

Data, the new frontier.

Rob Mollard
March 2019


Hewlett Packard
Enterprise



Courtesy Nasa.gov




knowledge
is power!



Convergence of HPC, Big Data, AI & IoT for an intelligent future

Prepare for a world where everything computes





Technology will be
embedded
everywhere

Everyone and
everything
will be connected

Real-time will drive
computing to the
edge

Extreme “Computing at the Edge”

HPE Spaceborne Supercomputer



Spaceborne Computer
similar compute nodes as NASA's
premier supercomputer, Pleiades


The example of an autonomous vehicle

Each autonomous car will generate 4TB of data per hour and incorporate LOTS of sensors i.e.

- 16+ 4K video cams
- 12 radar sensors
- 6 ultrasonic sensors
- 5 LIDAR sensors

It is estimated that each car will require 300 TFlops of computing power and a network pipe of 25 to 40GB/s to enable data to be moved from edge to core

Data is essential in this race – *Full ecosystem data*



AI – driven by the data explosion



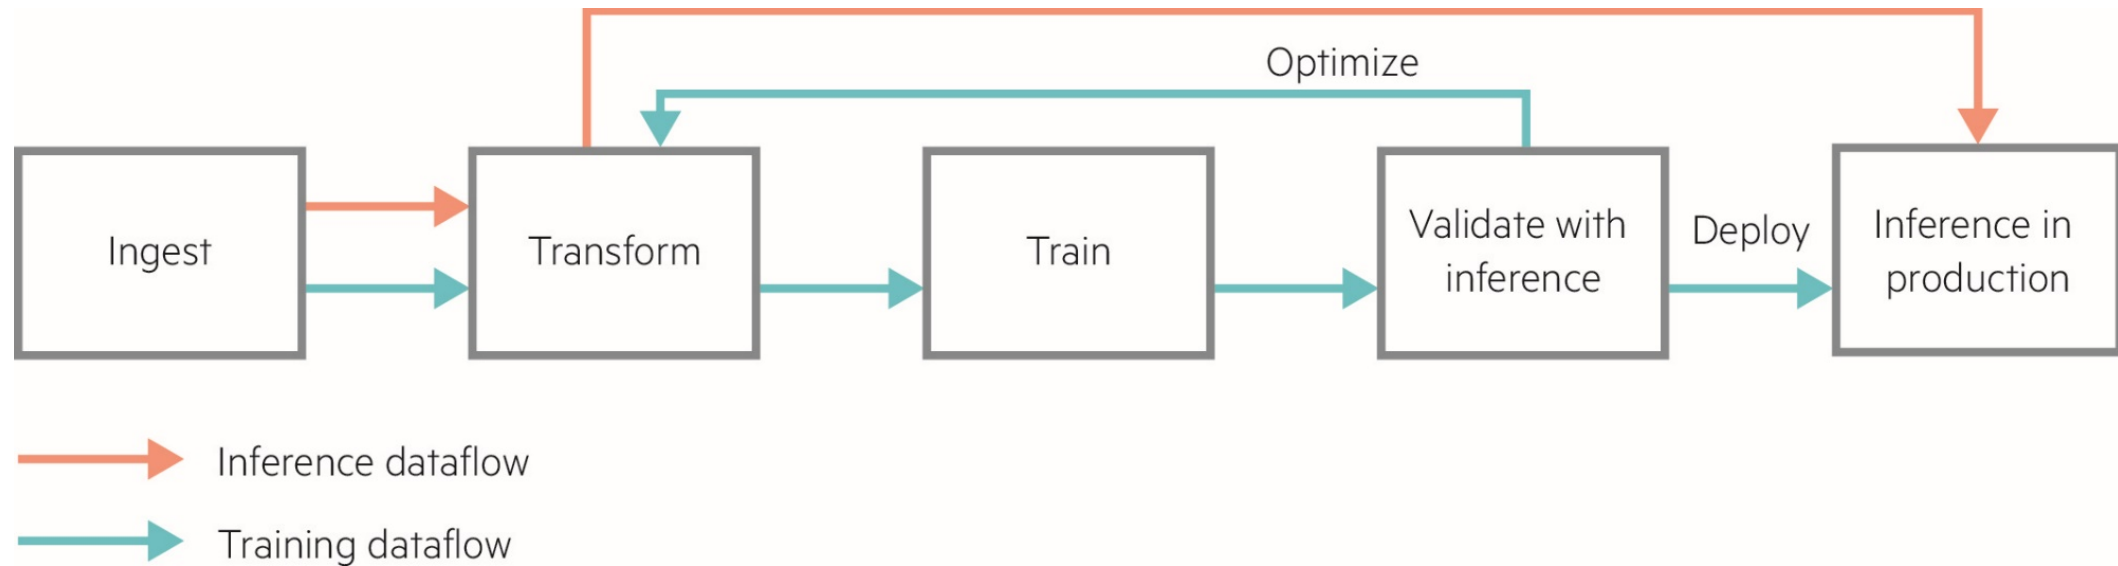
You can't stop the signal!

