



The Machine: The future of technology

patrick.demichel@hp.com

Hyperscale Division EMEA

© Copyright 2014 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice.

Agenda

The background is a vibrant blue with a complex network of white lines and dots. A prominent feature is a globe-like structure composed of interconnected nodes and edges, positioned in the center-left. To the right, there are wavy, horizontal lines that resemble signal waves or data paths, also composed of nodes and edges. The overall aesthetic is futuristic and technological.

1: Vision

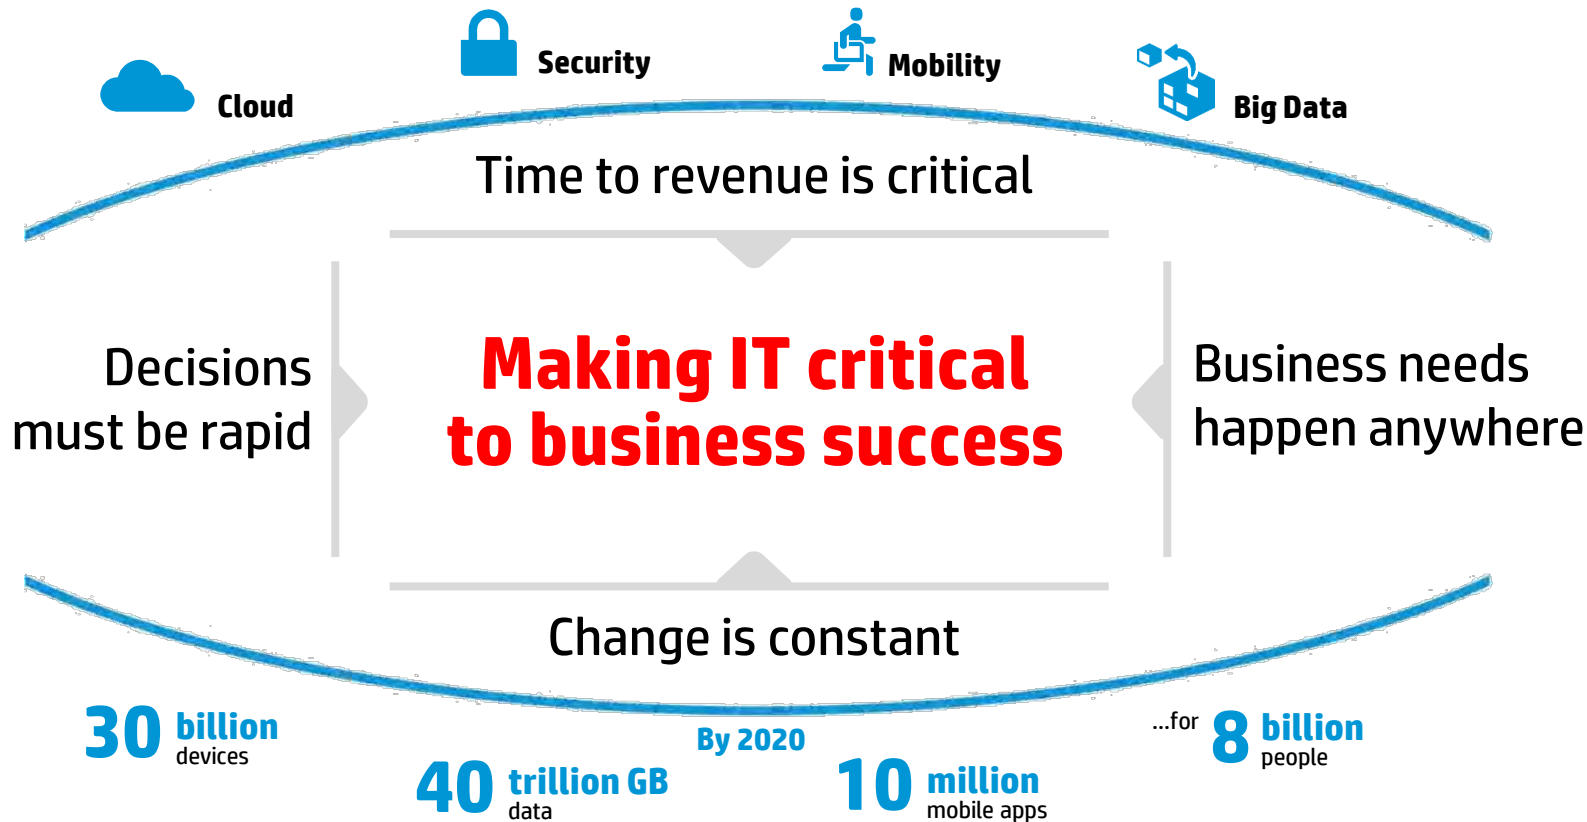
2: The core technologies

The memristor

The photonic

3: The MACHINE

The most exciting shifts of our time are underway



Today's computing infrastructure unable to keep up



You won't be able to get more capacity for less



Big Data will be too big to extract meaning from



You won't be able to move your data from where it's created – useful data may get ignored or discarded



