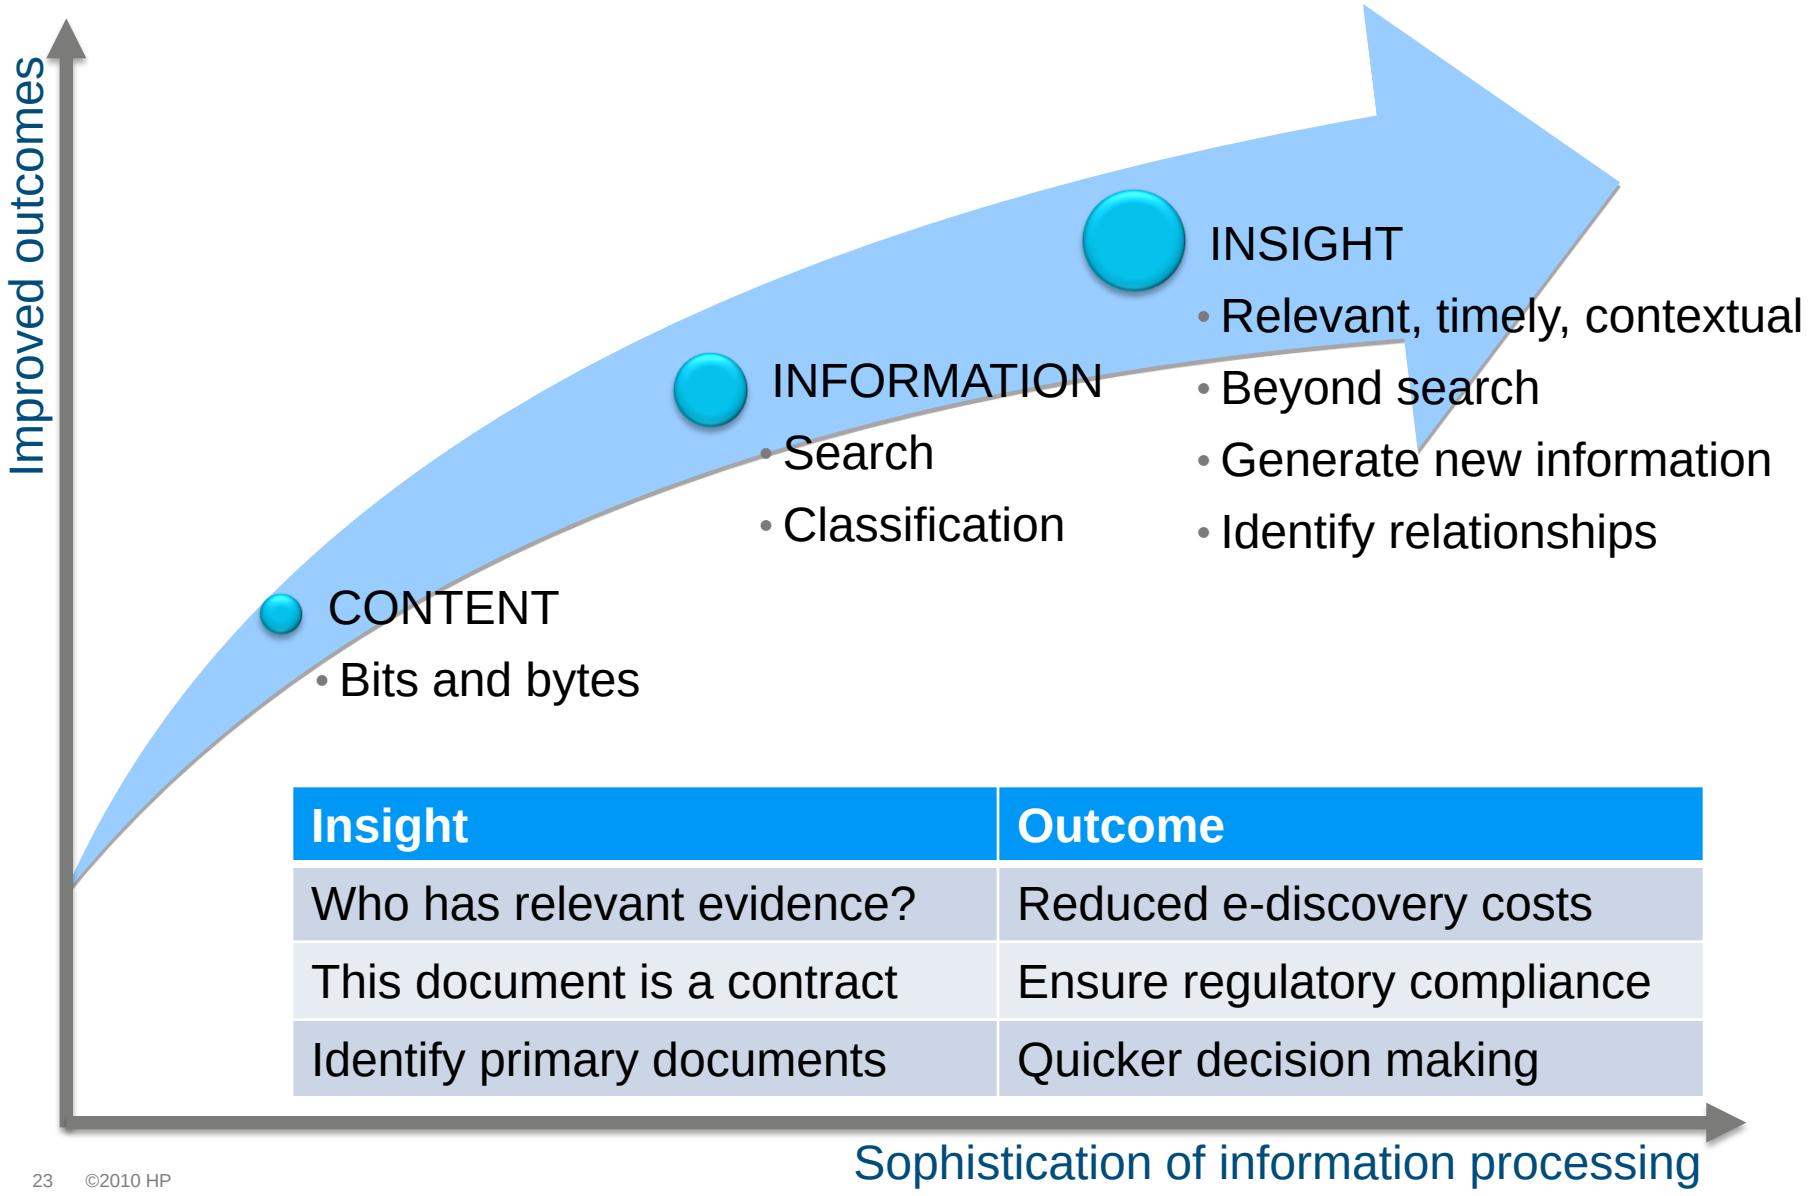

Insight	Outcome
Who has relevant evidence?	Reduced e-discovery costs
This document is a contract	Ensure regulatory compliance
Identify primary documents	Quicker decision making

Federated analytics

Principal Documents

Context analysis

Clustering & Classification

Provenance & freshness

Taxonomy builder

Metadata stores

Document analytics

Entity extraction

Keyword extraction

Document summaries

Data Stores

Web

Laptops / desktops

File servers / email / Application servers

Wikis

Research challenges: Systems research

- Information-aware data stores
  - Develop novel per-document analytics
  - Pluggable SCAN-lite analytics (Soules et al., Eurosys 2009)
- Metadata store
  - 100K+ users, 1M+ data sources, 1T+ metadata records
  - Continuously ingest metadata and execute analytics
  - LazyBase (Keeton et al., HotStorage 2009)
- Federated analytics
  - Develop novel, scalable, pan-document analytics

Research challenges: Analytics “transform information into insight”

- Clustering
  - Reveals the “inherent structure within the collection”
- Classification
  - Uses training data to automatically classify new documents
- Provenance
  - Track evolution of content over time
- Freshness
  - Finds latest document which covers the same concepts
- Summarization
  - Key sentences from document
- Principal documents identification
  - Identify *authoritative* documents

