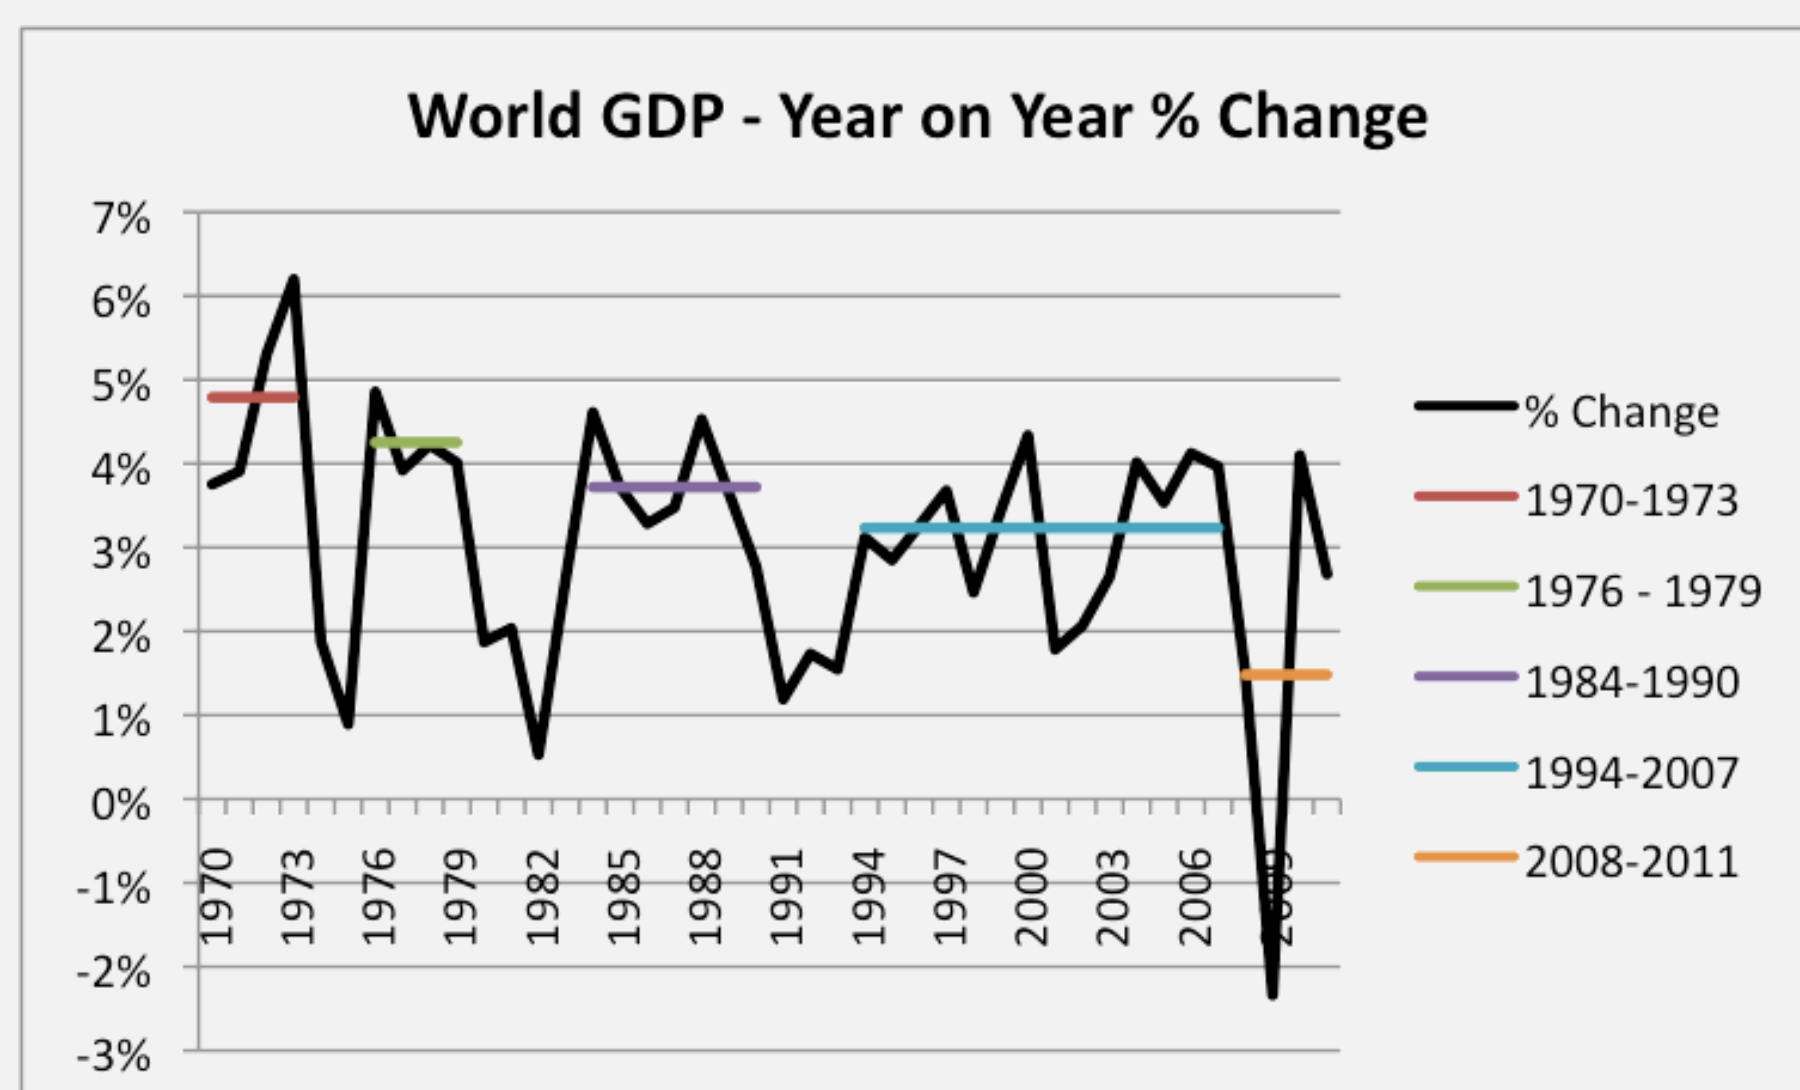# Adapting to Thrive in a New Economy of Memory Abundance

Processing technology has eclipsed memory technology for the past six decades, but processor-centric architectures are reaching their terminal efficiency. We can reboot computing on the basis of abundant memory enabled by emerging device physics, which will make computation, communication, and memory more efficient. This approach also provides a unique opportunity to address novel security threats with modern, systemic solutions.