By the time you've analyzed your data – it will be out of date



Your infrastructure will require more resources than you can get



Securing your enterprise will take more computing resources than you have

Internet of Things




Pervasive Connectivity


Smart Device Expansion

Explosion of Information

By 2020

 **200 Billion**⁽¹⁾
IoT "Things"

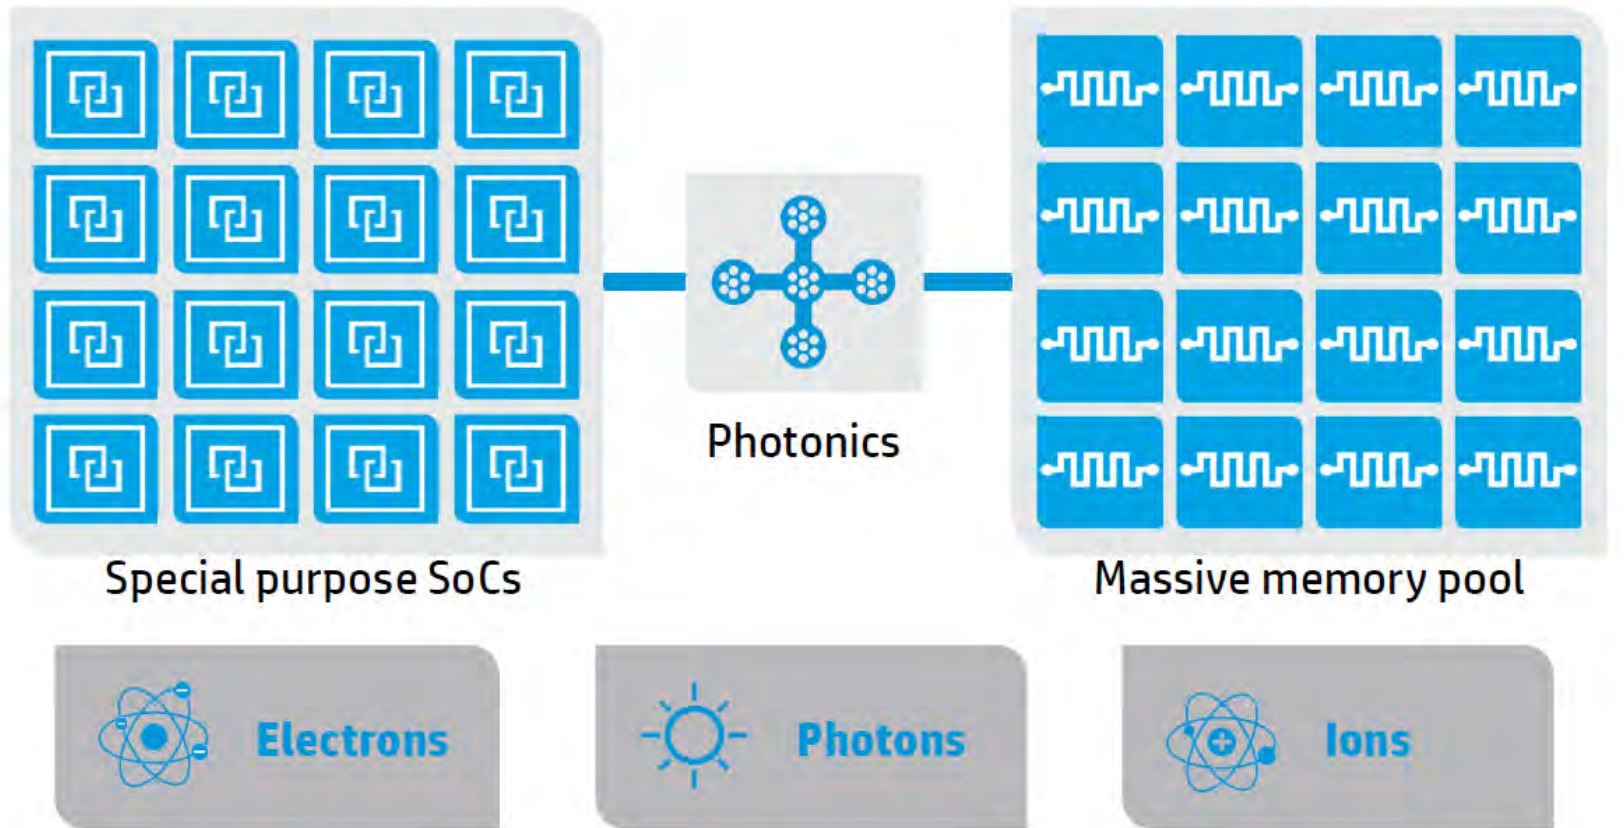
 **30 Billion**⁽²⁾
Connected Devices

 **1 Billion**⁽³⁾
Smart Meters


... for 8 Billion⁽⁴⁾ 

(1) IDC "Worldwide Internet of Things (IoT) 2013-2020 forecast" October 2013. (2) IDC "The Digital Universe of Opportunities: Rich Data and the Increasing Value of the Internet of Things" April 2014 (3) Global Smart Meter Forecasts, 2012-2020. Smart Grid Insights (Zypryme), November 2013 (4) <http://en.wikipedia.org>

3 disruptive technologies to the rescue



The memristor





Disruption #1: Non-volatile memories

Breakthrough in storage and memory technology

Scientists Create First Memristor: Missing Fourth Electronic Circuit Element

By Bryan Gardiner | April 30, 2008 | 10:03 am | Categories: Uncategorized
 Like Send 54 likes Sign Up to see what your friends like.

Researchers at HP Labs have built the first working prototypes of an important new electronic component that may lead to instant-on PCs as well as analog computers that process information the way the human brain does.



Store large amounts of data permanently like hard disks, but 100,000 times faster, and at much lower energy

Technology	Density (μm ² /bit)	Bandwidth (GB/s)	Latency Read (ns)	Latency Write (ns)	Energy Read (pJ/b)	Energy Write (pJ/b)
Hard Disk	N/A	0.5	3,000,000	3,000,000	2500	2500
Flash SSD [3] [6]	0.0021	1.0	25,000	200,000	250	250
DRAM [6] [30]	0.0038	51.2	55	55	24	24
PCRAM (22nm) [30]	0.0058	variable	48	150	2	19.2
Memristor (22nm) [8]	0.0048	variable	100	100	1-3	1-3