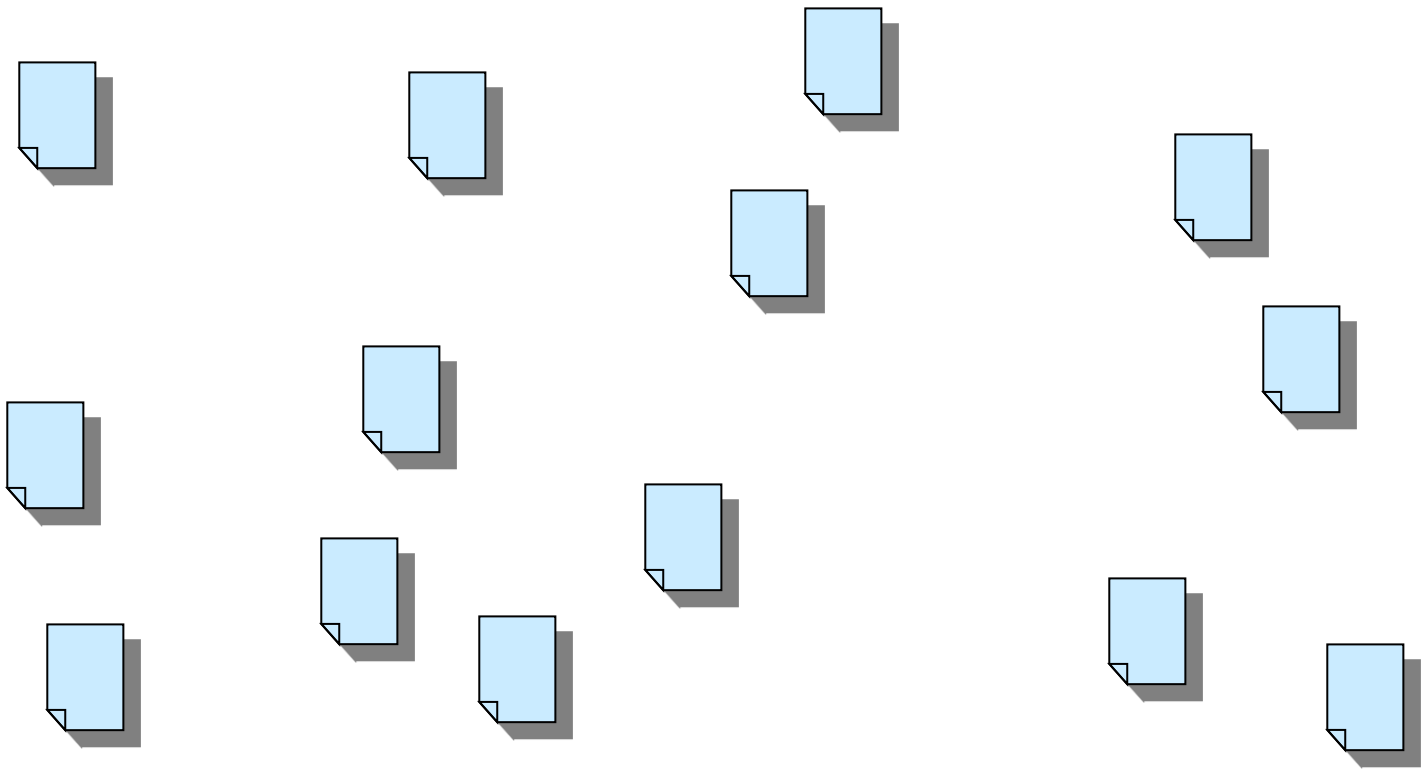

Documents in data stores

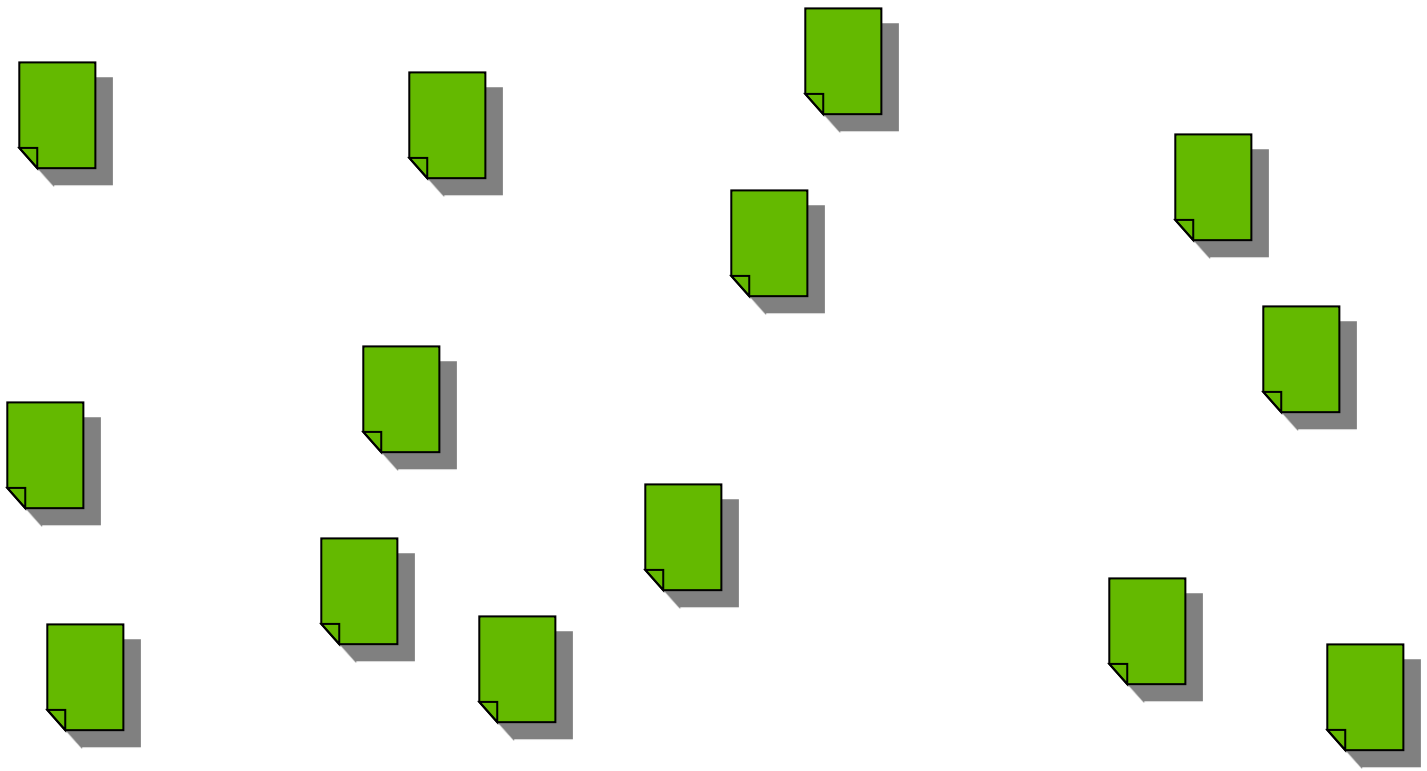

Metadata in metadata stores