 Hewlett Packard
Enterprise



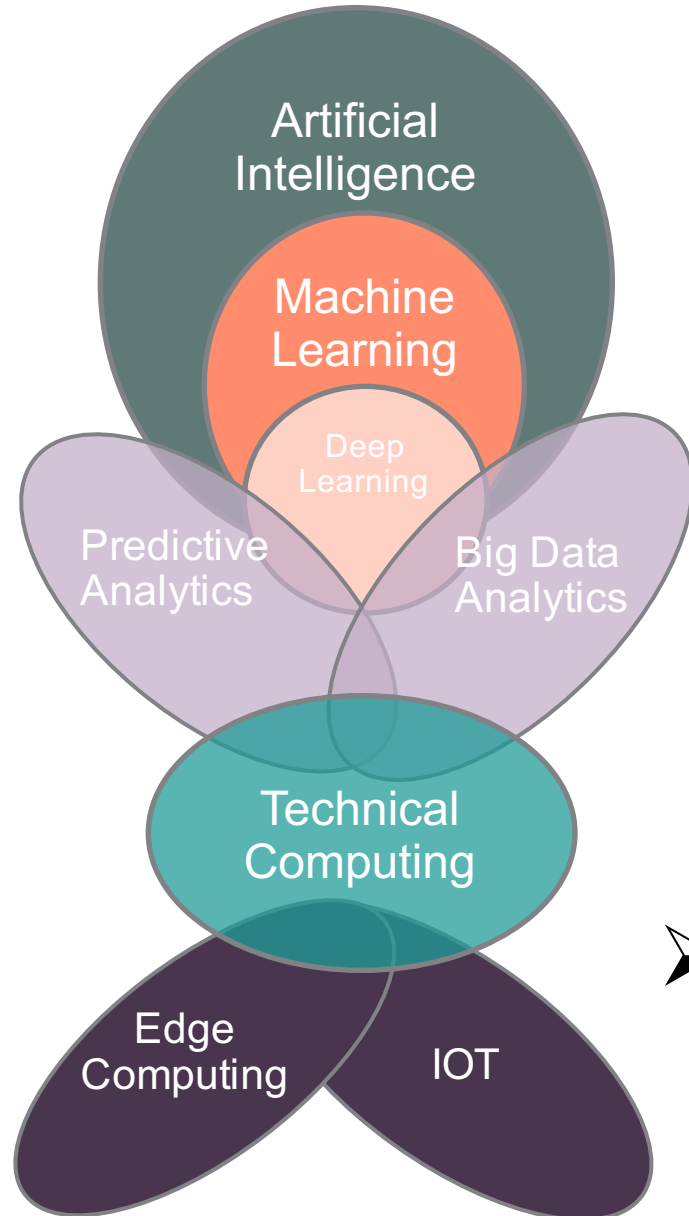
Mr. Universe, Serenity

Inferencing in the Deep Learning Development Cycle



- Market focused on training because we’re learning how to build AI applications
- Inferencing is used for model validation during the development cycle to assess factors such as performance, accuracy, and reliability, similar to software QA.
- After completion, models are refined further with inferencing
 - Ex. Autonomous vehicles need hundreds of millions of miles in testing to be as good as a human driver. Miles can be simulated via inferencing.

Convergence of HPC, AI, Big Data & IOT




Gartner forecasts that, **by 2020;**

- **20.4 billion connected devices**
 - **Up from 8.4 billion in 2017**

➤ **As high as 100 billion by 2020...**


Data Trends and their Challenges

Explosive Growth of Data



More than 37% of total data generated in 2020 (40 ZB) will have significant business value!

Need Answers ... FAST!




Increasing amounts of data to be analyzed But businesses demanding real-time insight!

Massive data with unpredictable growth, conventional systems can't keep up



Supercomputing
Conference 1996

sgi



Hewlett Packard
Enterprise

Big Graphics

Big Compute


Live 128-Processor
CRAY Origin2000™

- LPP: The Leading
SMP Architecture
- Pushing the Bar on
SMP Scalability

Big Data

Effortless Data Handling For
Complex Simulation Problems

Big Data, it's not new. We've been solving these challenges for decades!