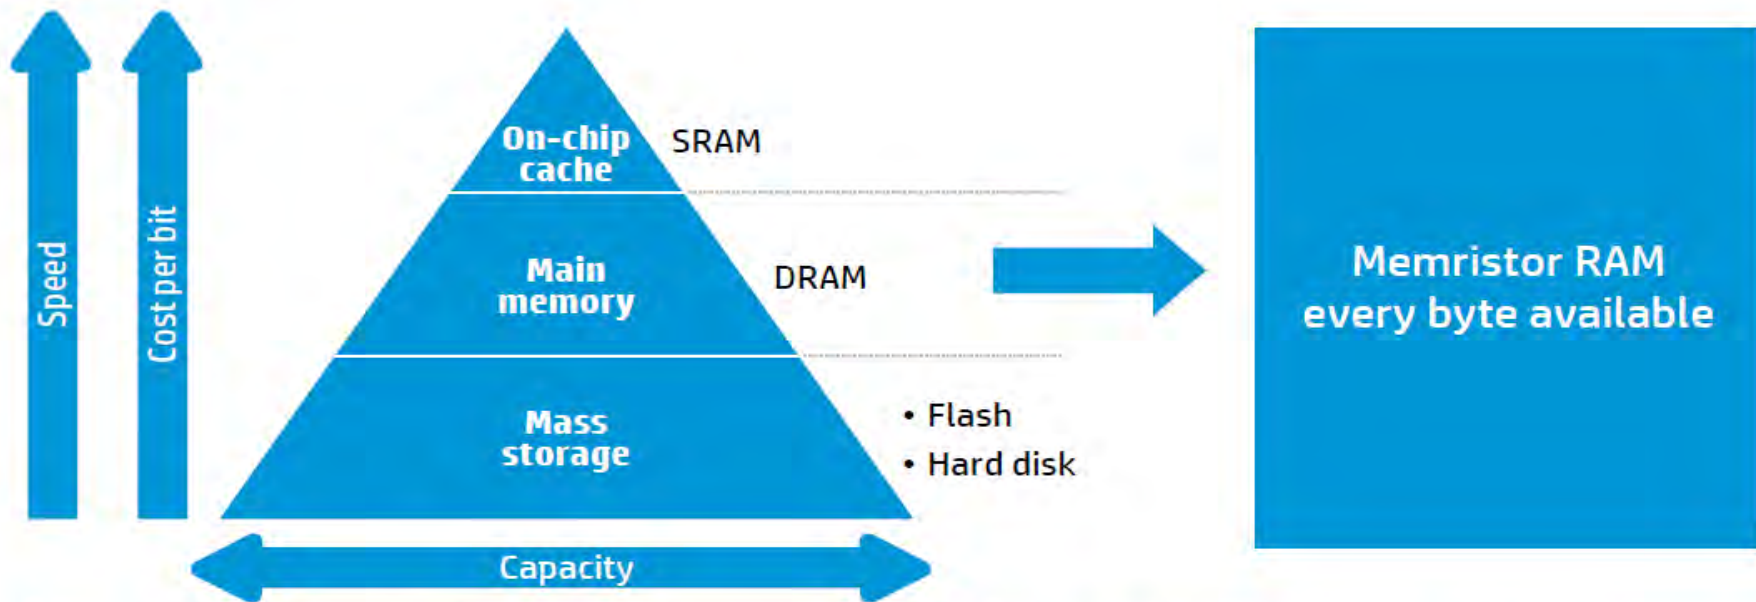
UNIVERSAL MEMORY

Massive memory pool



A drastic reduction of the memory stack complexity and cost


But requires a complete software stack redesign to leverage the full potentiality of the new architecture



Memristors change how and where data are stored

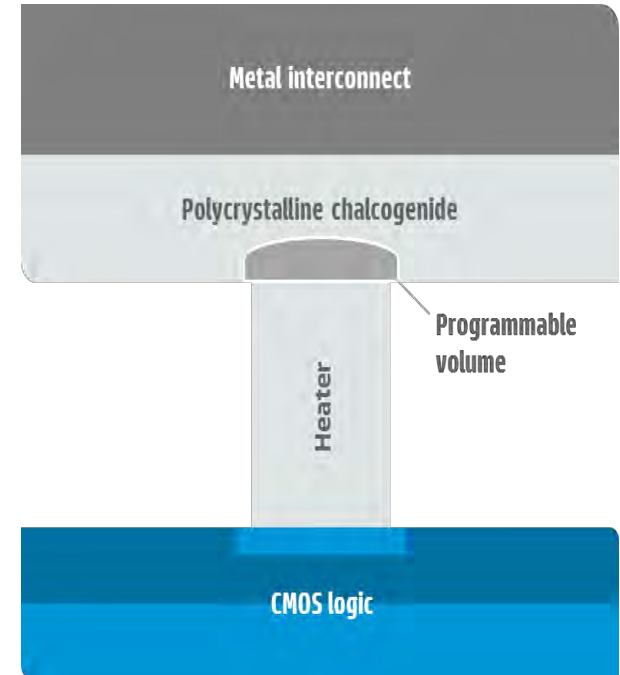


Other NVM Technologies



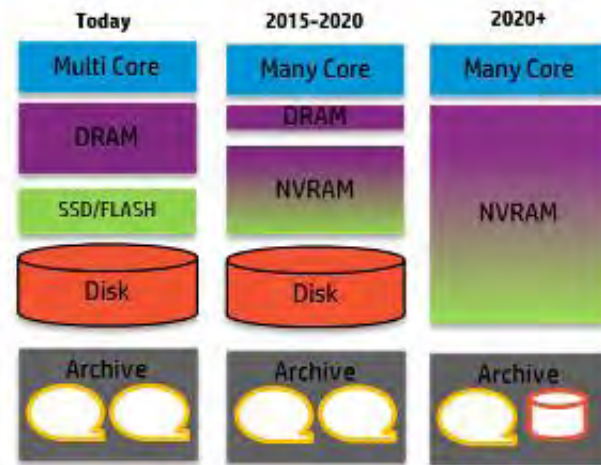
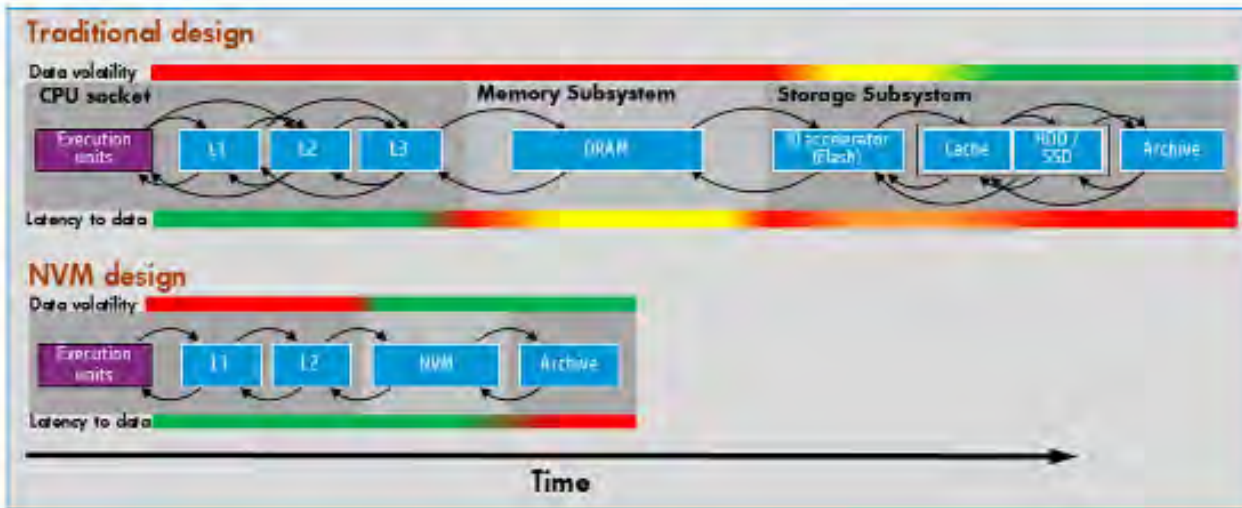
Spin-Transfer Torque (STT-RAM)

Phase-Change (PCM)



