From Edge to Core

Memory-Driven Hardware and Software Co-design
for the intelligent enterprise

The Second Coming

– THE SECOND COMING – W. B. Yeats - 1919

Turning and turning in the widening gyre
The falcon cannot hear the falconer;
Things fall apart; the centre cannot hold;
Mere anarchy is loosed upon the world,
The blood-dimmed tide is loosed, and everywhere
The ceremony of innocence is drowned;
The best lack all conviction, while the worst
Are full of passionate intensity.

Surely some revelation is at hand;
Surely the Second Coming is at hand.
The Second Coming! Hardly are those words out
When a vast image out of Spiritus Mundi
Troubles my sight: a waste of desert sand;
A shape with lion body and the head of a man,
A gaze blank and pitiless as the sun,
Is moving its slow thighs, while all about it
Wind shadows of the indignant desert birds.

The darkness drops again but now I know
That twenty centuries of stony sleep
Were vexed to nightmare by a rocking cradle,
And what rough beast, its hour come round at last,
Slouches towards Bethlehem to be born?

Things fall apart;
the centre cannot hold;

Oh, Inverted World!

Compute Abundance	Memory Abundance
Moore’s Law Scaling	Photonic Scaling
Imperative	Declarative
Programming	Training
Centralized	Distributed
General Purpose	Precision
Data as a Cost Center	Data as a Profit Center

The New Normal: Compute is not keeping up

Microprocessors

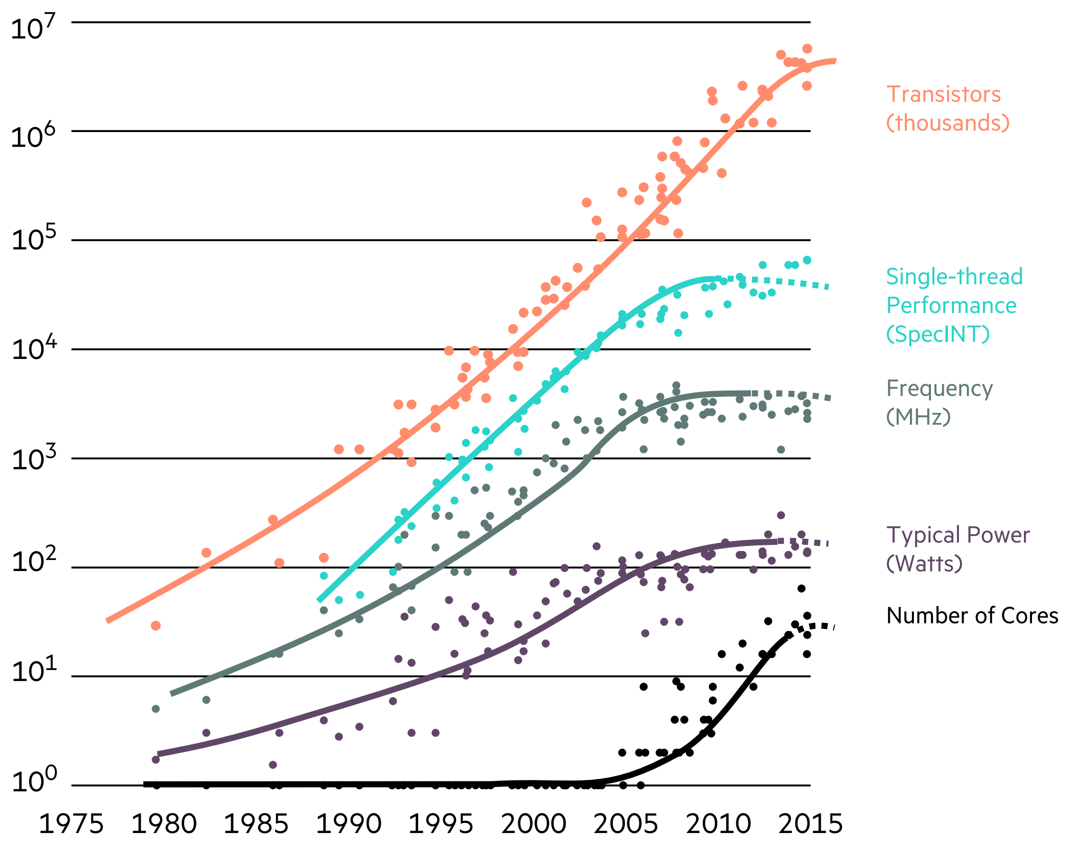

Data growth

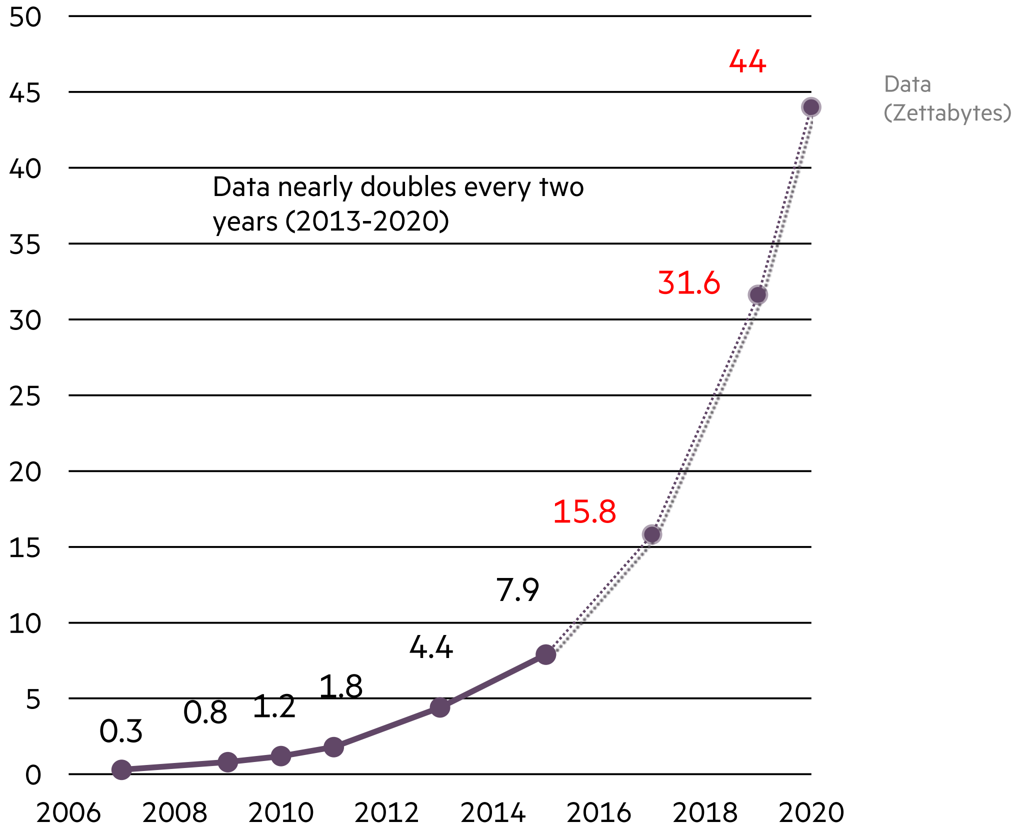

Hewlett Packard Enterprise

The end of scaling at just the wrong or just the right time …

(**8B** people x **20B** mobile devices x **100B** social infrastructure x **1T** apps)