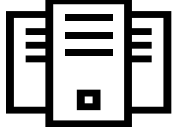
“Data always lives longer than the hardware it’s stored on.”

“Forward migration to new technologies should never adversely impact the users.”


-- Dr Robert Bell

HPC National Partnerships, CSIRO

HPC & AI | Data Management



HPC – AI – Machine Learning Cluster



Software



DDN STORAGE
Hewlett Packard Enterprise
Lustre

High-Performance Storage



XFS



Tier Zero



WEKA.io
Radically Simple Storage



All-Flash

Defined



HPE Data Management Framework

- Metadata management & data provenance
- Policy-based data migration with job scheduler integration
- Data protection, repair and disaster recovery



Storage



Tape






Zero Watt Storage





Object Storage & Cloud




Bring the applications to the data!

Provide the correct IO to the applications!

Memory Centric Computing Exascale

Scale-Out HPC

CFD
FEA
CEM
Chemistry
Physics
Scalable applications
Diverse mixed workload throughput

I/O intensive jobs


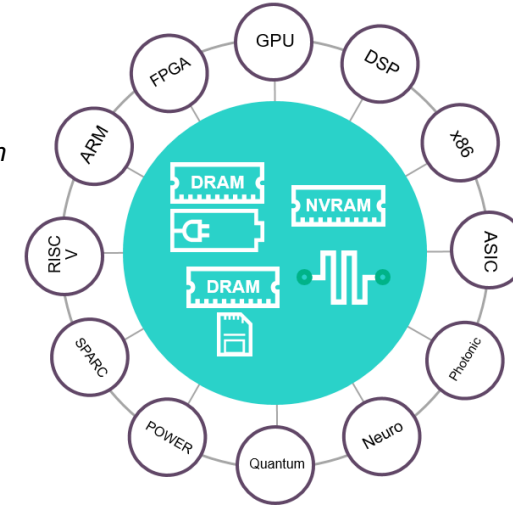
Local scratch
Out-of-core
Per-job private file systems
Dynamic shared scratch
Tier Zero

Scale-UP HPC

Bioinformatics
Genomics
Prototype / discovery
Algorithm development
Analytics
Visualisation

Science Cloud


IaaS
HPCaaS
Virtual laboratories
The long tail of HPC
Industry Outreach
Portals / collaboration