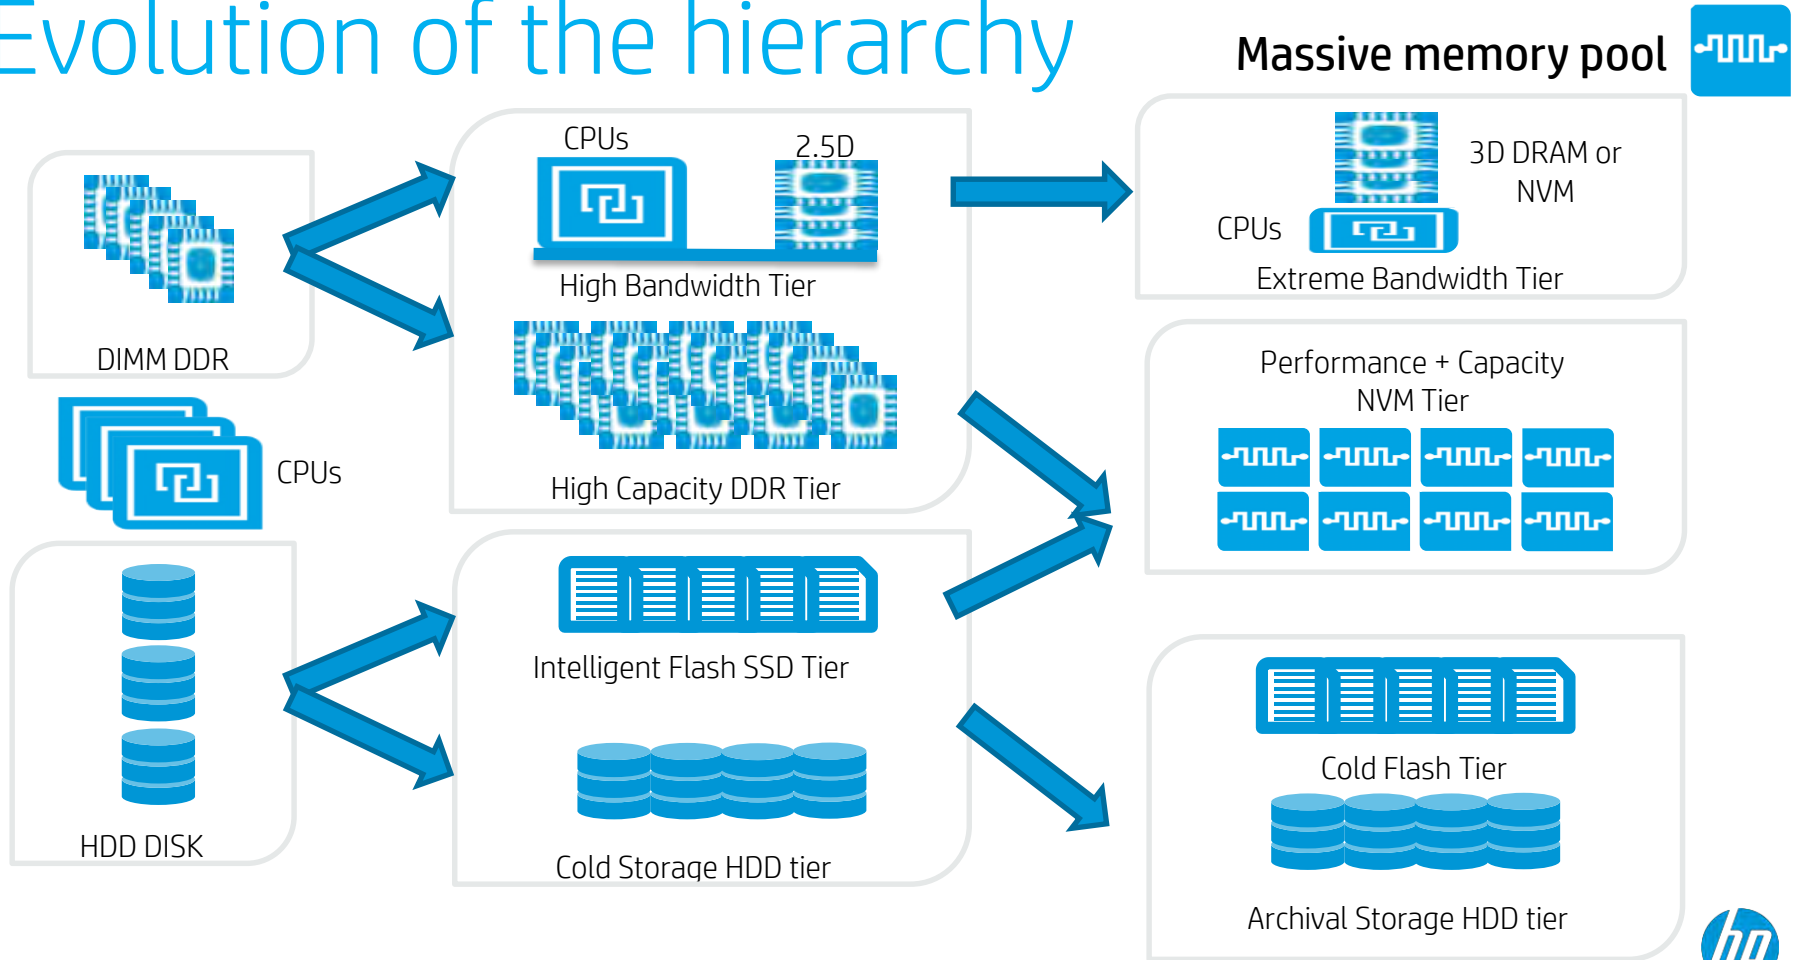
Memory Hierarchy As NVM Replaces DRAM

Step wise memory evolution to NVM



- Traditional API's designed to hide long device latencies and complex memory hierarchies will become obsolete. NVM Data Analytics can be done in-memory
- Applications such as relational databases that are structured to manage the long latencies to disk and the volatility of DRAM will be replaced by new technologies such as in-memory databases.