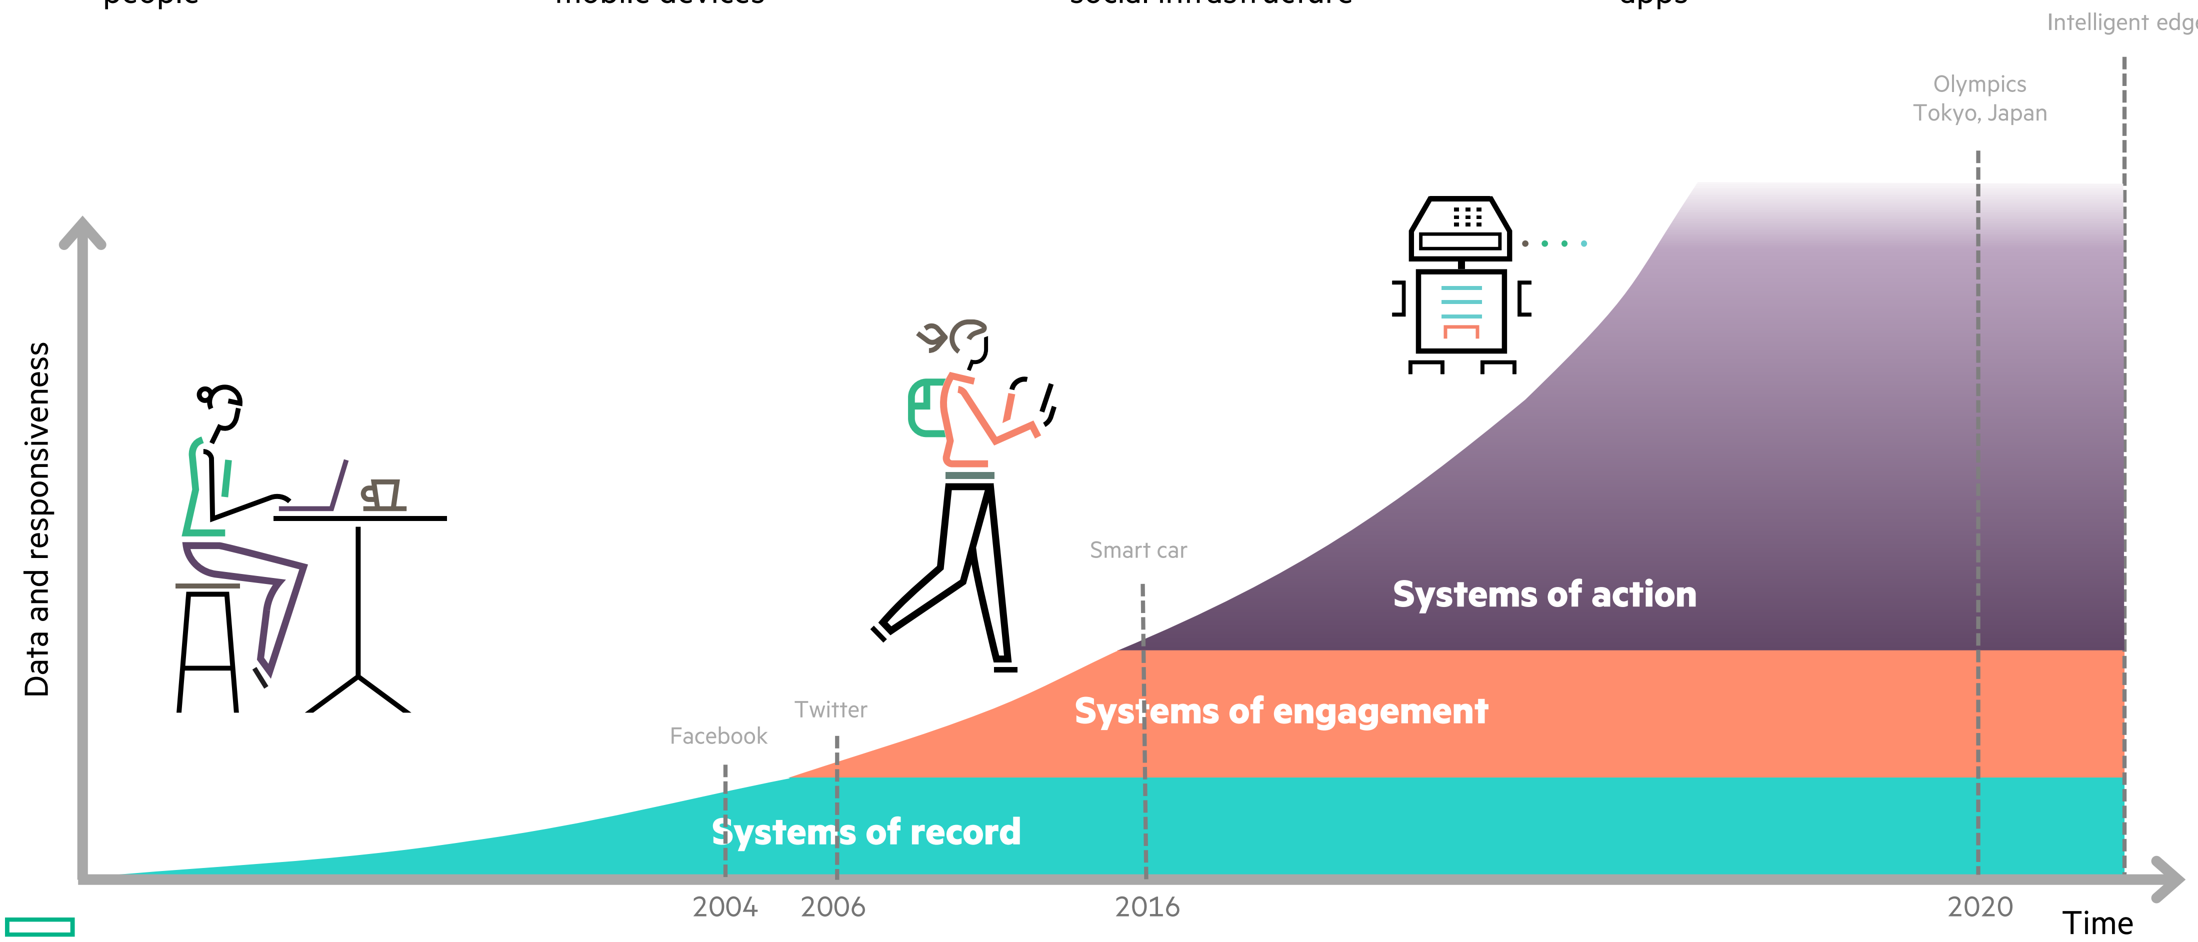

Hewlett Packard Enterprise

Opportunity - the hyper-competitive digital enterprise

The hyper-competitive digital enterprise:

- Understands data is the new source of competitiveness and economic value creation, of equal or greater value than the underlying commodity or process
- Instruments every physical or digital product, every manufacturing process in the factory, every business process in the enterprise to produce data
- Pushes analytic and machine learning capability as close as possible to the edge for real time insight and action
- Forges a continuum from the enterprise core to the intelligent edge
- Relentlessly and remorselessly turns raw data to economic advantage via process improvement, investment strategy, customer satisfaction, market expansion, warranty reduction, direct monetization