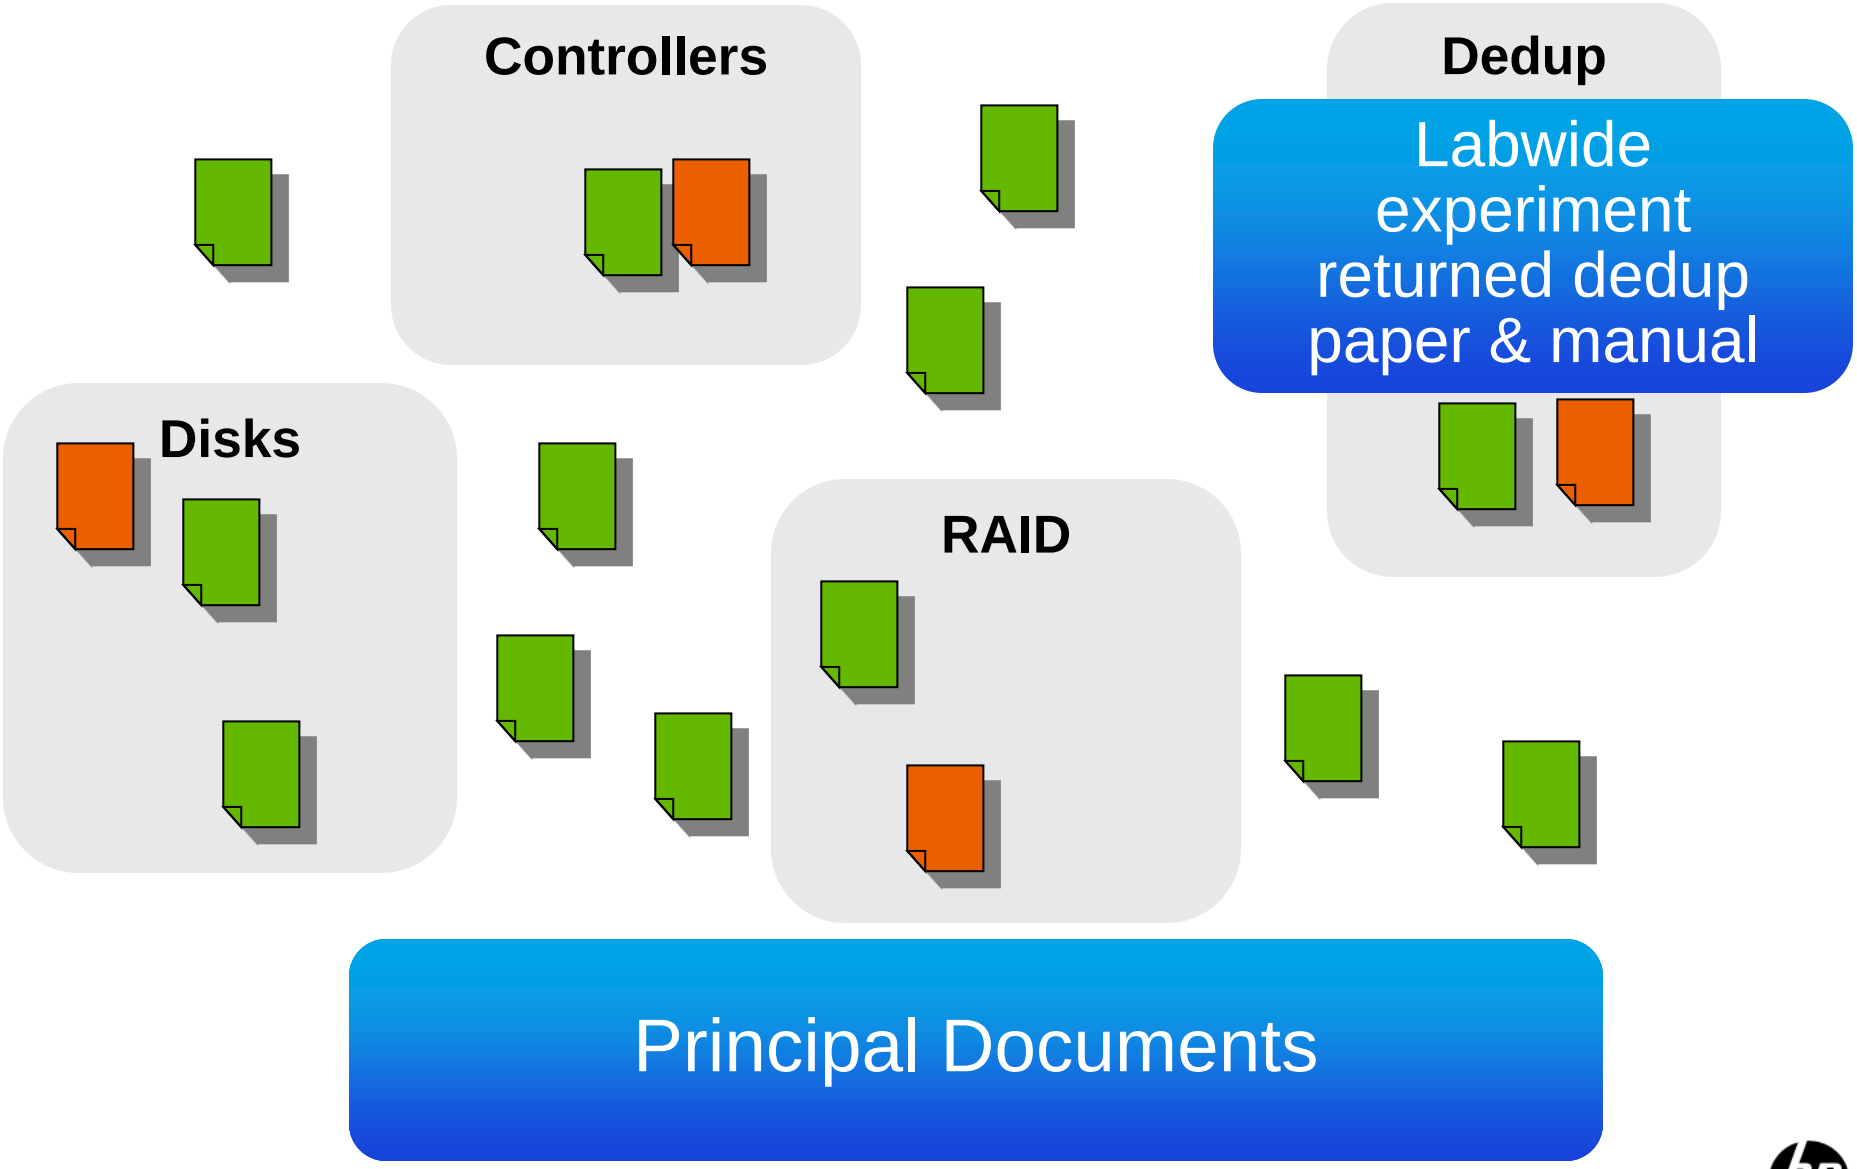
Controllers
Dedup
Labwide experiment returned dedup paper & manual
Disks
RAID
Principal Documents

Information-aware

Global-scale storage

NVMem & data centers

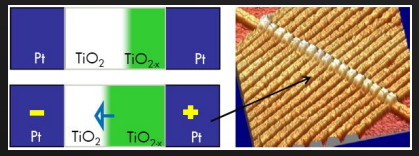

Memristors

Discussion