Evolution of the hierarchy



Benefits of universal memory

Massive memory pool

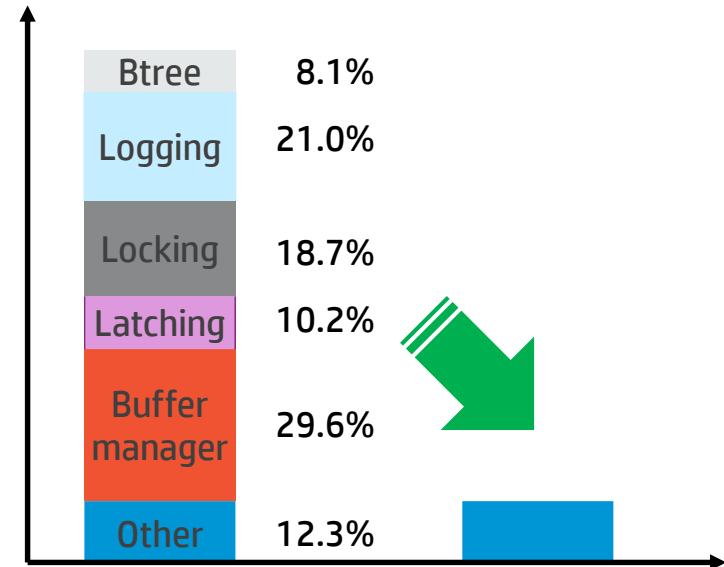


Example: a database transaction

Traditional databases struggle with big & fast data

90% of a database transaction is overhead

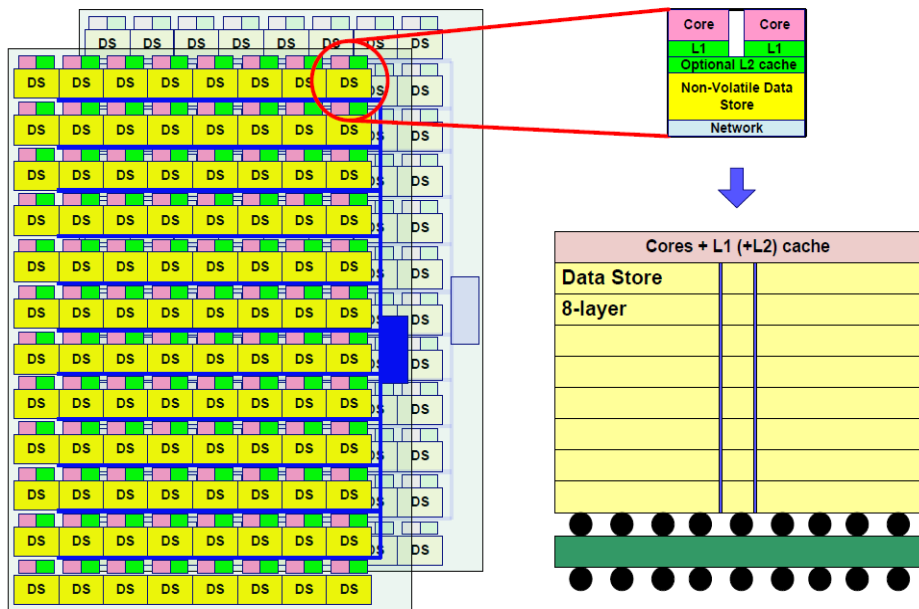
Memory-semantics nonvolatile memory: up to 10x improvement



Source: S. Harizopoulos, D. Abadi, S. Madden, and M. Stonebraker, "OLTP Through the Looking Glass, and What We Found There," *Proc. SIGMOD*, 2008.



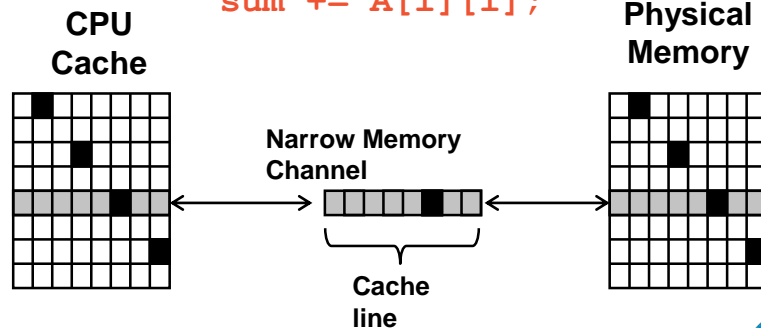
Nanostores: in-memory compute



Flat converged storage hierarchy with compute colocation for **10x-100x** improvement in performance per Watt

Example: Matrix Computation

```
for (i = 0; i < n; i++)  
    sum += A[i][i];
```




Technologies for Check-point Restart

www.nd.edu/~rich/SC09/tut157/SC2009_Jouppi_Xie_Tutorial_Final.pdf




Architecture

PCRAM





The schematic view of a PCRAM cell with NMOS access transistor (BL=Bitline, WL=Wordline, SL=Source line)

	HDD	NAND Flash	PCRAM
Taille cellule	-	4-6F ²	4-6F ²
Cycle lecture	~4ms	5us-50us	10ns-100ns
Cycle écriture	~4ms	2ms-3ms	100-1000ns
Watt à arrêt	~1W	~0W	~0W
Endurance cycles	10 ¹⁵	10 ⁵	10 ⁸



Memristor

CMOS chip avec des composants memrésistifs

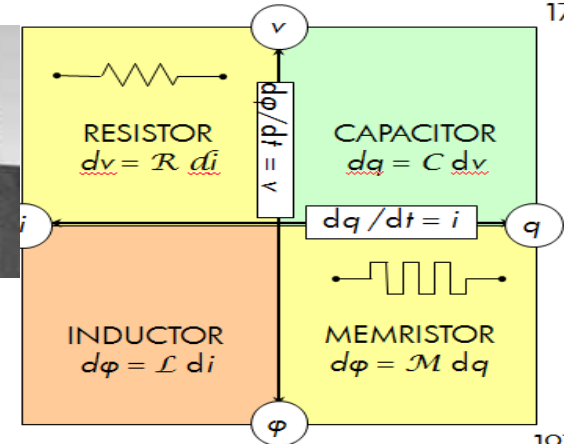


Ohm
1827

Von Kleist
1745




L. O. Chua, (1971)



1831
Faraday

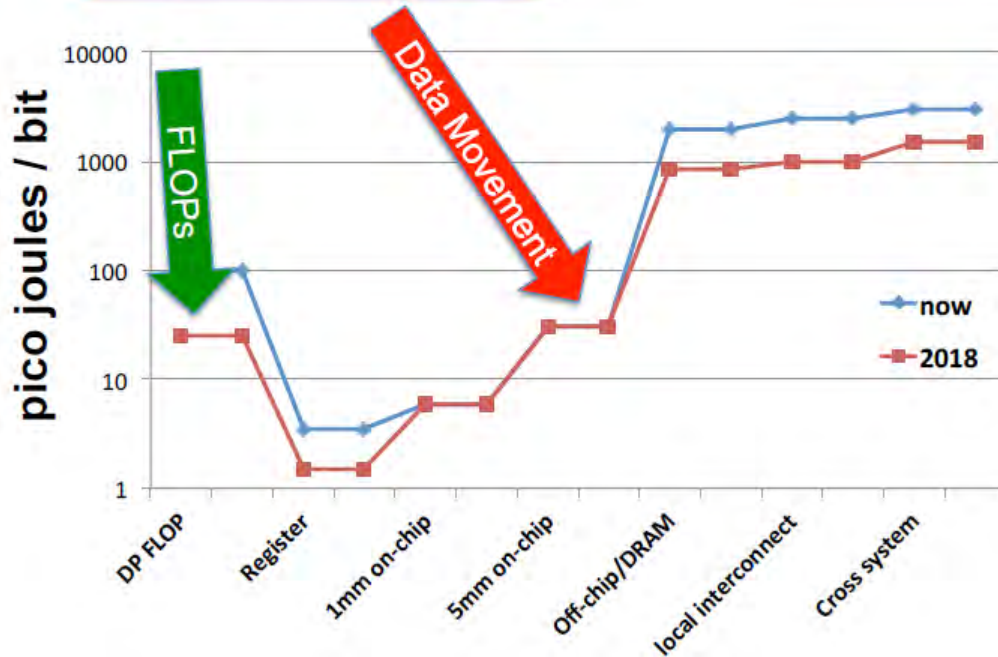
1971
Chua

Photonic



Why photonics?