Re-establishing innovation horizons

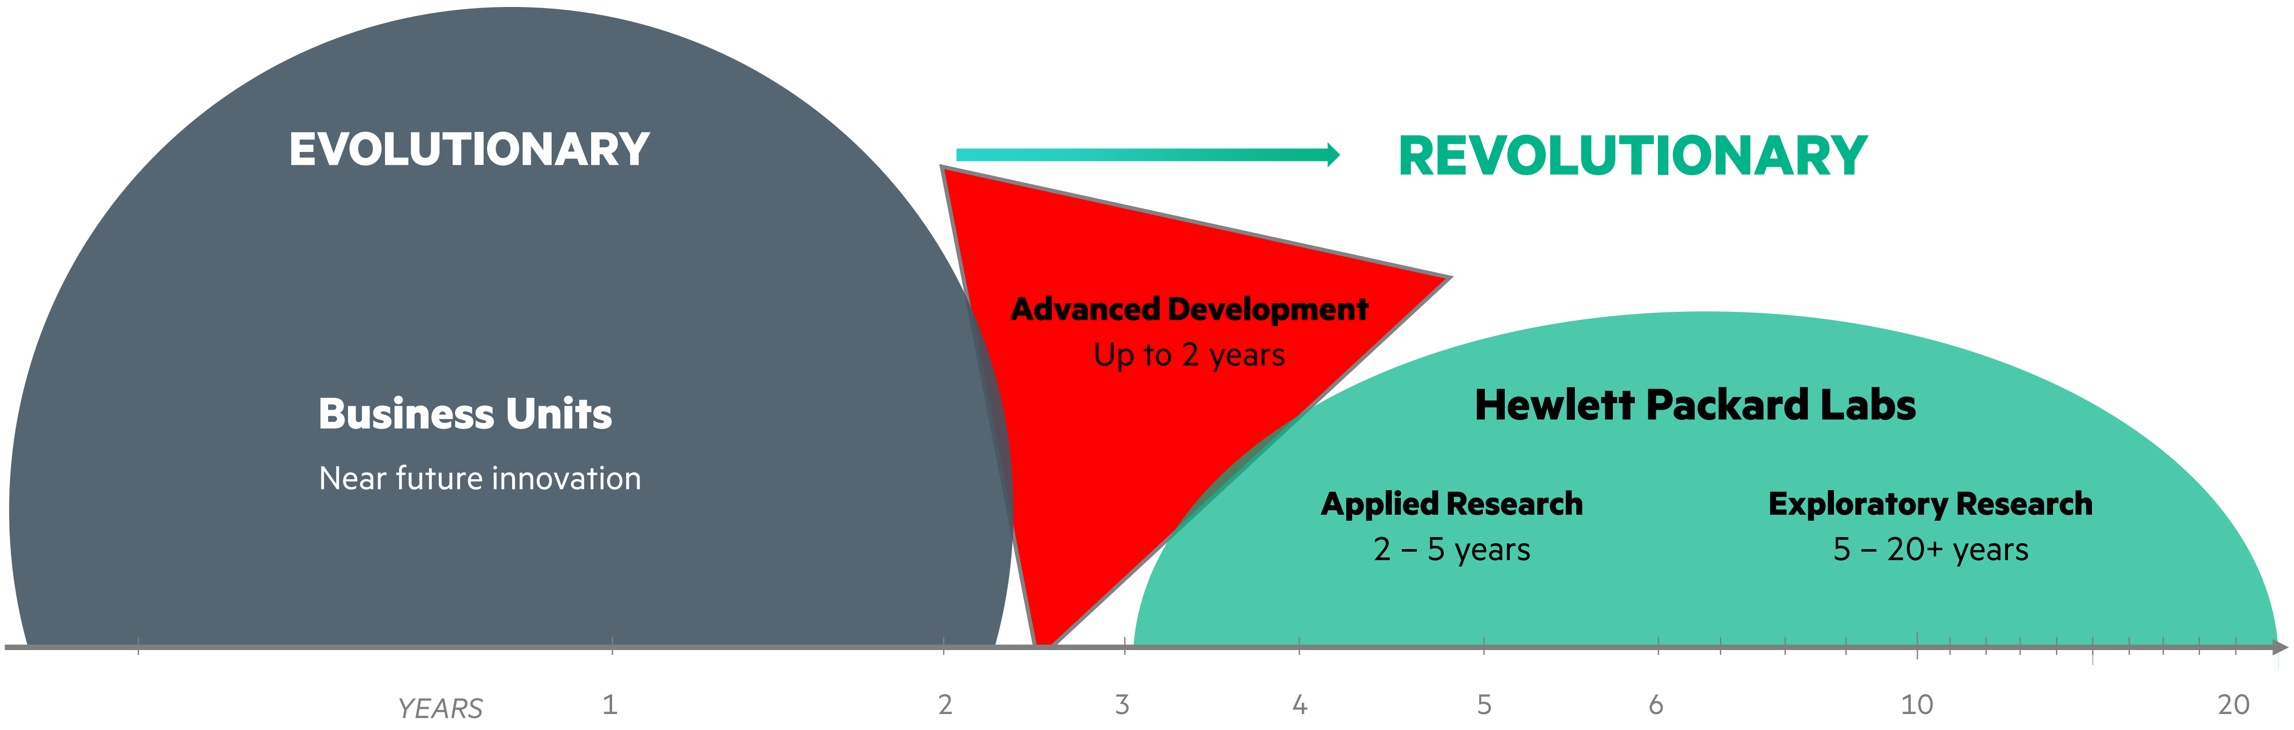

Inversion as Architecture

Today's architecture

From processor-centric computing

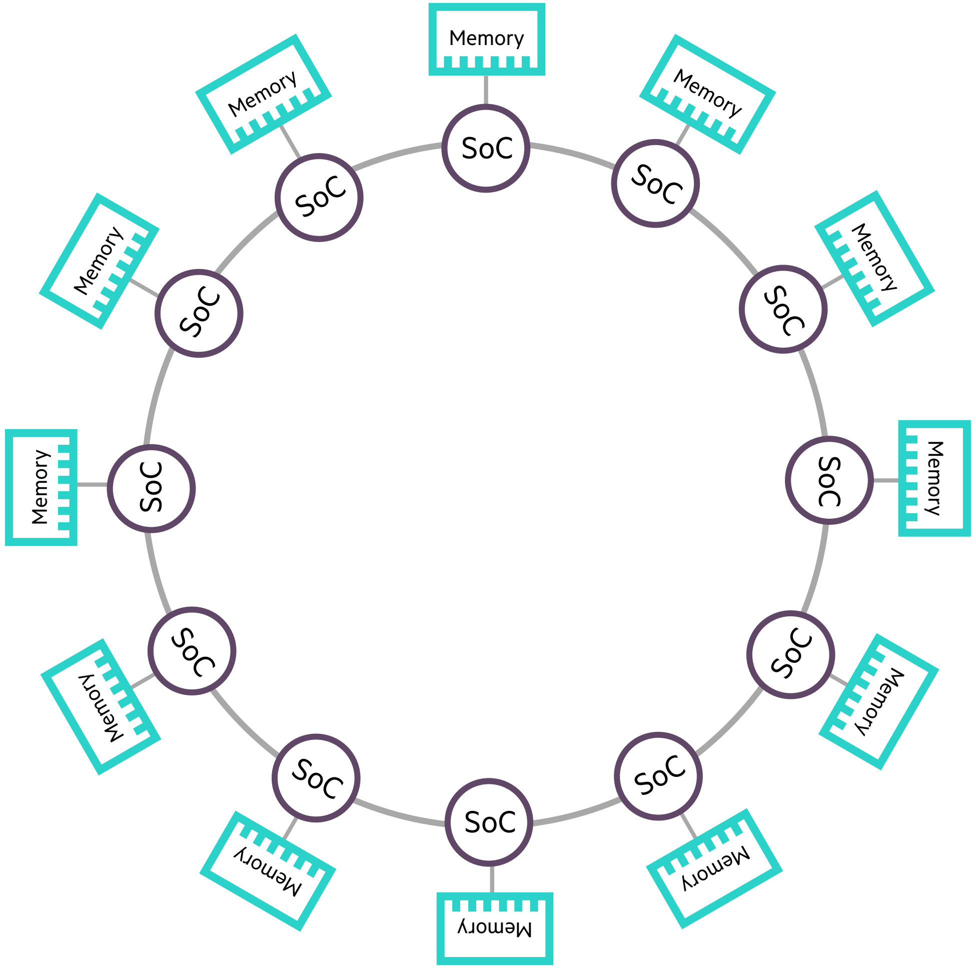

Future architecture

Memory-Driven Computing

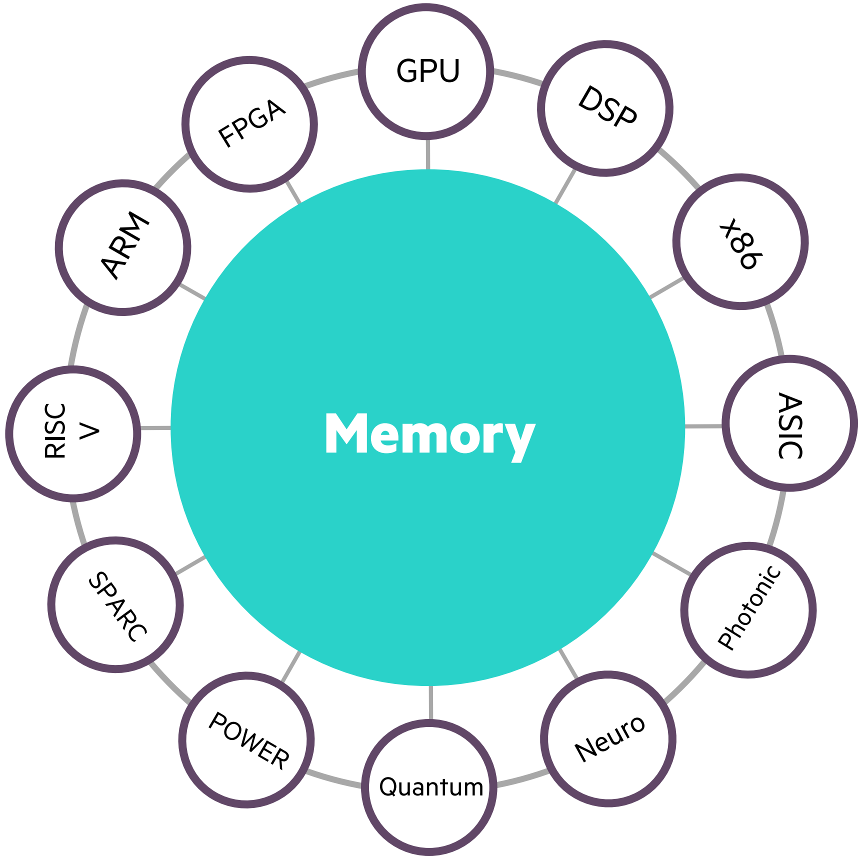

HPE introduces the world's largest single-memory computer

The prototype contains 160 terabytes of memory

- 160 TB of shared memory spread across 40 physical nodes, interconnected using a high-performance fabric protocol.

- An optimized Linux-based operating system running on ThunderX2, Cavium's flagship second generation dual socket capable ARMv8-A workload optimized System on a Chip.

- Photonics/Optical communication links, including the new X1 photonics module, are online and operational.

- Software programming tools designed to take advantage of abundant of persistent memory.

