Hewlett Packard Enterprise

Hewlett Packard Labs

THE MACHINE & MEMORY-DRIVEN COMPUTING

711. WE-Heraeus-Seminar, Bad Honnef 15.01.2020

Martin Brennecke

WE TALK ABOUT…

- Power consumption – A problem statement
- The new normal: we are not keeping up…
- Re-thinking computing fundamentals
- The Machine
- Where are we today?
- Wrapping Up

PROJECTED POWER CONSUMPTION – A PROBLEM STATEMENT

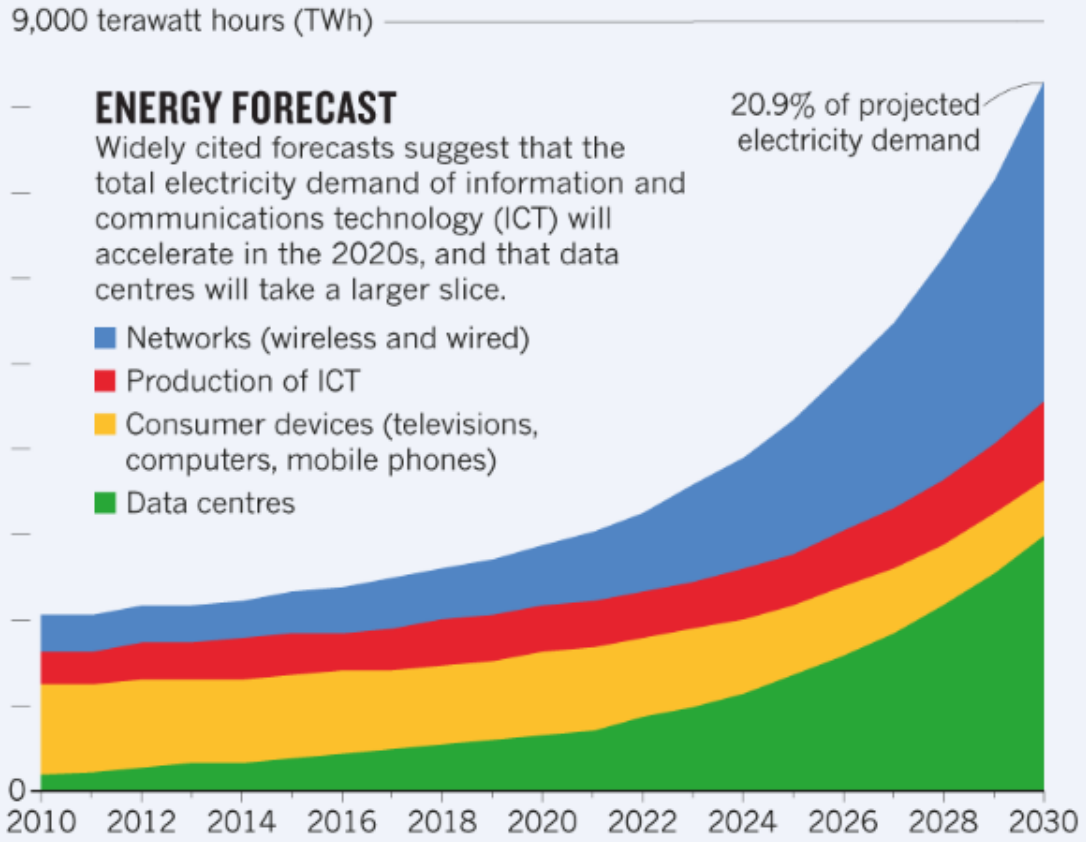

The chart above is an 'expected case' projection from Anders Andrae, a specialist in sustainable ICT. In his 'best case' scenario, ICT grows to only 8% of total electricity demand by 2030, rather than to 21%.

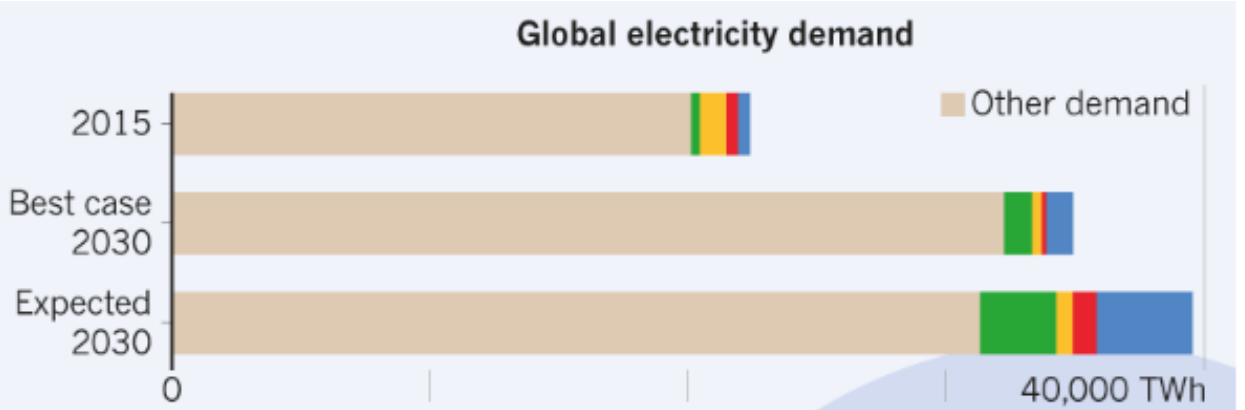

INTERNET EXPLOSION

Internet traffic* is growing exponentially, and reached more than a zettabyte (ZB, 1×10^{21} bytes) in 2017.

*Traffic to and from data centres.
†TB, terabyte (10^{12} bytes); PB, petabyte (10^{15} bytes); EB, exabyte (10^{18} bytes).