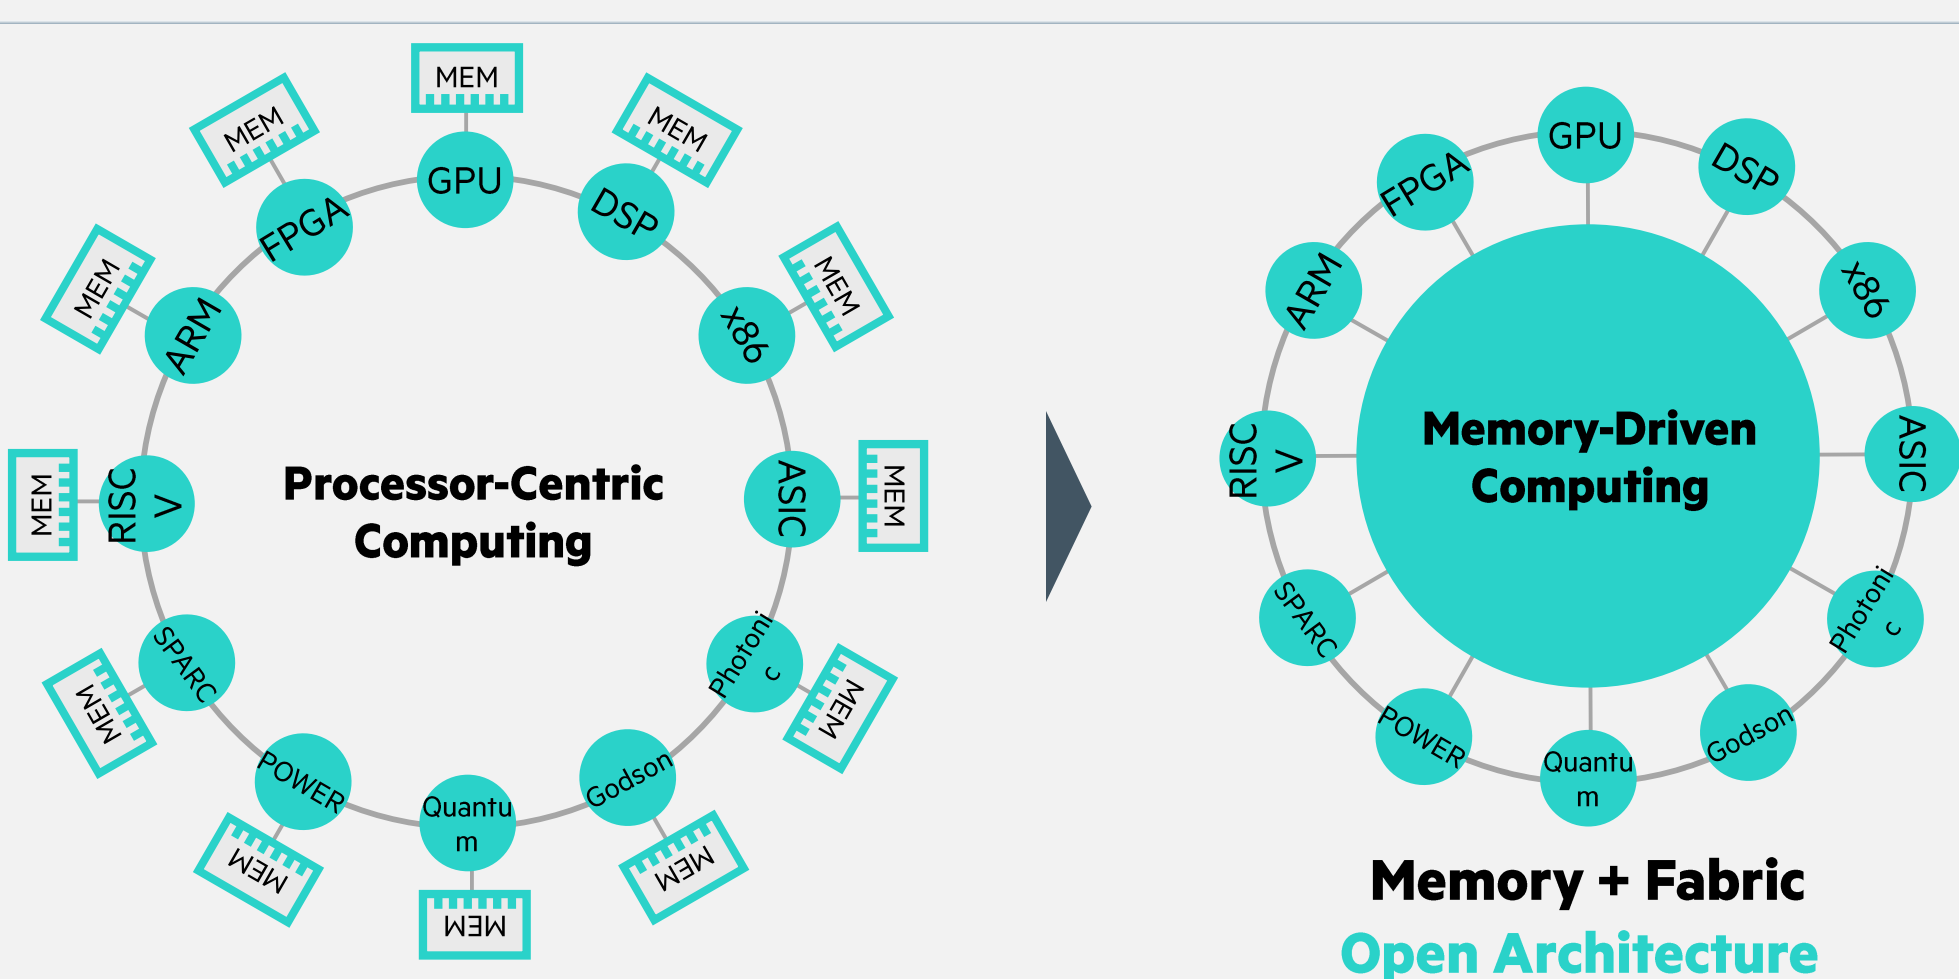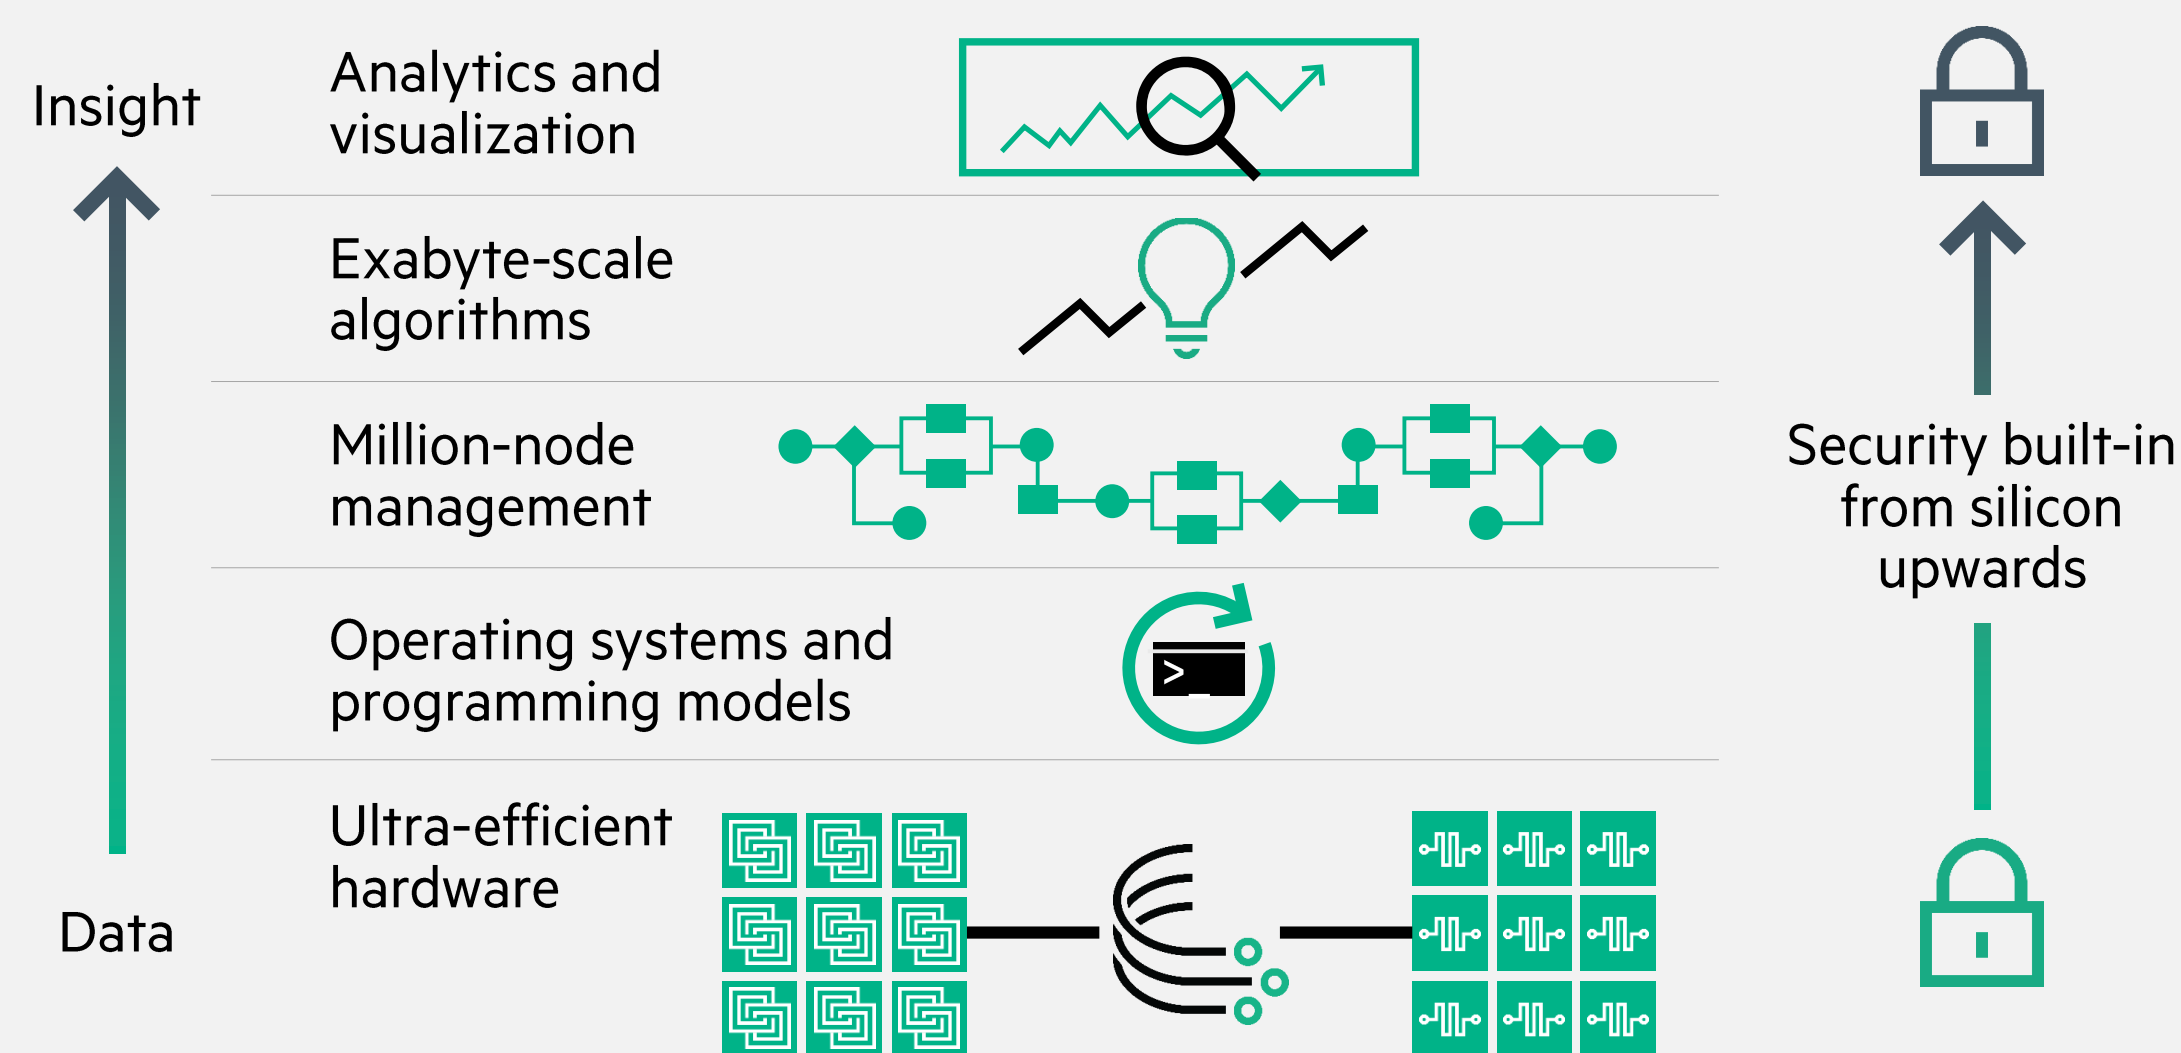## Moore's Law kept up with data creation in the traditional economy