Hewlett Packard Enterprise

HPE's X1: Fully integrated photonics interconnect chip module

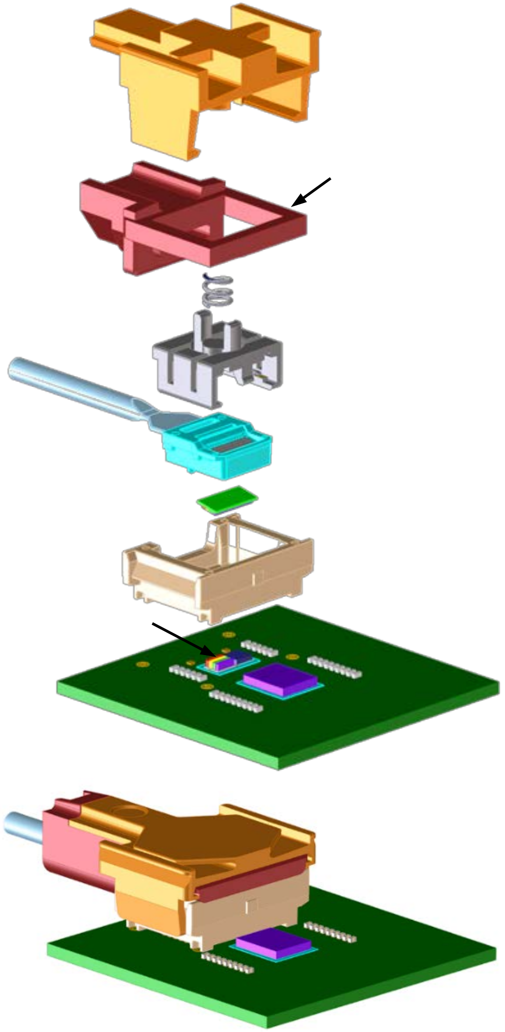

The current prototype

Gen-Z
Proprietary

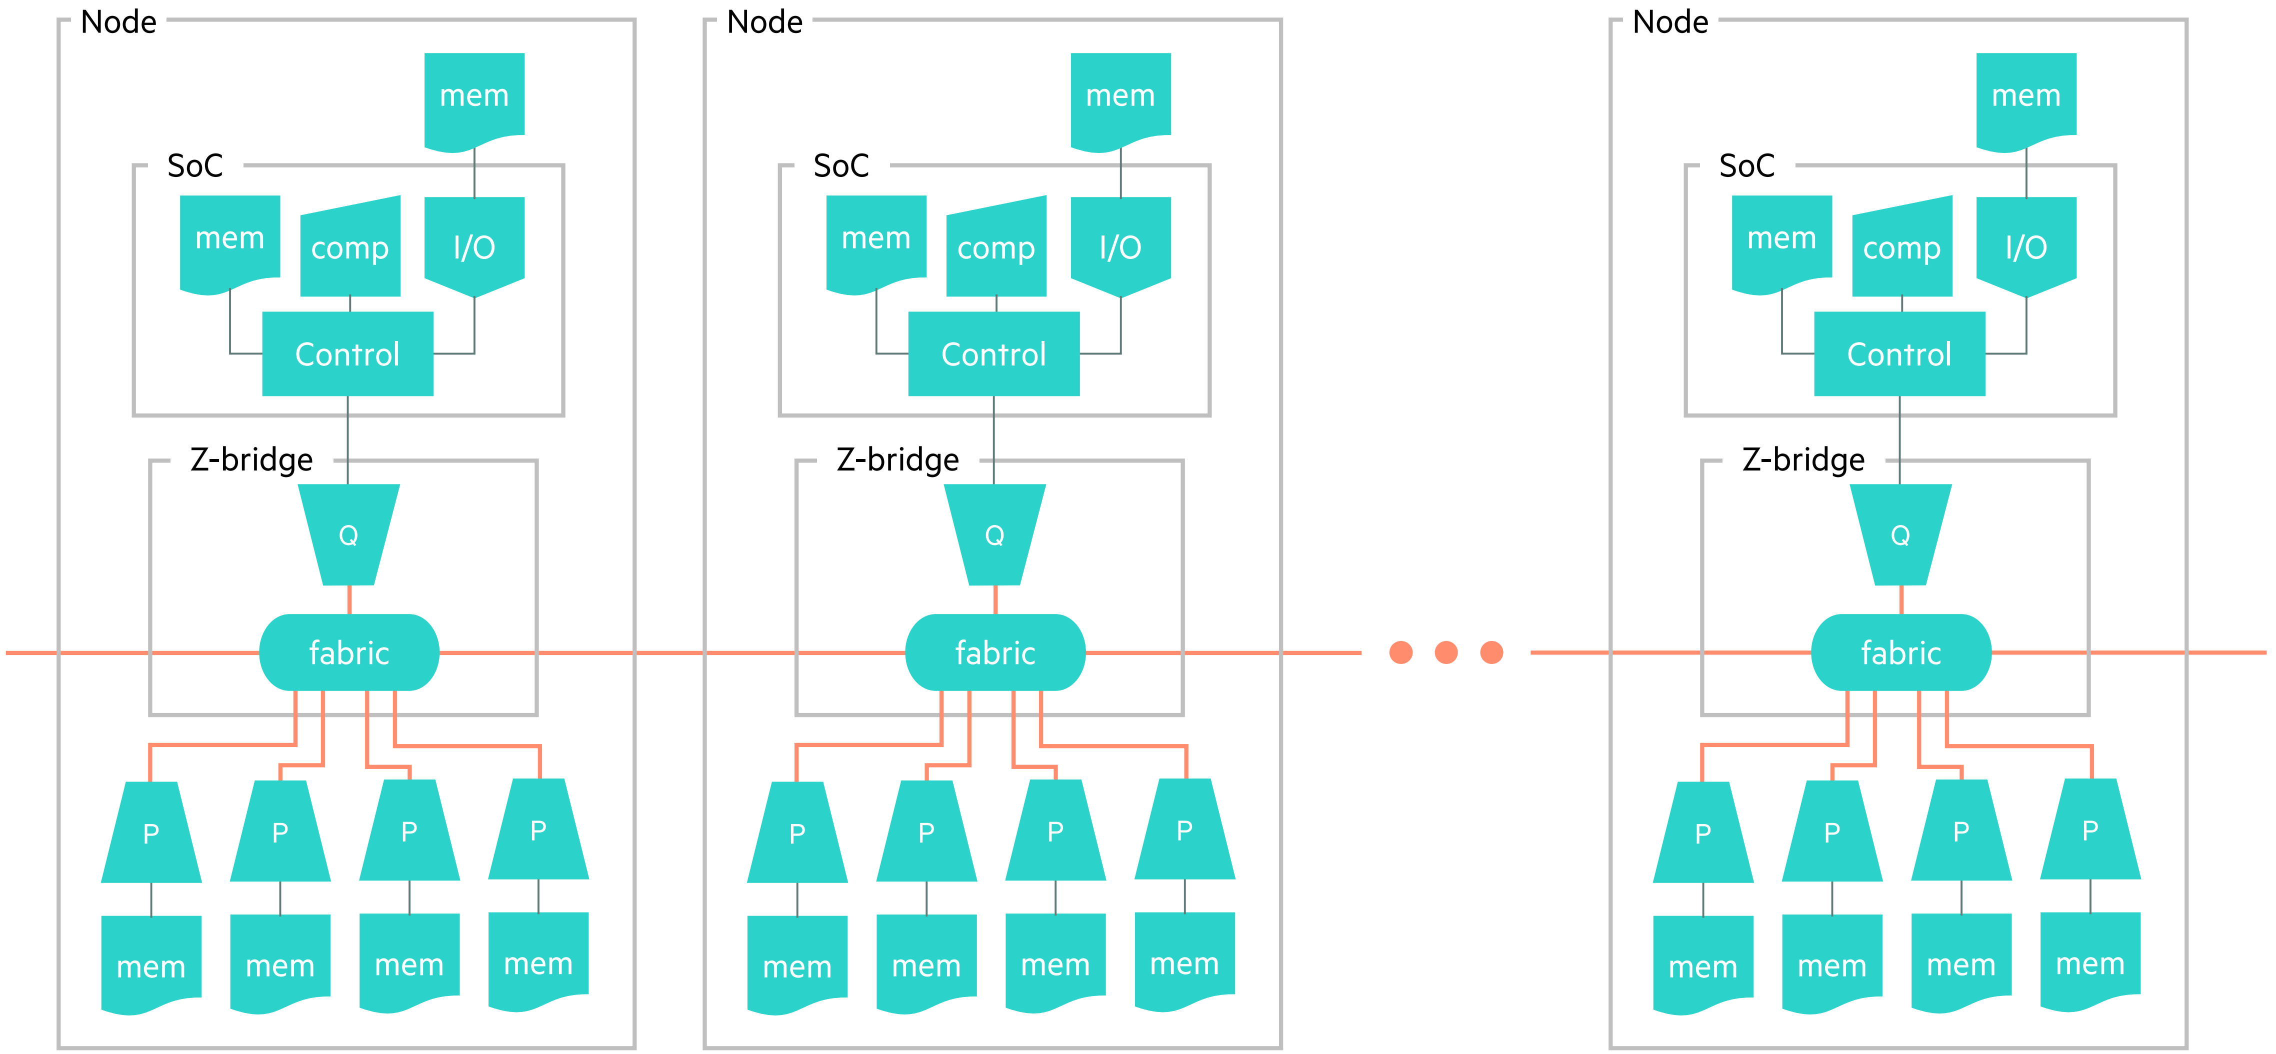

Mature Gen-Z devices

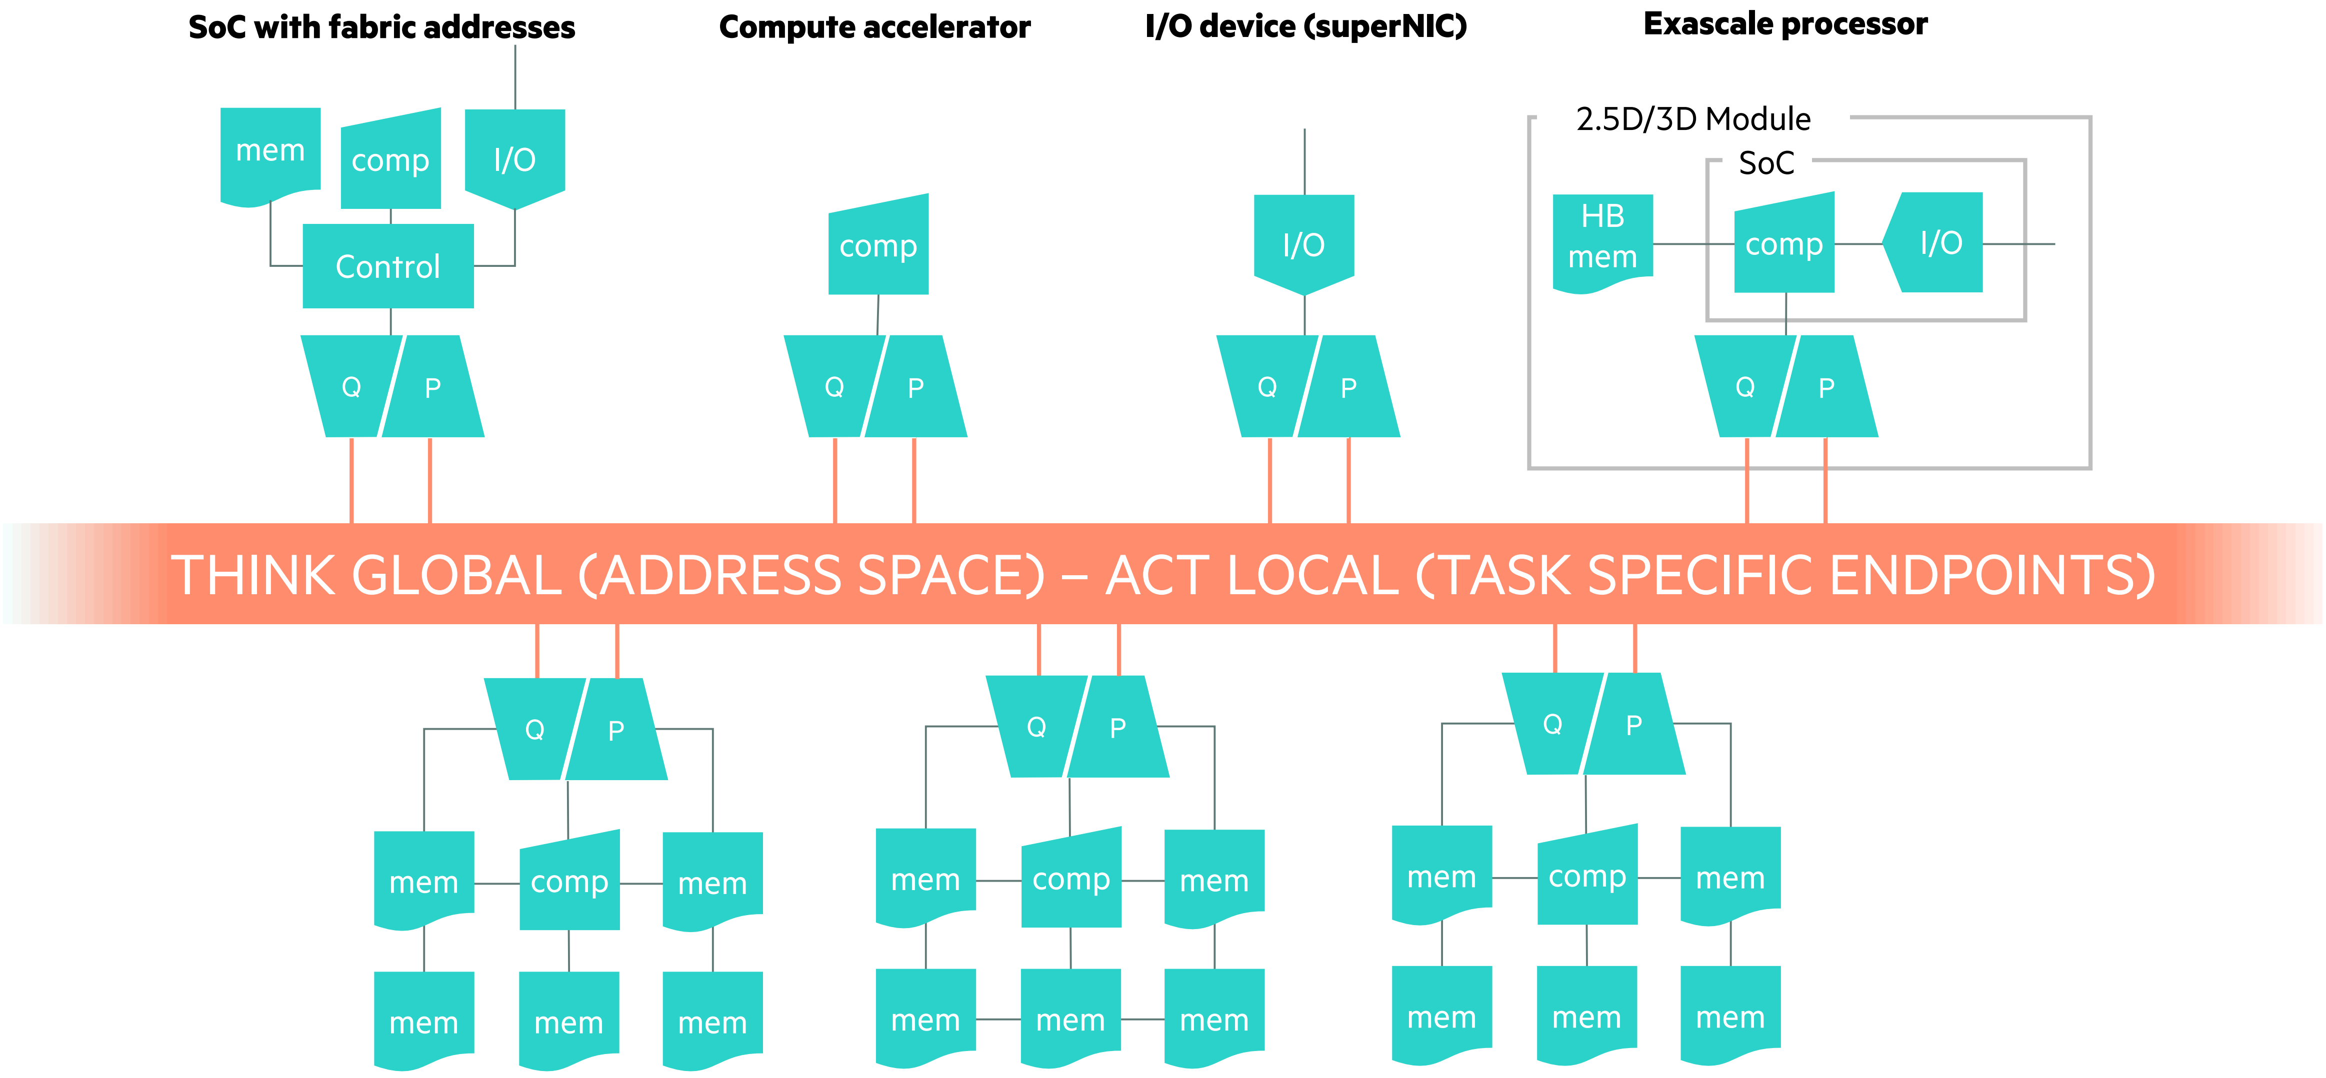

Memory node
with memory-side accelerator

What are core Memory-Driven Computing components?

Fast, persistent memory	Fast memory fabric	Task-specific processing	New and Adapted software
Combining memory and storage in a stable environment to increase processing speed and improve energy efficiency	Using photonics where necessary to eliminate cost of distance and create otherwise impossible topologies	Optimizing processing from general to specific tasks and embrace novel computational techniques	Radically simplifying programming and enabling new applications that we can't even begin to build today

How does Memory-Driven Computing enable new applications?

One architecture scales from the dense data center to the intelligent edge

Memory Abundance

- Similarity search
- Search space optimization
- Financials futures modeling

Non-volatility of Memory

- Scalable transactional key value stores
- Managed data structures
- Energy scalability and retention

Memory shared with just the right compute

- Spark in-memory Hadoop
- Deep neural net training
- Network function virtualization

Dynamic Range

- Memory-Driven Computing edge
- Consistent node, enclosure, rack, row, data center