©nature

The Information Factories: https://www.nature.com/articles/d41586-018-06610-y

PROJECTED POWER CONSUMPTION – A PROBLEM STATEMENT

Expected Case

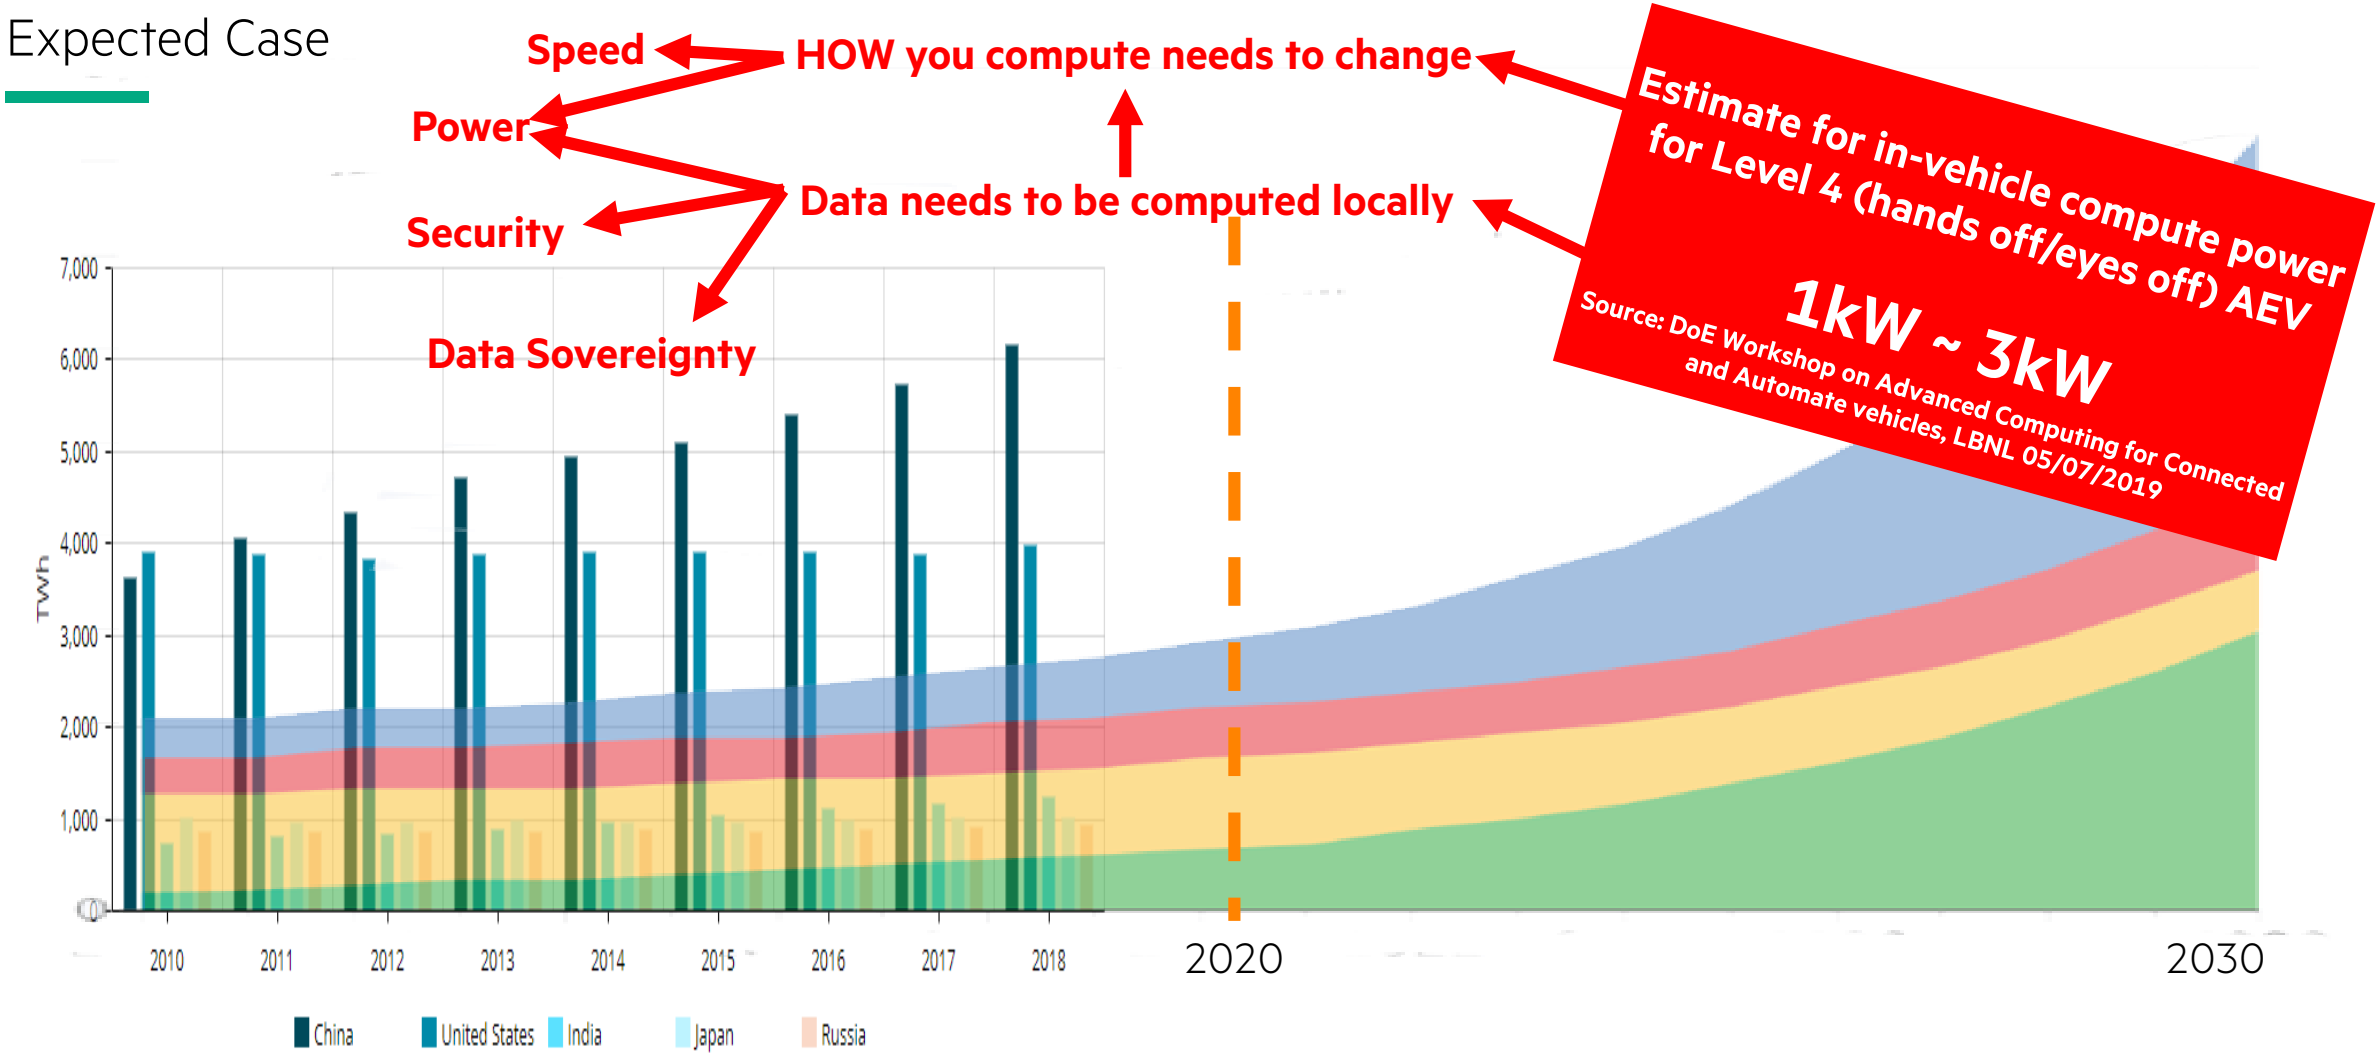

Global Energy Statistical Yearbook 2019

The Information Factories: https://www.nature.com/articles/d41586-018-06610-y

THE NEW NORMAL: WE ARE NOT KEEPING UP

Data quantity

Data locality

Remember yesterday's session from Michiel van Haarlem

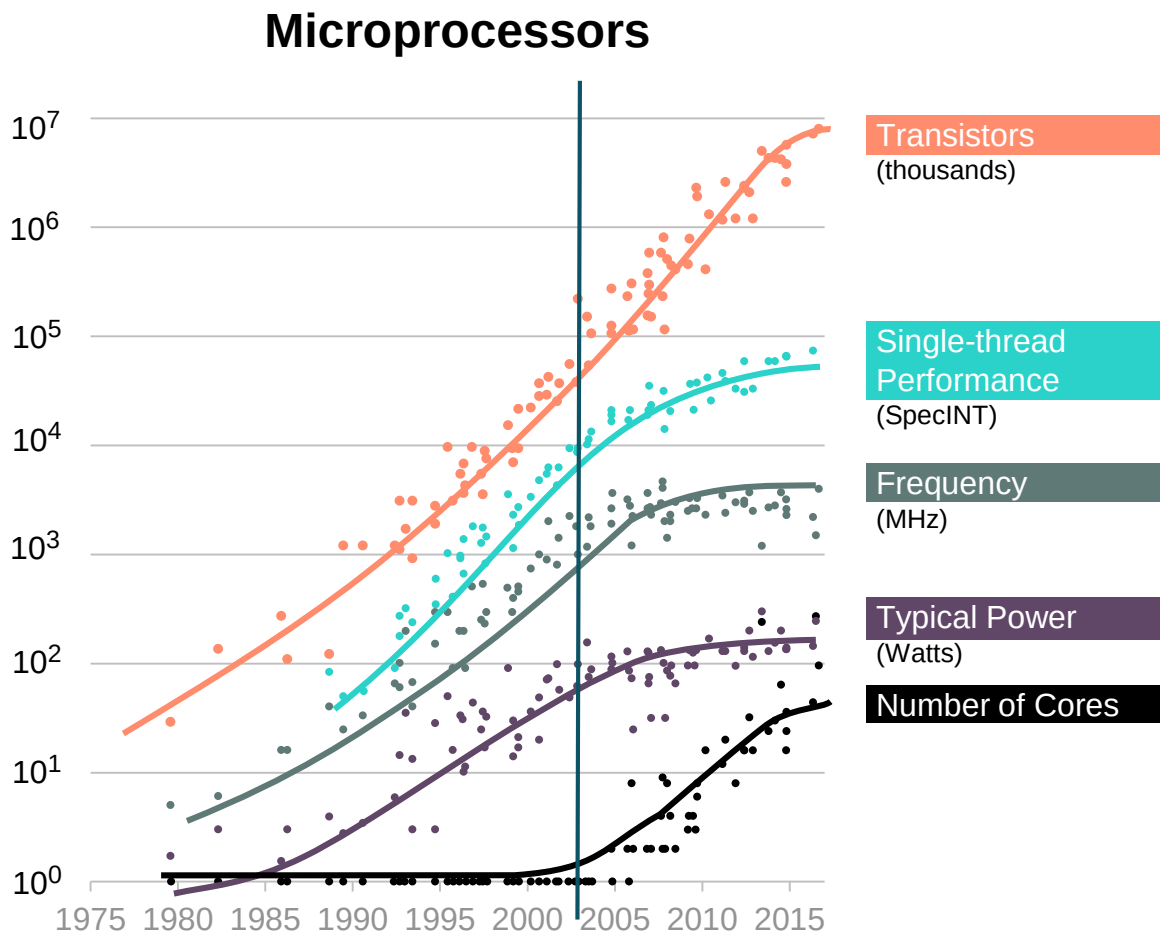

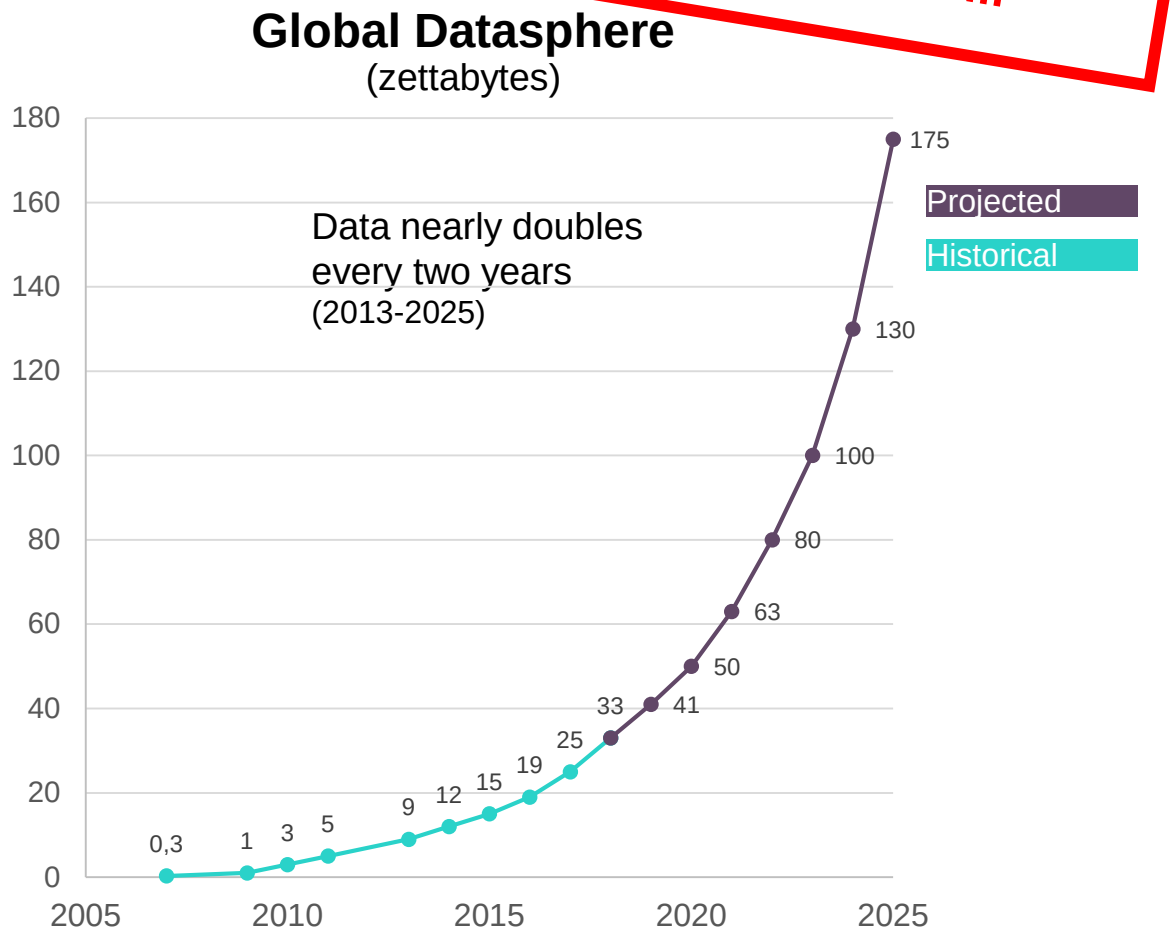