Distributed memory parallel



Distributed memory
Local flash



Large shared memory
Large local flash



Private Cloud

GEN Z



HPE Data Management Framework



Parallel File Systems

High bandwidth shared storage
Streaming workloads




Shared flash

High performance shared storage
IOPs workloads




Bulk disk / MAID

Fast recall
Spin down



Object Storage

Geographic replication
Multi-site locality



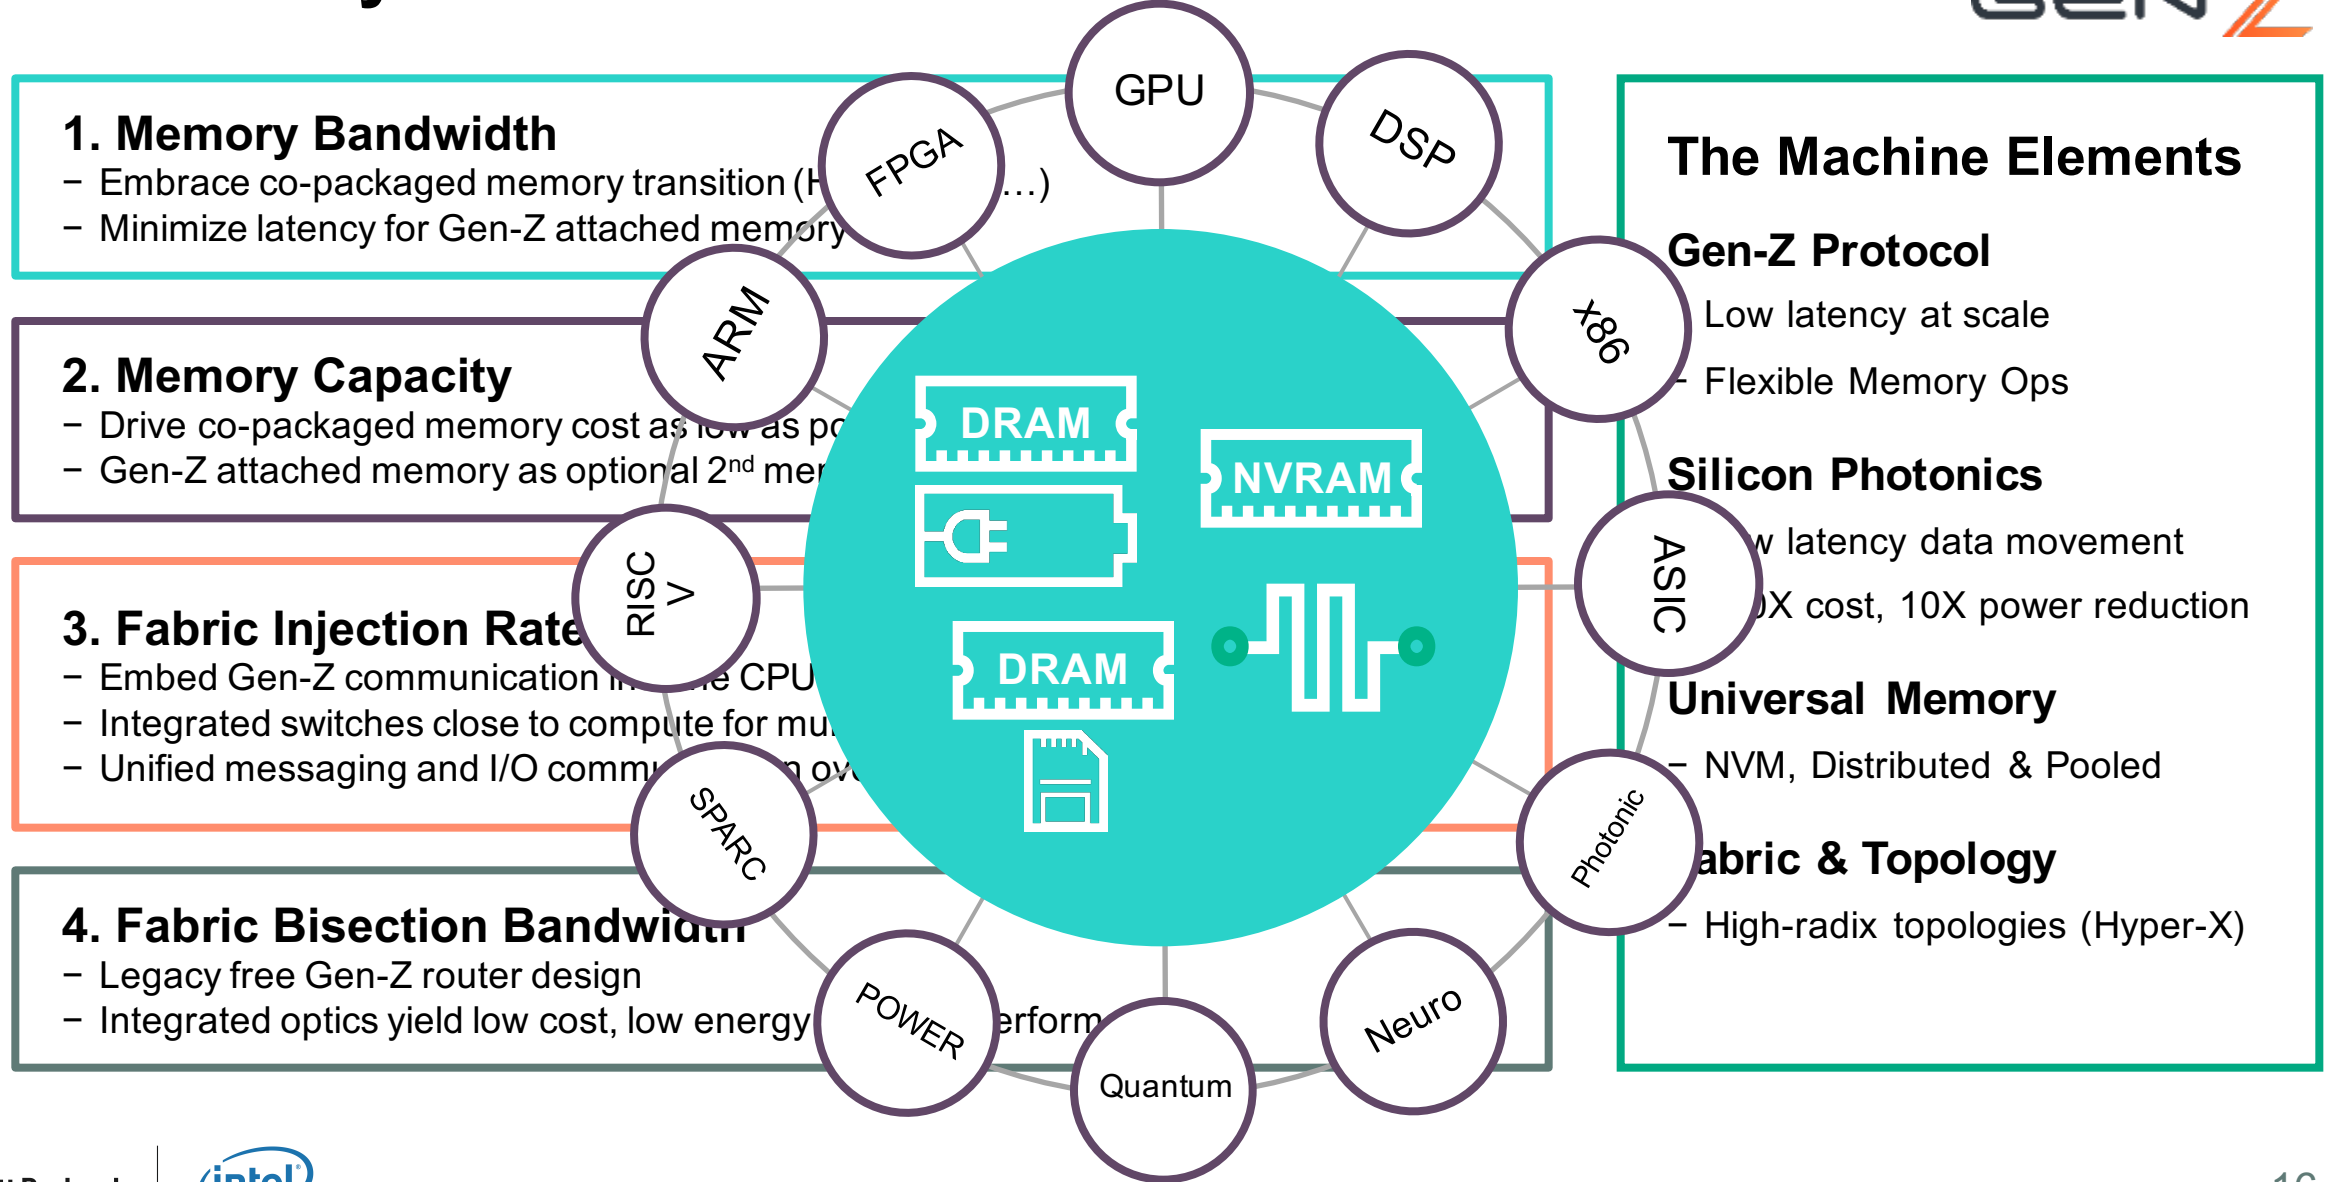
Tape

Secure
Reliable
Encrypted
On-site and off-site

Metadata

Version control
Scheduler integration
Workflow and automation
Prestaging of data
Query and discovery
Analytics and reporting
Fabric volume management
File system orchestration

A better system balance for Exascale




Future: PathForward Exascale R&D



3 Year Program
\$80M Budget
Co-Design & Prototype

100+ Engineers
5 Technology Areas
39 Deliverables