Hewlett Packard Enterprise

Transform performance with Memory-Driven programming

Modify existing frameworks

In-memory analytics

15x
faster

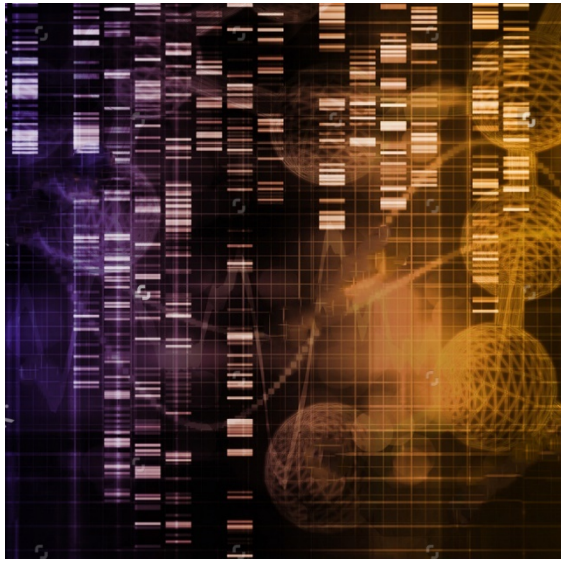

Similarity search

40x
faster

New algorithms

Large-scale graph inference

100x
faster

Completely rethink

Financial models

10,000x
faster

Hewlett Packard Enterprise

Open, open, open

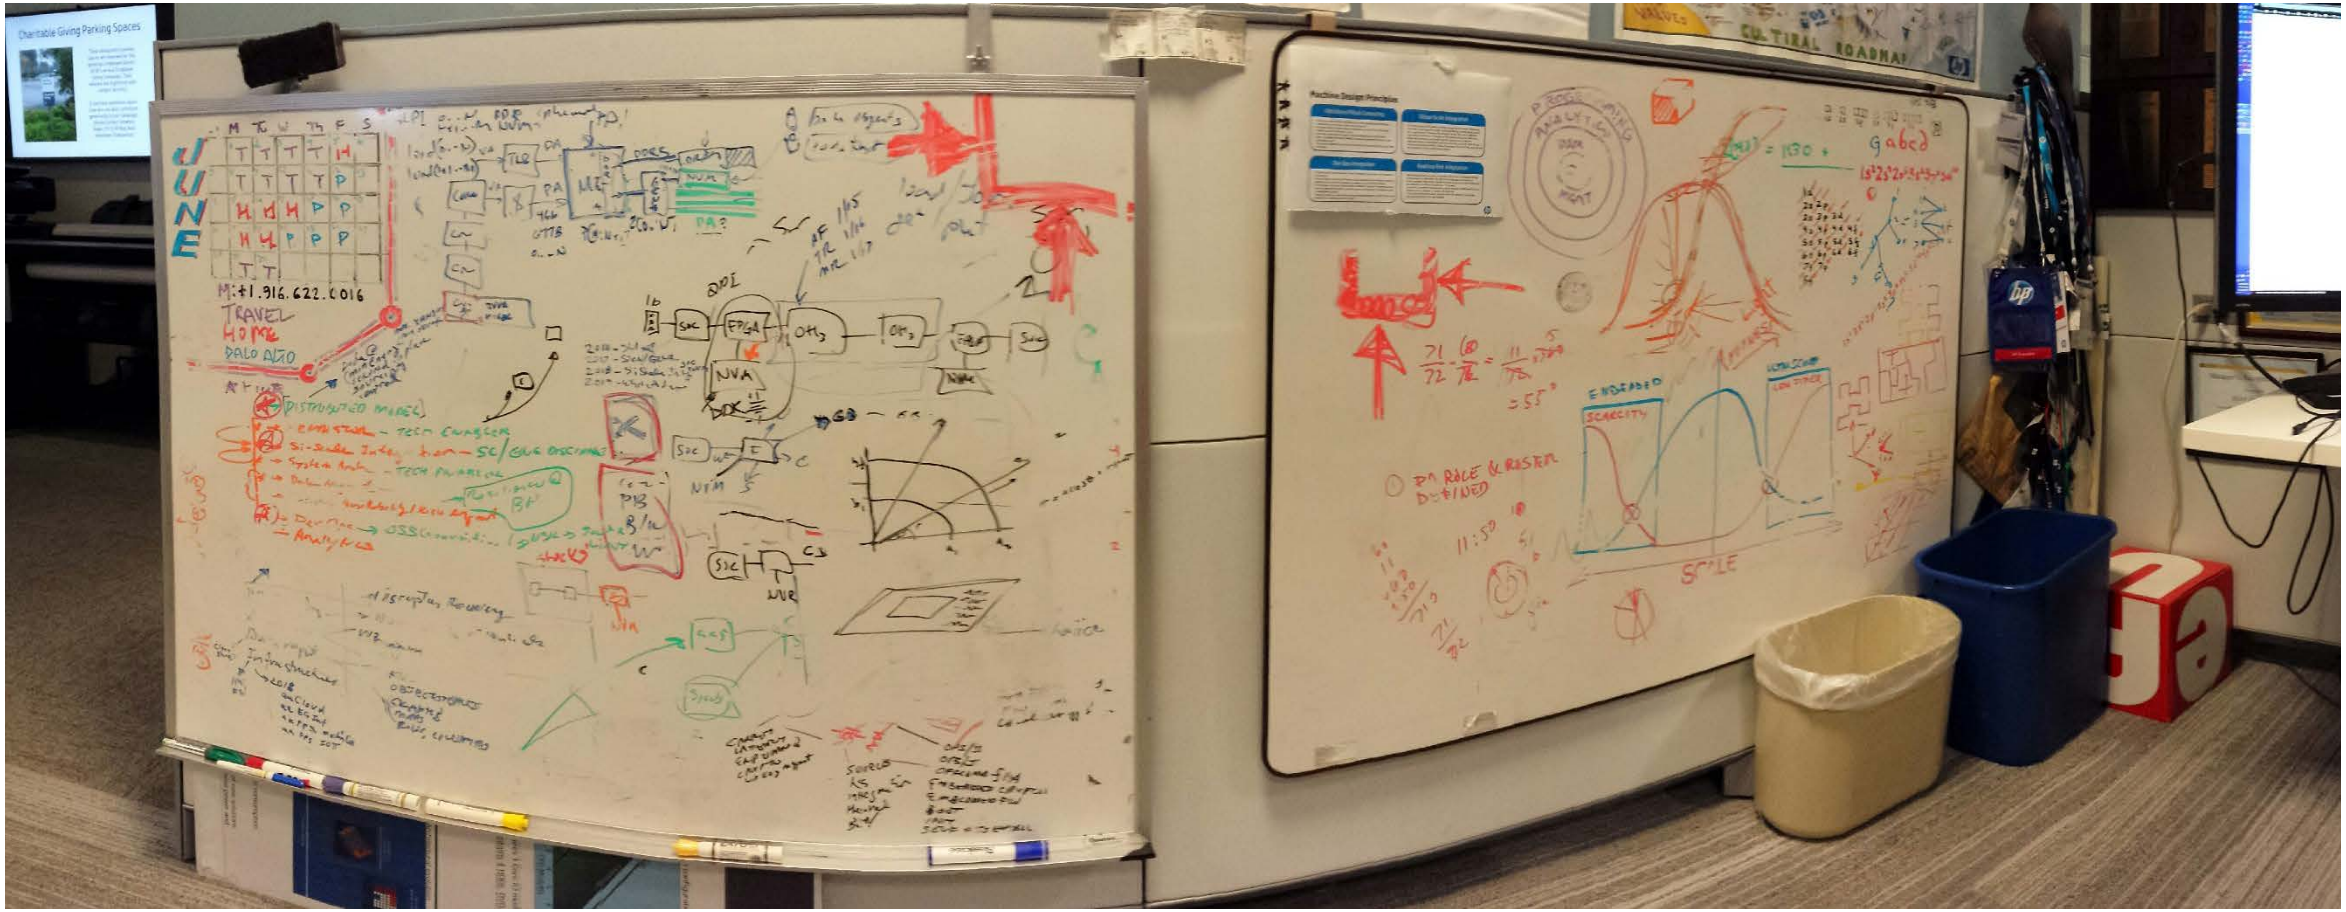

The state of the MFT prototype when we announced at HP Discover, June 2016

Gen-Z: new open interconnect protocol

Key enabler of the Memory-Driven Computing open architecture

High Bandwidth
Low Latency

Advanced workloads & technologies
Scalable from IoT to exascale

Compatible
Economical

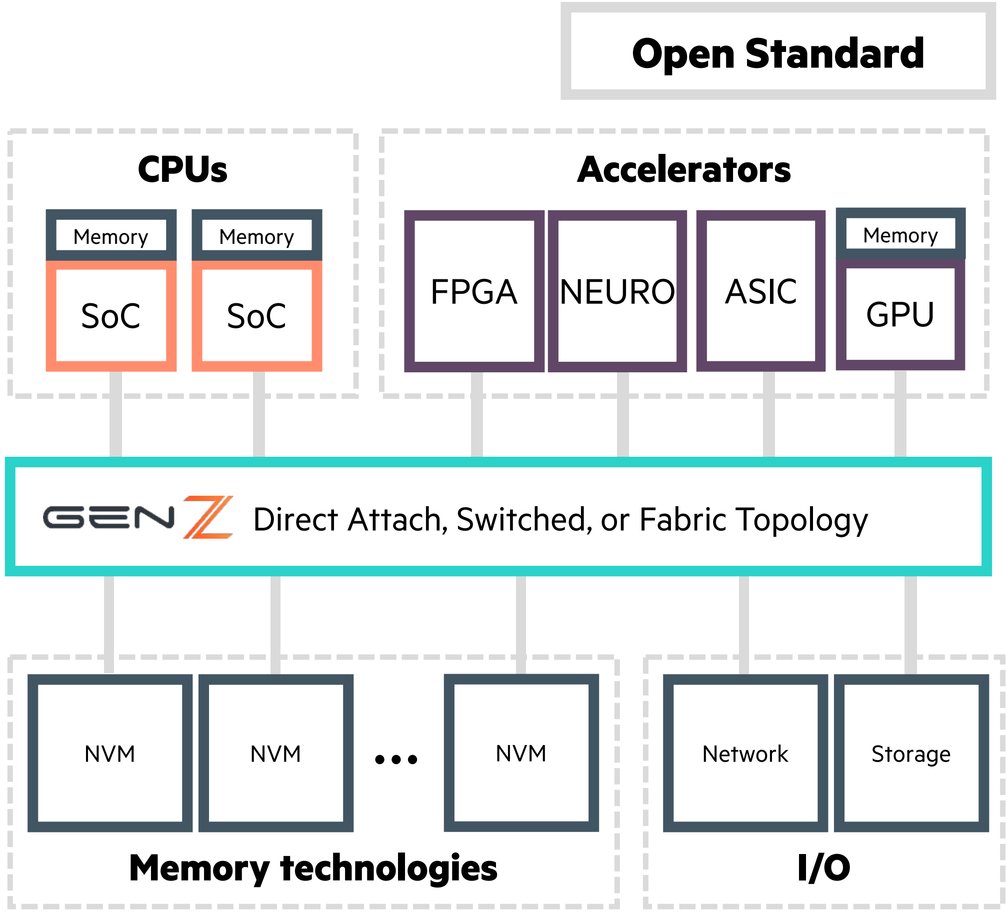

Industry collaboration on interconnect technology

GEN Z

Industry leaders developing a next generation, memory-semantic interconnect

www.genzconsortium.org

Current Membership

GEN Z

- Alpha Data
- AMD
- Amphenol Corporation
- ARM
- Broadcom
- Cavium Inc.
- Cray
- Dell EMC
- Everspin
- FoxxConn Interconnect
- HPE
- Huawei
- IBM
- IDT
- Intelliprop
- Jabil
- Lenovo