Source: K. Rupp. 40 Years of Microprocessor Trend Data
https://www.karlrupp.net/2015/06/40-years-of-microprocessor-trend-data/

Source: IDC Data Age 2025 study, sponsored by Seagate, Nov 2018
https://www.seagate.com/files/www-content/our-story/trends/files/idc-seagate-dataage-whitepaper.pdf

WHY CAN'T COMPUTE KEEP UP WITH THIS DATA EXPLOSION?

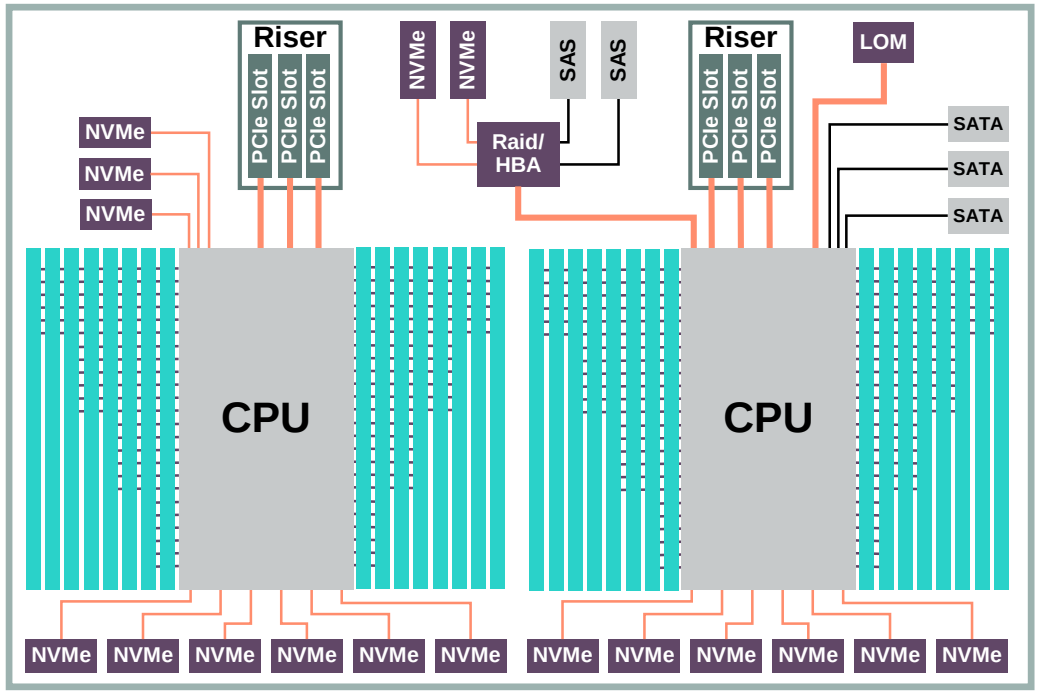

CPU core counts growing more rapidly than available processor memory and I/O bandwidth

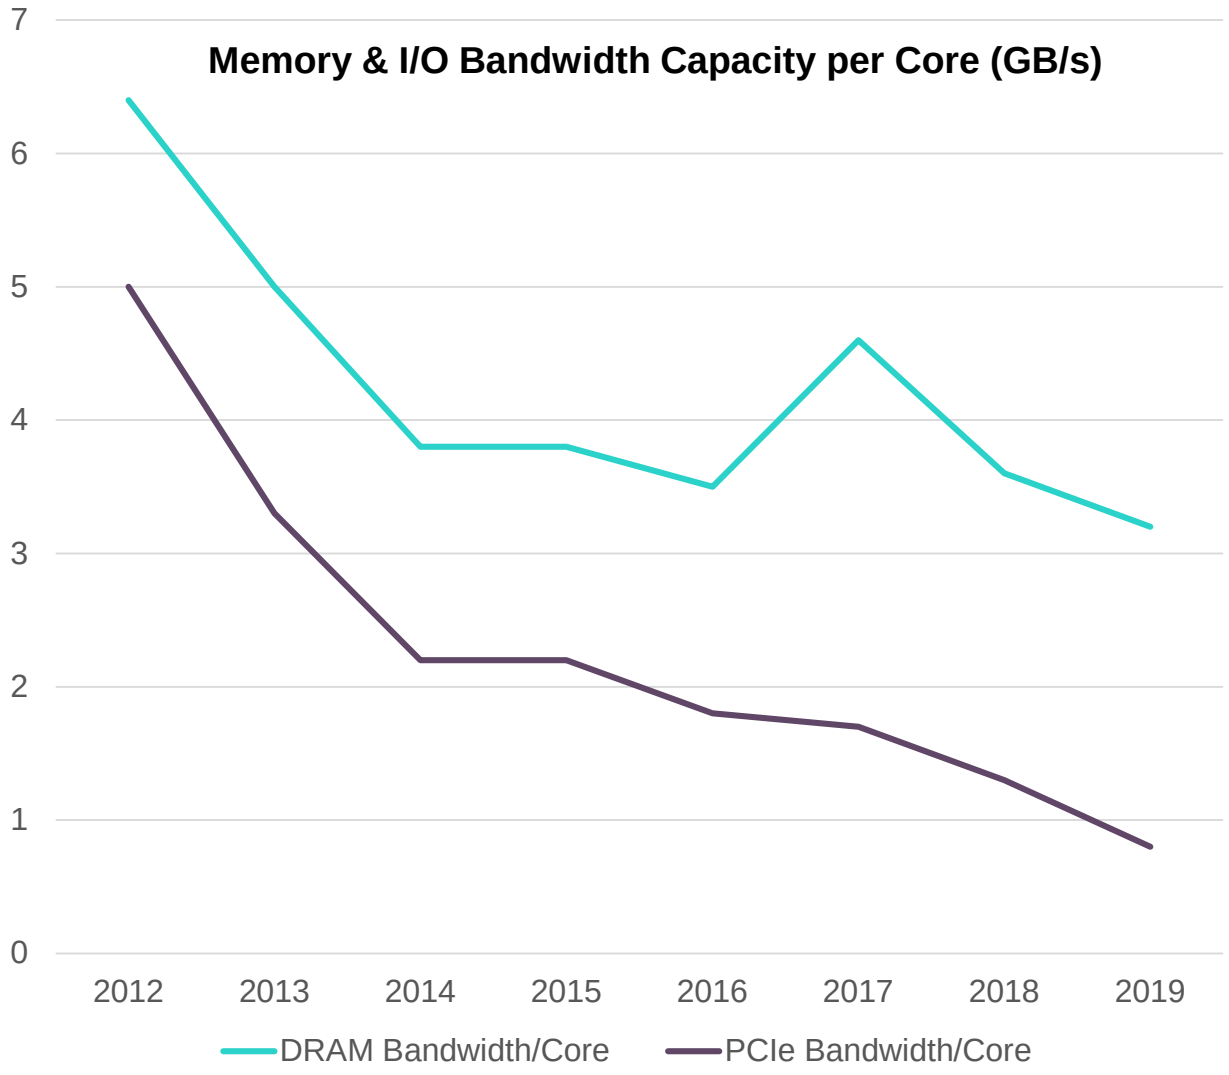

WHY CAN'T COMPUTE KEEP UP WITH THIS DATA EXPLOSION?