FLOPs will cost less than on-chip data movement! (NUMA)



10¹⁸ ops*
1 Byte/ops =
10¹⁹ bits*
1 pj/bit =
10 MWatts!!




Photonics technologies

Communication fire hose for memristor stores

Why Photonics?



- Huge increases in the *volume* of data
- Enables efficient access to that data
- Shrink *time* and *space* to gain immediate access without regard for location

Transmit data using light for 30-fold more bandwidth at one-tenth the energy



Short term:
short range, low cost VCSEL

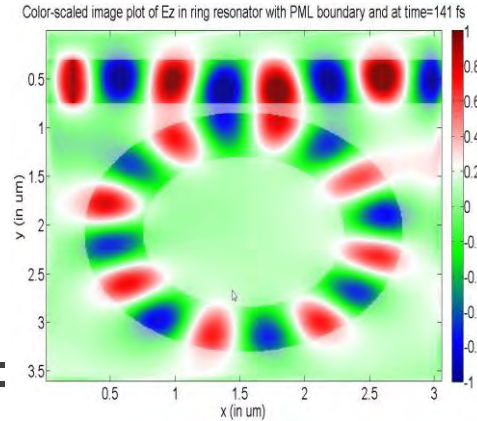

Color-scaled image plot of E_z in ring resonator with PML boundary and at time=141 fs



Long term: micro-ring resonator
(low cost, long distance, integrated on silicon)

Ring Resonators

One basic structure, 3 applications



A **modulator** – move in and out of resonance to modulate light on adjacent waveguide

A **switch** – transfers light between waveguides only when the resonator is tuned

A wavelength specific **detector** – add a doped junction to perform the receive function