- Lotes
- Luxshare ICT
- Mellanox Tech, Ltd
- Mentor Graphics
- Micron
- Microsemi
- Molex
- NetApp
- Nokia
- Numascale
- PLDA Group
- Red Hat
- Samsung
- Seagate
- Smart Modular
- SK Hynix
- Spin Transfer Tech
- TE
- VMware
- Western Digital
- Xilinx
- YADRO

39 Members (and growing)
12 Board Companies

Memory-Driven Computing Developer Toolkit

Software already available to you

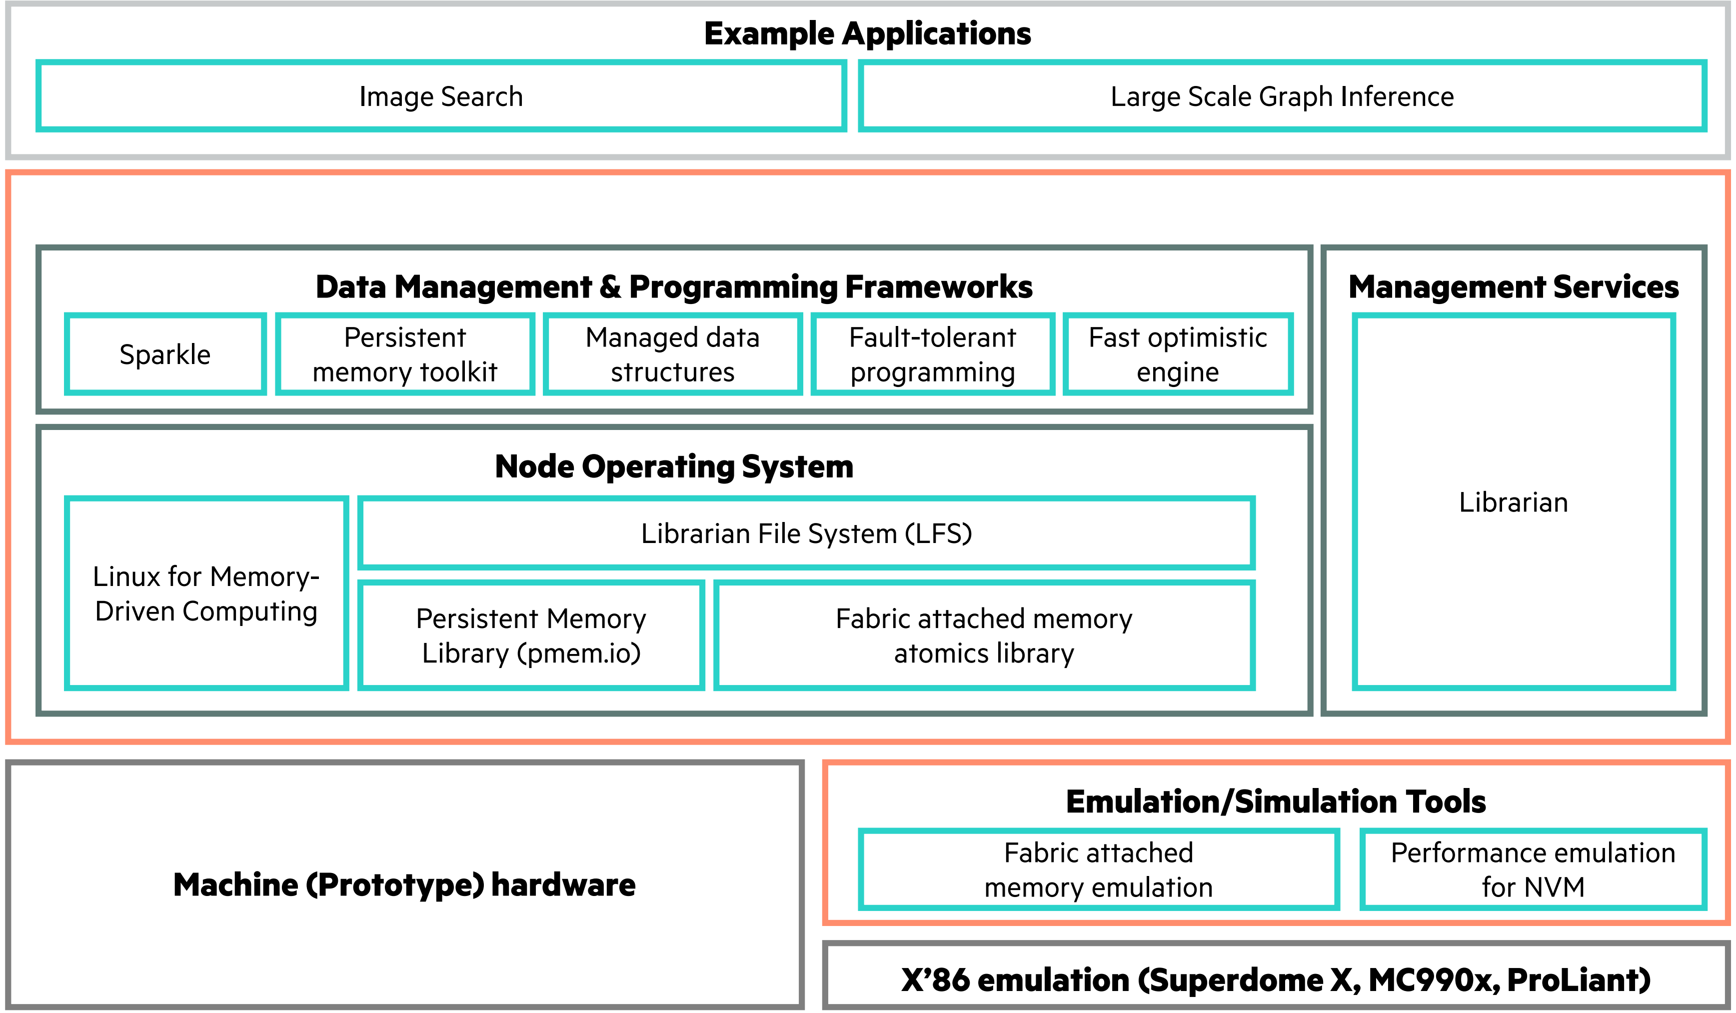

- Example Applications
- Programming and analytics tools
- Operating system support
- Emulation/simulation tools

Get access to the toolkit:

https://www.labs.hpe.com/the-machine/developer-toolkit

Open source components

What can you do in 300ns?

It's more than just bits per (meter • second • Watt • dollar) between transmit and receive

Budget for the endpoint: $300ns - 150ns = 150ns$

Budget for the hops: $150ns - 3 \times 25ns = 75ns$

Give the balance to fiber: $75ns \times \frac{1m}{5ns} \times \frac{1}{2} = 7.5\text{m}$

Describe the sphere: $\frac{4}{3}\pi(7.5m)^3 = 994m^3$

Convert to racks: $994m^3 \times \frac{1\,rack}{8m^3} = 124\,racks$

124 MFT racks = 19PB, 159K cores