Today's architecture is constrained by the CPU

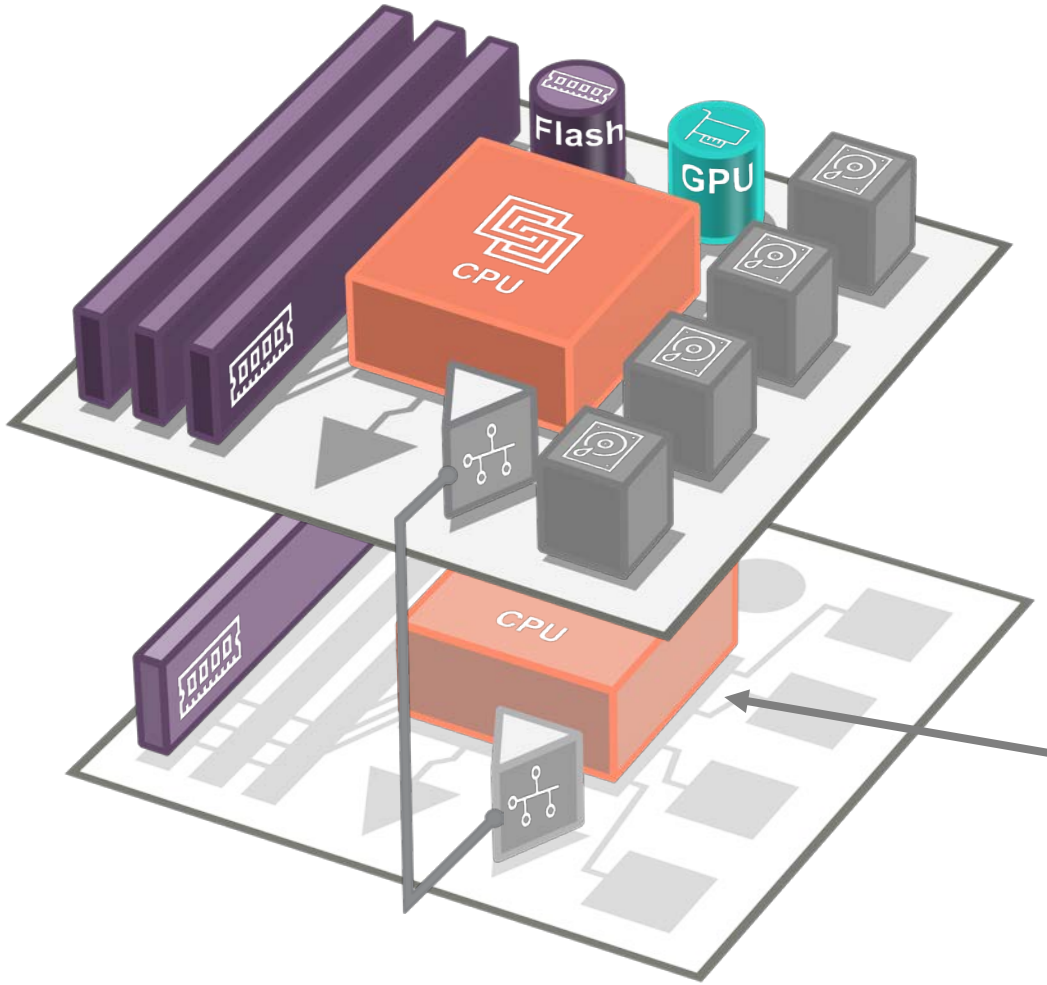

If you exceed the limit what can be connected to one CPU, you need another CPU (or server)