System Architecture




Silicon Design




Data Movement



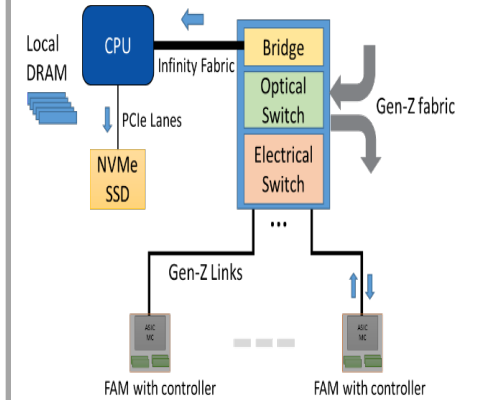
Optical Interconnects



Rx module



I/O Subsystem



Technology Development

- Exascale Architecture
- Application & Co-Design
- System Mgmt & Monitoring
- Power/Cooling
- System Prototyping
- Gen-Z Bridge Development
- Gen-Z Bridge Integration
- Network Modeling
- Router Architecture
- Topology (OZS ASIC)
- Gen-Z for HPC
- MPI over Gen-Z
- Photonics Sys Architecture
- Photonics Device Fab
- Electronic Device Fab
- Packaging of Photonics, Electronics, Fibers
- NVM Media Controller
- File and Data Access Library
- I/O Node Design & Proto



**Hewlett Packard
Enterprise**



Brought to you by Hewlett Packard Enterprise and Intel®.
Intel Inside®. Powerful Productivity Outside.

**Without supercomputers, we
would just be philosophers.**

–Stephen Hawking

R.I.P. – March 2018

Thank you