Microrings

Full link configuration

Advantages

Modulators wavelength specific, no additional mux

Same ring structure used for drop filters

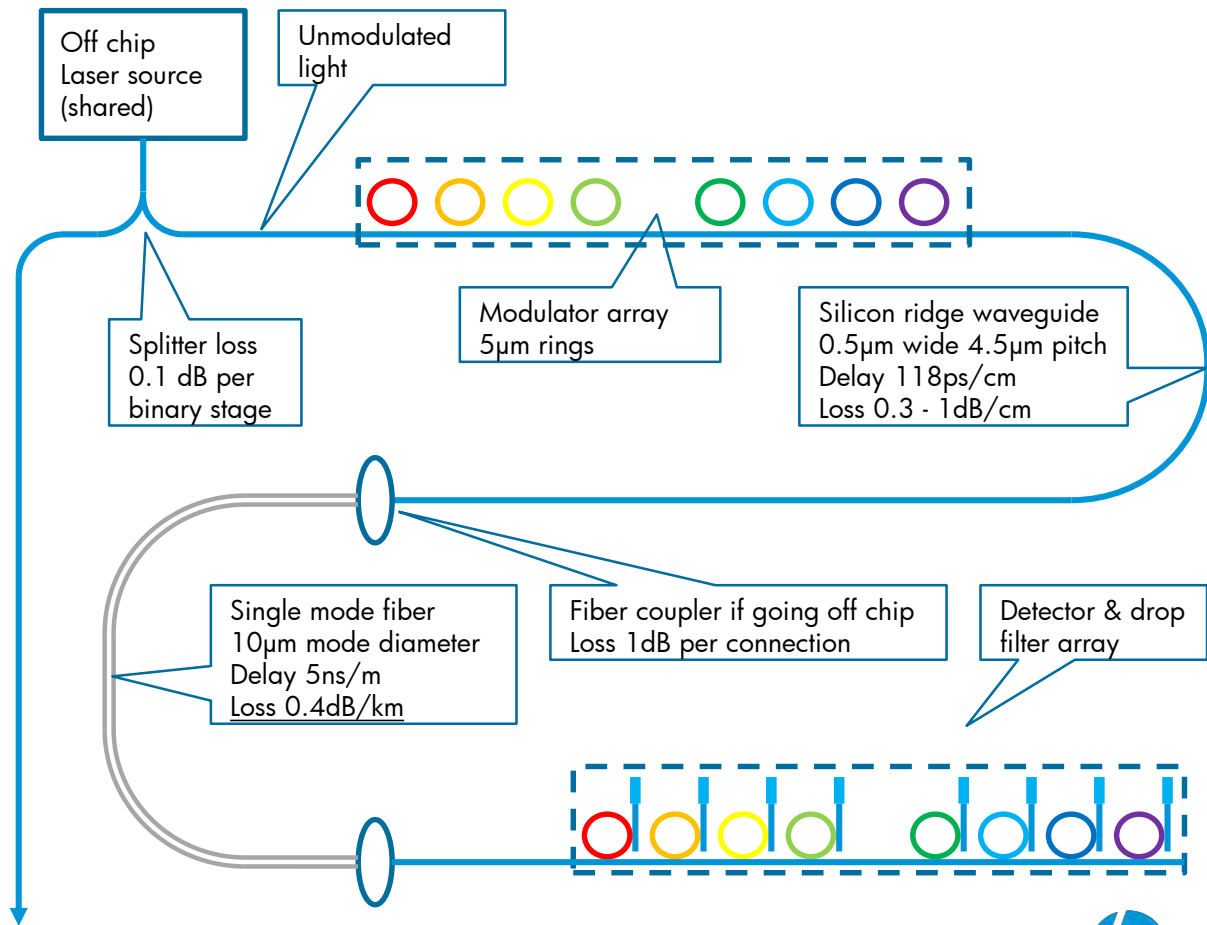
Loss budget dominated by cost

Up to 64 wavelengths

Outstanding issues

Ring tuning

Thermal stability





Photonics destroys distance

System On Chip

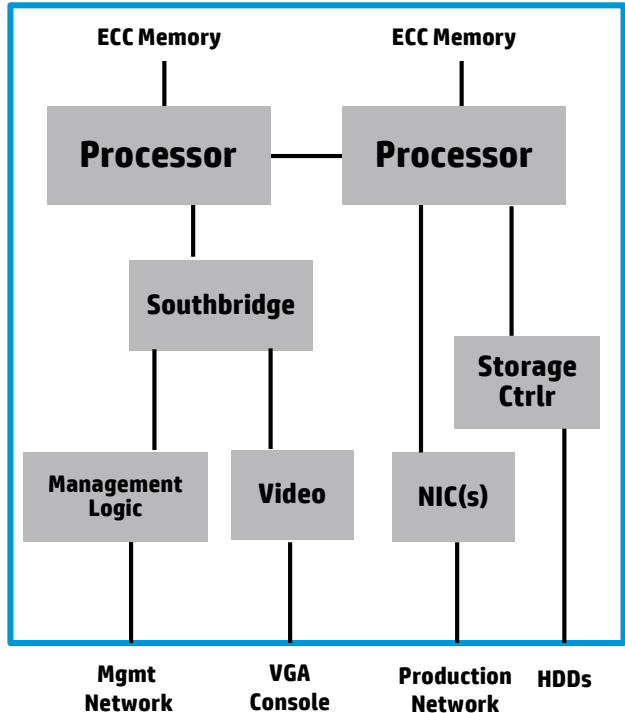


Application-focused silicon

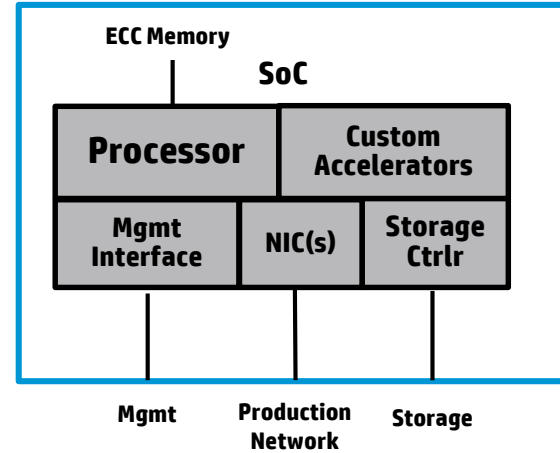
Special purpose SoCs



Traditional Server Motherboard



System on a Chip (SoC)-based Server



- Less general-purpose, more workload focused
- Dramatic reduction in power, cost, and space
- SoC vendors bring their own differentiated features and opportunities to disrupt markets