(VERY) SIMPLIFIED VIEW ON COMPUTERS SINCE THE 60‘S

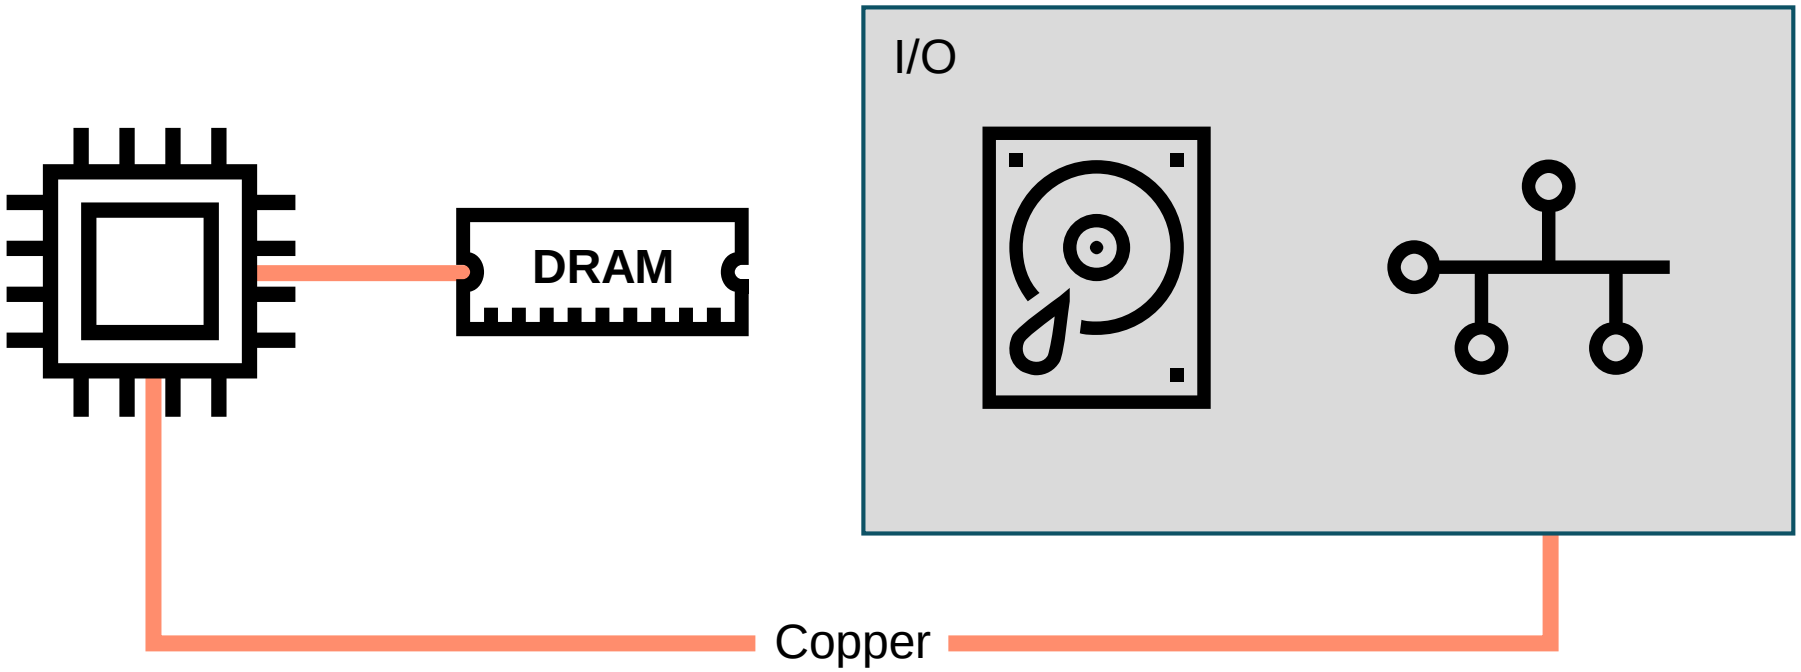

From processor-centric computing…

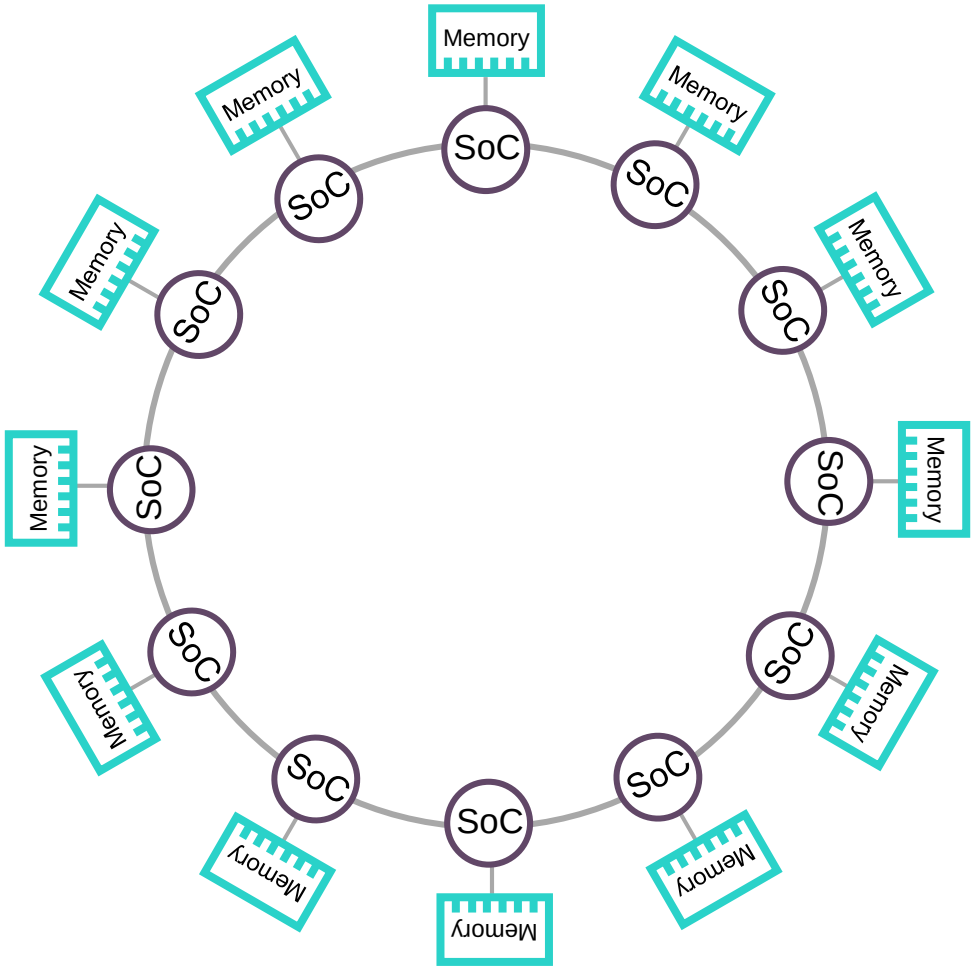

…to Memory-Driven Computing

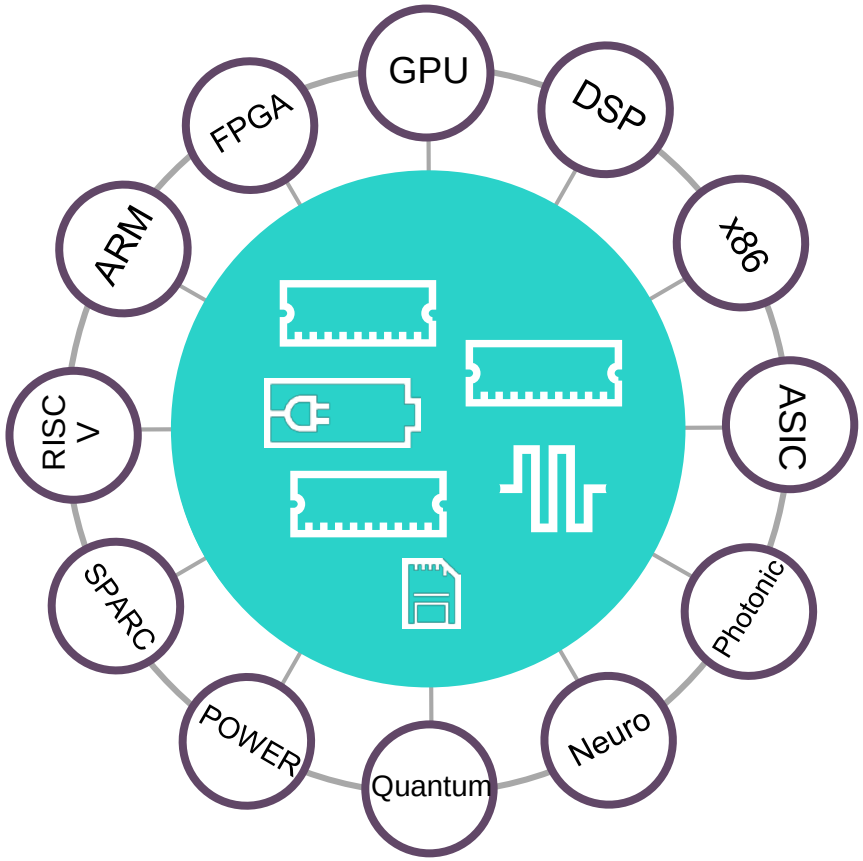

MEMORY-DRIVEN COMPUTING – SCALE TO ANY SIZE & WORKLOAD

Edge device

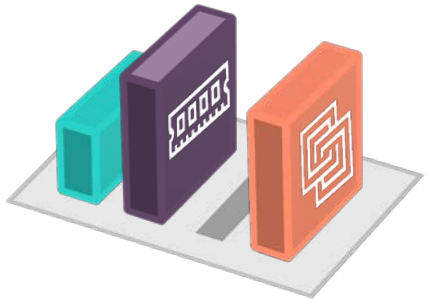

- Near-zero power
- Persistent memory
- AI task-specific accelerator

Cloud infrastructure

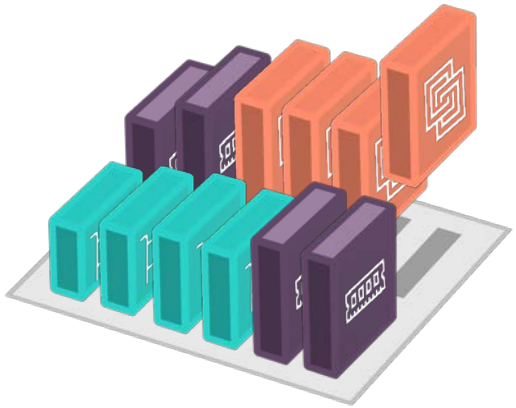

- Composable infrastructure from every edge to any cloud
- Microservices in microseconds at massive scale

Big memory machine

- High-performance data analytics
- Large shared memory
- Monte Carlo, graph analytics applications, etc.

Exascale computer

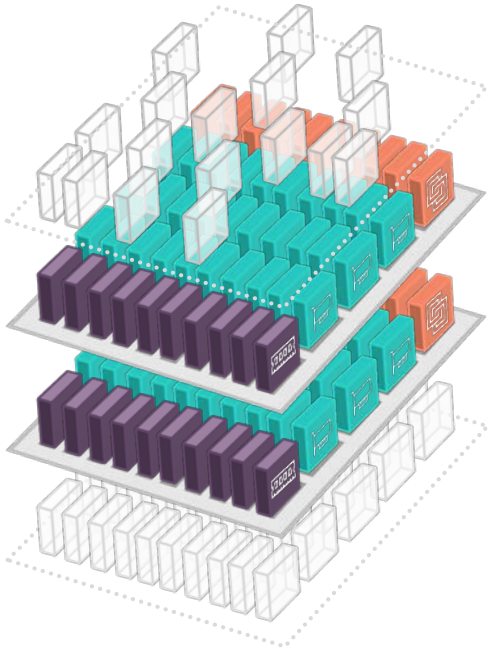

- 100,000+ components
- Ultra-fast message passing and checkpointing
- 20x more energy-efficient than state-of-the-art

CORE COMPONENTS OF A MEMORY-DRIVEN ARCHITECTURE

Fast, persistent memory

Combining memory and storage in a stable environment to increase processing speed and improve energy efficiency

Fast memory fabric

Using photonics where necessary to eliminate distance and create otherwise impossible topologies

Task-specific processing

Optimizing processing from general to specific tasks

New and adapted software

Radically simplifying programming and enabling new applications that we can't even begin to build today

FAST MEMORY-SEMANTIC FABRIC USING PHOTONICS

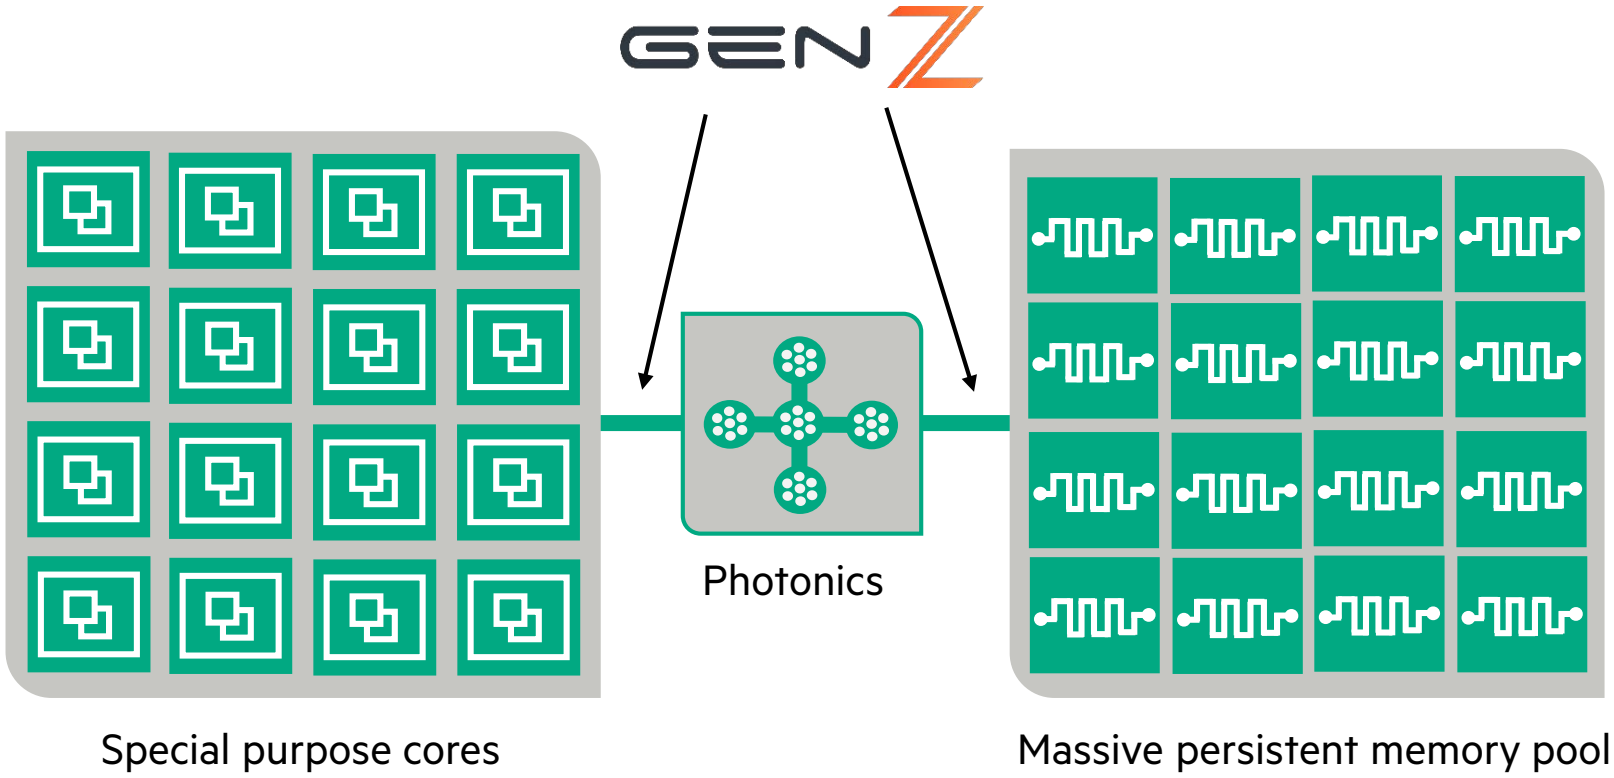

GEN-Z: NEW OPEN INTERCONNECT PROTOCOL

Key enabler of the Memory-Driven Computing open architecture

- High bandwidth & Low latency
- Supports electrical or optical interconnects
- Advanced workloads & technologies
- Scalable from IoT to Exascale
- Economical & Open standard
- Security built-in at the hardware level

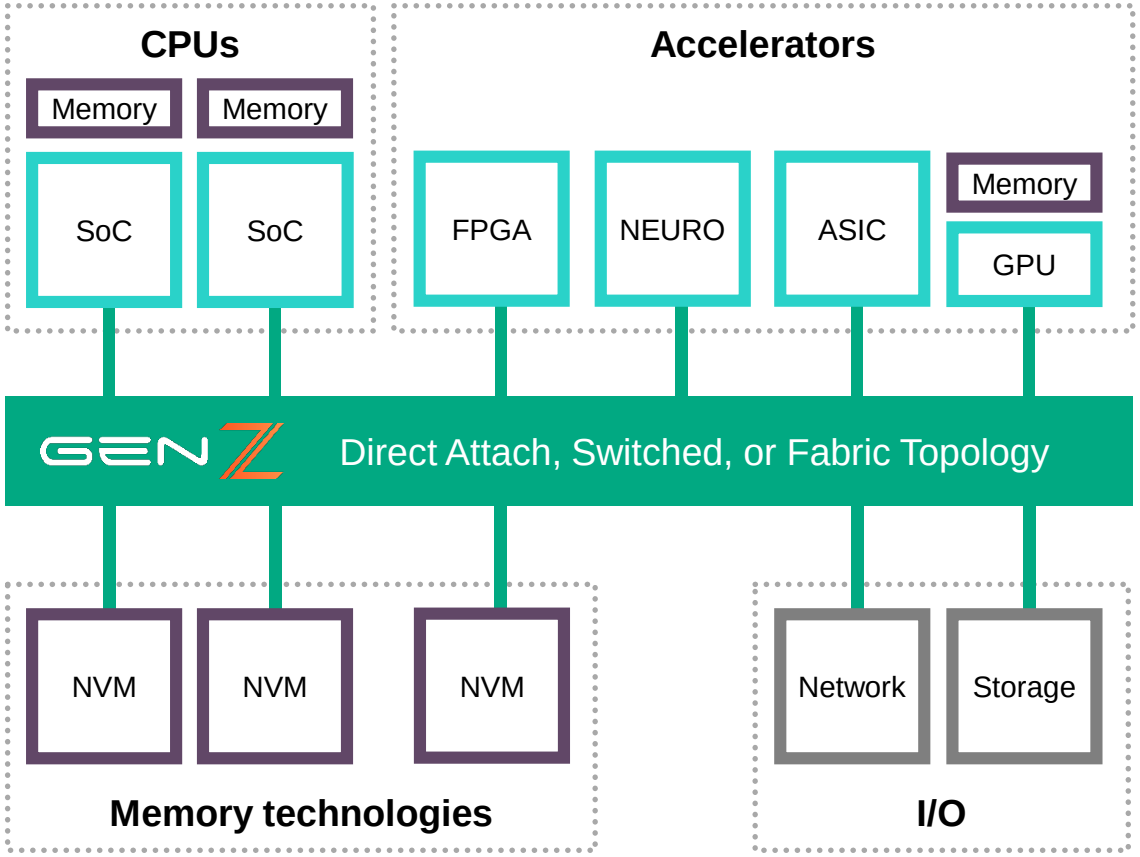

SCALING UP USING PHOTONIC FABRICS

Photonics

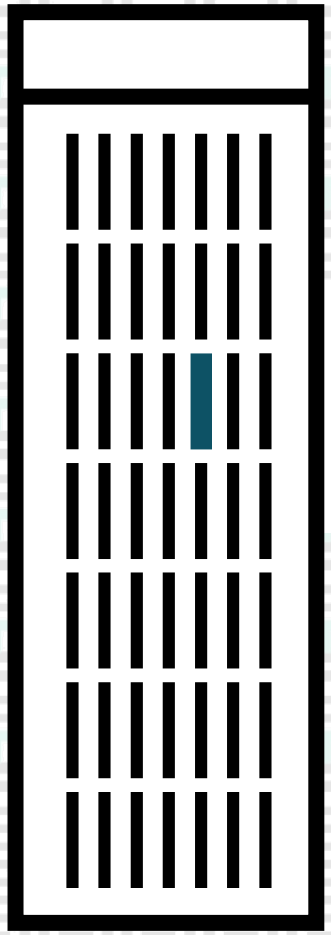

Hundreds of racks can behave as a **single** server

2^{92} bytes = 4096 yottabytes (=$4096*10^{24}$ Bytes)

OPEN CONSORTIUM WITH BROAD INDUSTRY SUPPORT

GEN Z Consortium Members

System OEM	CPU/Accel	Mem/Storage	Silicon	IP	Connect	Software
Cisco	AMD	Everspin	Broadcom	Avery	Aces	Redhat
Cray	Arm	Micron	IDT	Cadence	AMP	VMware
Dell EMC	IBM	Samsung	Marvell	Intelliprop	FIT	
H3C	Qualcomm	Seagate	Mellanox	Mentor	Genesis	**Govt/Univ**
Hitachi	Xilinx	SK Hynix	Microsemi	Mobiveil	Jess Link	ETRI
HP		Smart Modular	Sony Semi	PLDA	Lotes	Oak Ridge
HPE		Spintransfer		Synopsys	Luxshare	Simula
Huawei						
IBM		Toshiba			Molex	UNH
Lenovo		WD			Samtec	Yonsei U
NetApp					Senko	ITT Madras
Nokia		**Technology SP**		**Eco/Test**	TE	
Yadro		Google		Allion Labs	3M	
		Microsoft		Keysight		
		Node Haven		Teledyne LeCroy		

WHERE ARE WE TODAY?

HPE INTRODUCED RUNNING PROTOTYPE OF THE MACHINE

The world's largest single-memory computer, launched 2018

- 160 TB of shared memory spread across 40 interconnected physical nodes.
- An optimized Linux-based operating system running on Cavium ThunderX2 ARMv8-A System on a Chip.
- Photonics/Optical communication links, including the new X1 photonics modul
- Software programming designed to take advantage of abundant of persistent memory.

the MACHINE

MDC SANDBOX ON SUPERDOME FLEX – 64 SOCKETS, 48TB MEMORY

MEMORY DRIVEN COMPUTING SOFTWARE DEVELOPMENT

Frictionless migration to future platforms – start today!

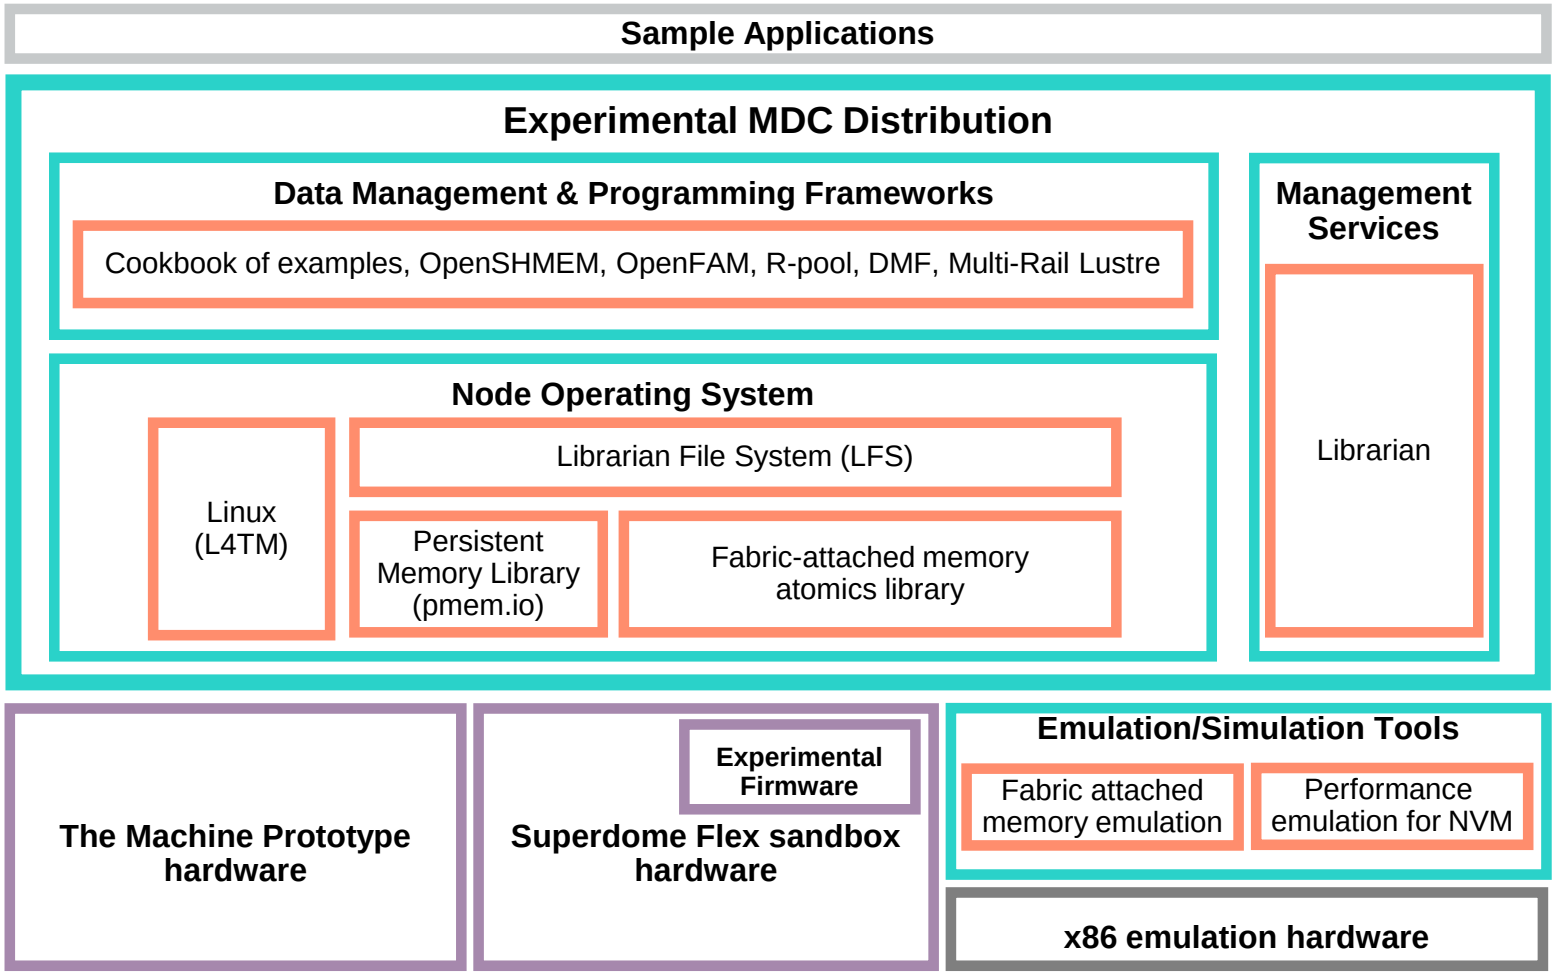

Start re-think your algorithms & apps:

- Eliminate data movement
- Change Mindset how to create code
- Think of having ubiquitous memory
- Think of access to memory having no interference from other compute elements

Open sourced components

https://github.com/HewlettPackard/mdc-toolkit
https://www.labs.hpe.com/the-machine/the-machine-distribution

TRANSFORM PERFORMANCE WITH MEMORY-DRIVEN PROGRAMMING

Modify existing frameworks

New algorithms

Completely rethink

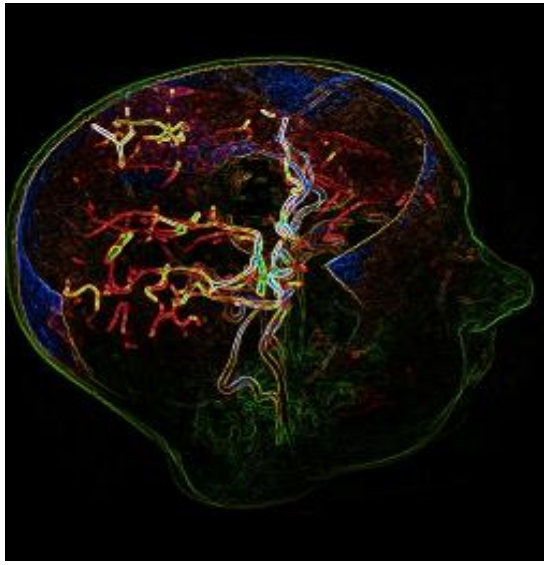

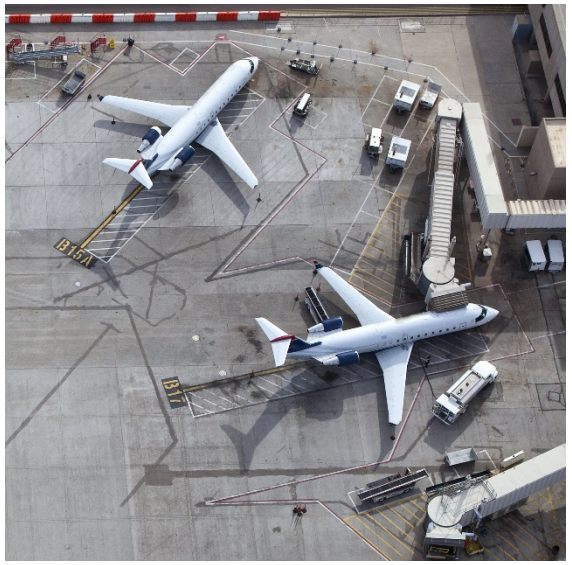

In-memory analytics	Genome comparison	Large-scale graph inference	Simulation e. g. Monte Carlo
15x faster	**100x** faster	**100x** faster	**10,000x** faster

COMPLEX MODELS CONVERGE IN MINUTES NOT DAYS

Store intermediate results of simulation models, simply look up results instead of recalculating