HP Moonshot

Special purpose SoCs



HP Moonshot System

**Optimize your
application performance**



A system that truly fits the
workload it is running

**Realize breakthrough
economics**



Redefining the economics of
your data center

**Accelerate business
innovation**

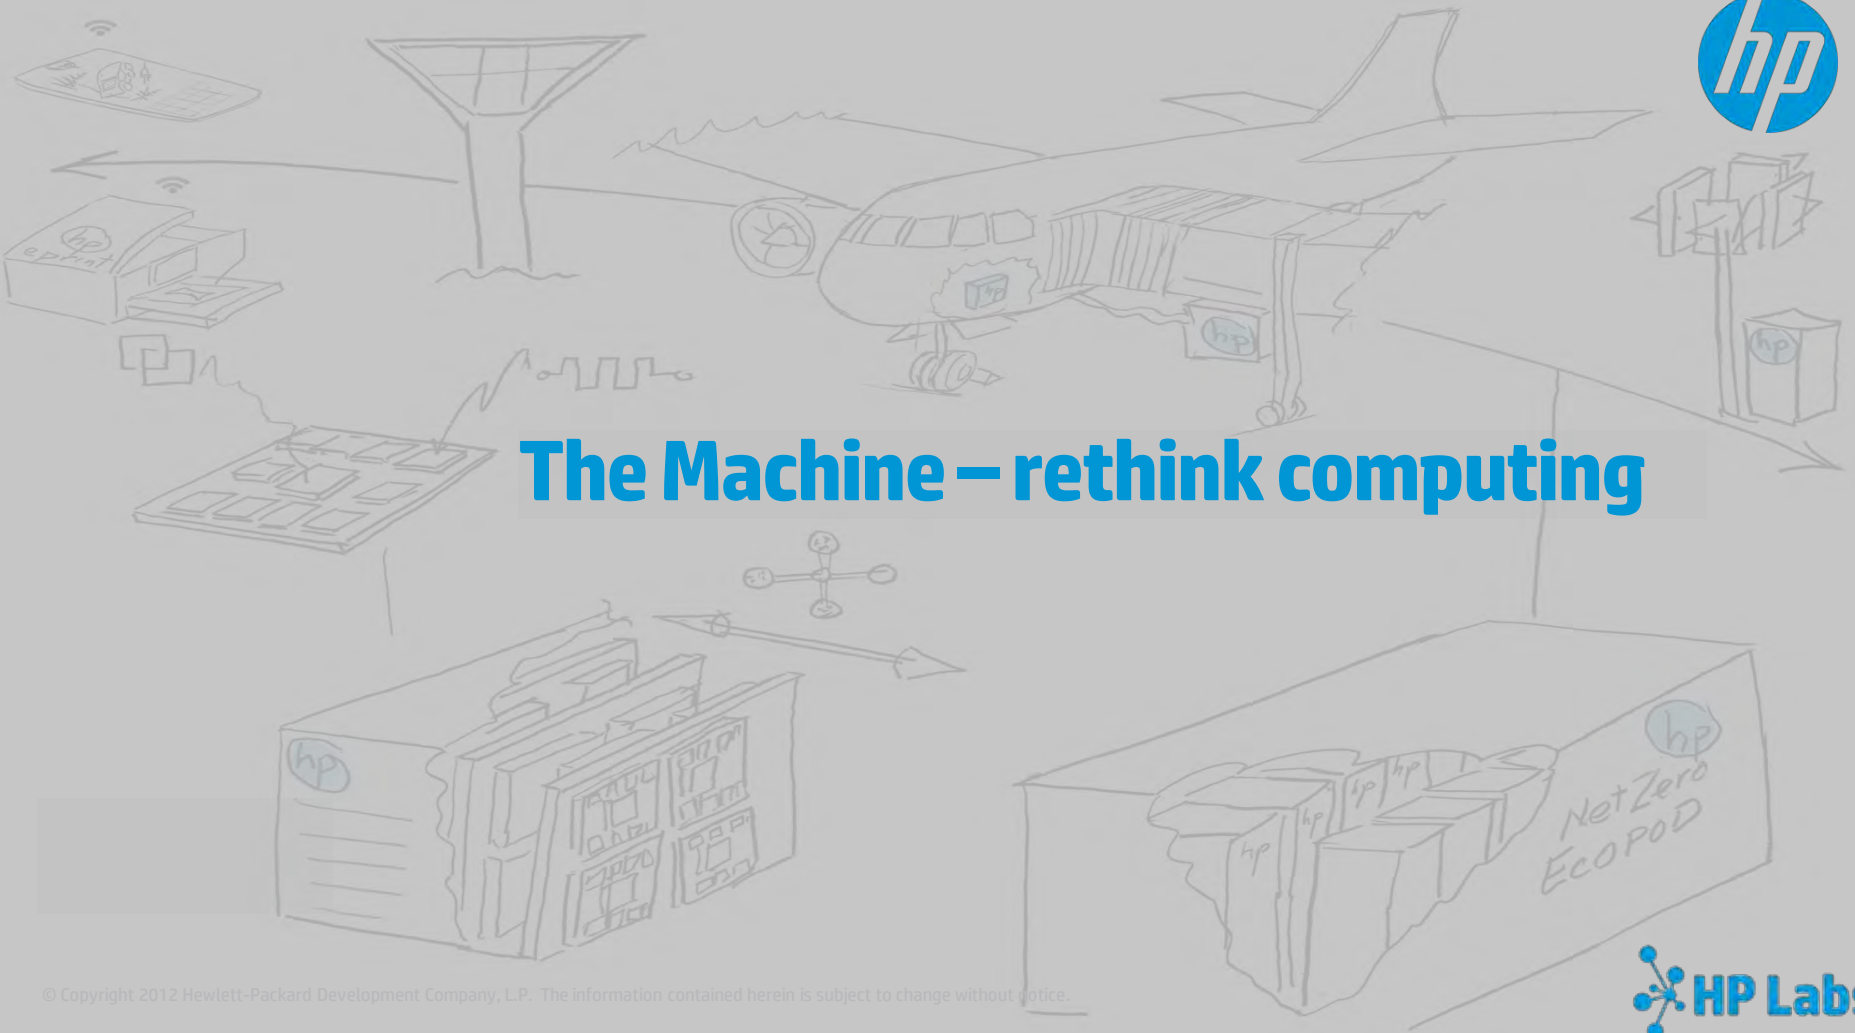


Enabling you to be more
competitive





The Machine – rethink computing



6 words to summarize the vision



Electrons



Compute



Photons



Communicate



Ions



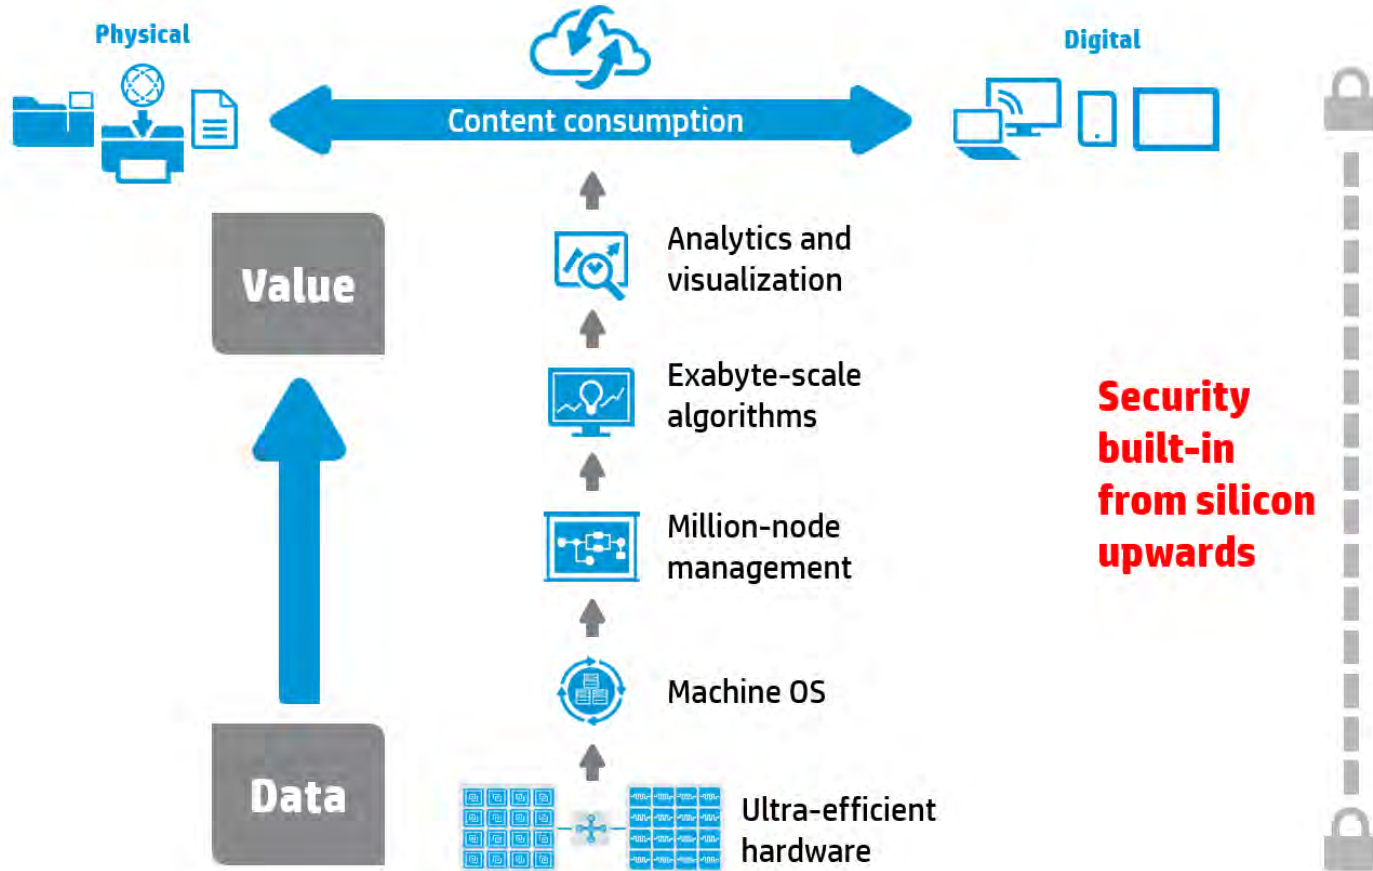
Store

Not substitutional technologies

Holistic re-architecting to get all benefits



The Machine: towards a new computing paradigm




The Machine use case



Use case: aircraft sensors

Internet-of-Things big data affects all industries



$$20 \text{ TB} \times 2 \times 3 \times 25,000 \times 365$$

20 terabytes of information per engine per hour

twin-engine aircraft


three-hour average flight duration

commercial flights per day (USA)


days in a year

$$= 1,095,000,000 \text{ TB}$$
$$= 1 \text{ ZB}$$

Use case: a mesh of connected aircrafts ...



Aspirational History



- SoC Partners selected for co-development
- Machine OS development begins

- Memristors begin sampling
- Physical infrastructure of Core prototypes established
- Open Source Machine OS SDK and emulators released
- ISV Partner collaborations begin

- Edge devices ship in volume
- Core Machines running real-world workloads at scale
- Machine OS released

- Core devices at volume
- Machine available as product, service, and as a business process transformation



2014

2015

2016

2017

2018

2019

2020


- Memristor DIMMs launched
- Integrated core technologies demonstrated

- Edge devices begin sampling
- Machine OS enters public beta

Distributed mesh compute goes mainstream



The MACHINE



[The Machine Webpage](#)

[The Machine 3 min video](#)

[Memristor Lab Tour](#)

[Photonics Lab Tour](#)

[HP Analytics Lab](#)

[HP Security and Cloud Lab](#)

This changes everything



Any questions?