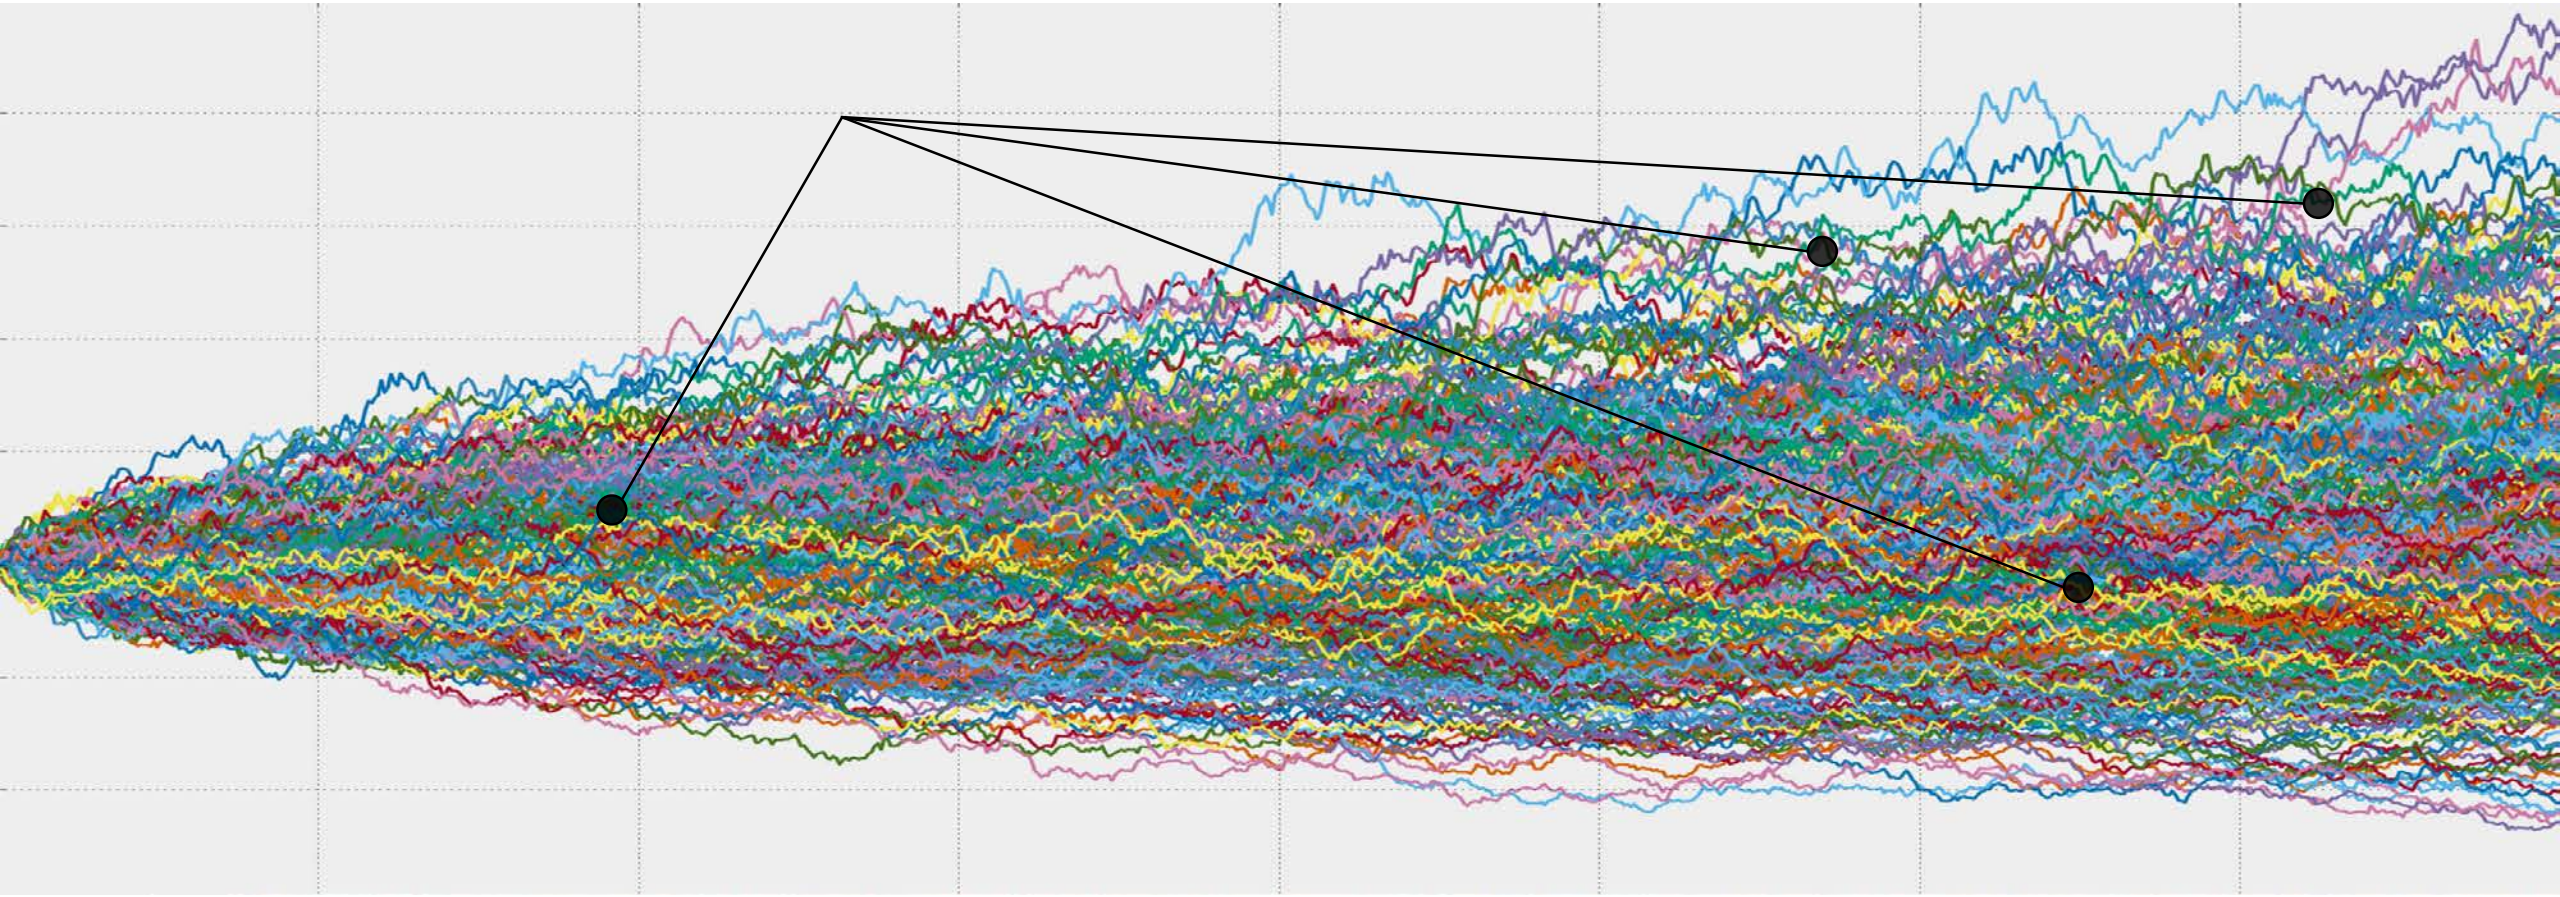

MEMORY-DRIVEN COMPUTING HELPS DZNE OUTPACING THE GLOBAL TIME BOMB OF NEURODEGENERATIVE DISEASE

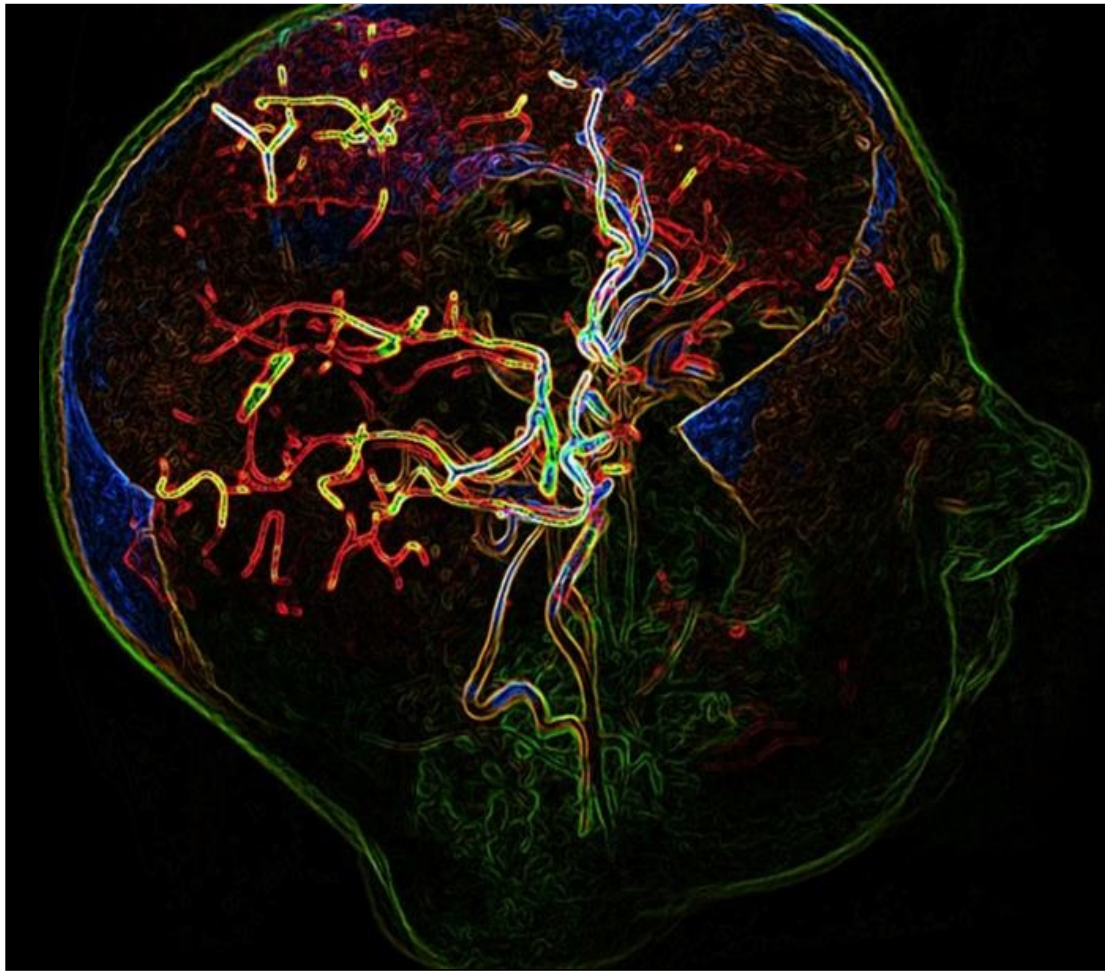

DZNE
Deutsches Zentrum für
Neurodegenerative Erkrankungen
in der Helmholtz-Gemeinschaft

DZNE discovered HPE's Memory-Driven Computing — and saw unprecedented computational speed improvements that hold new promise in the race against Alzheimer's

60% power reduction cuts research costs

100x increase in analytics speed blasts research bottlenecks, leading to shorter processing time — from 22 minutes to **13** seconds

WRAPPING UP

THE JOURNEY TO MDC

- MDC & Gen-Z provide fast load/store access to large shared pool of fabric-attached non-volatile memory
- De-coupling system components by using a memory-semantic open fabric and photonics
- Many opportunities for implementation of task-specific processors or any future technologies
- Many opportunities for software innovation – Re-think how to create code
  - Think of having ubiquitous memory
  - Eliminate data movement
  - Operating system
  - Data stores
  - Analytics platforms
  - Applications & Algorithms
- Non-MDC and MDC architectures will co-exist for the foreseeable future
- MDC technology will be phased in over the next generation of systems
- **Start now** with experimenting with your applications and algorithms!

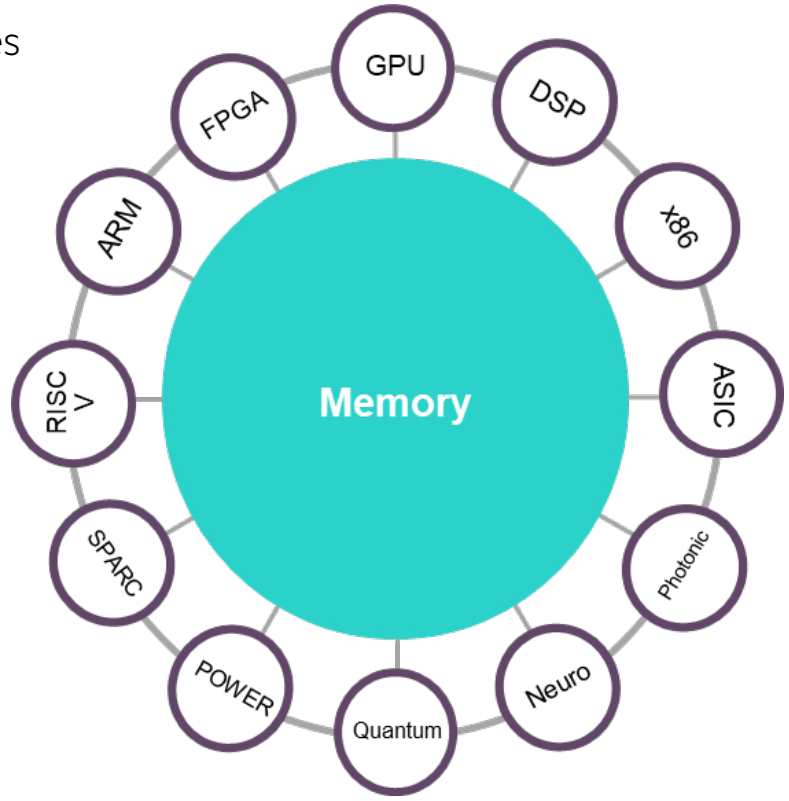

LINKS & SOURCES

- Entry point: https://www.labs.hpe.com/memory-driven-computing
- Memory-Driven Computing toolkit: https://github.com/HewlettPackard/mdc-toolkit
- Linux for The Machine: https://github.com/FabricAttachedMemory
- Fabric Attached Memory Emulation: https://github.com/FabricAttachedMemory/Emulation

More detailed stuff here:

- Spark for The Machine: https://github.com/HewlettPackard/sparkle
- Fast optimistic engine for data unification services: https://github.com/HewlettPackard/foedus
- Fault-tolerant programming model for non-volatile memory:
  - Atlas: https://github.com/HewlettPackard/Atlas
  - NVthreads: https://github.com/HewlettPackard/nvthreads
- Managed Data Structures: https://github.com/HewlettPackard/mds
- Performance emulation for NVM latency and bandwidth: https://github.com/HewlettPackard/Quartz

Hewlett Packard Enterprise

Hewlett Packard Labs

MEMORY DRIVEN

martin.brennecke@hpe.com
@mbrennecke
+49 170 92 444 69