World GDP - Year on Year % Change

Legend:
- % Change
- 1970-1973
- 1976 - 1979
- 1984-1990
- 1994-2007
- 2008-2011

The end of scaling at just the wrong time …

8B × **20B** × **100B** × **1T**

People × Mobile Devices × Social Infrastructure × Apps

Systems of Record ➜ Systems of Engagement ➜ Systems of Action
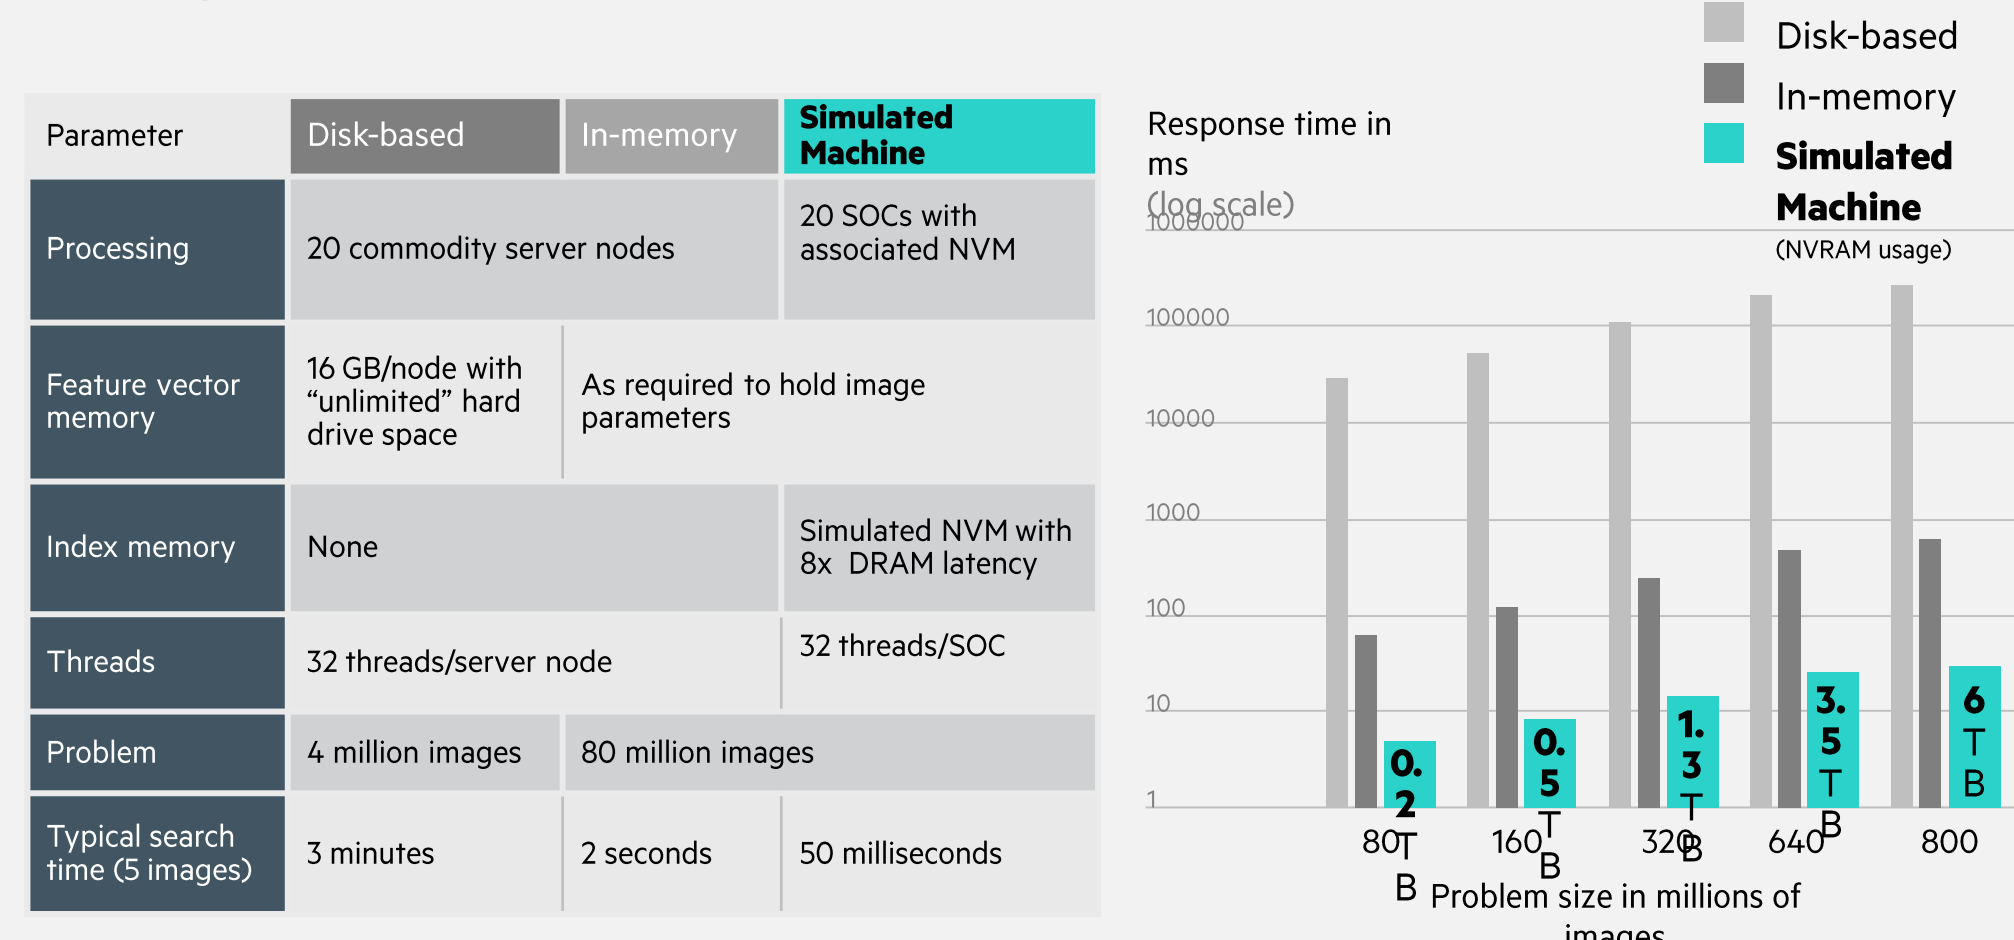
Hardware + software stack

Insight

- Analytics and visualization
- Exabyte-scale algorithms
- Million-node management
- Operating systems and programming models
- Ultra-efficient hardware

Data

Security built-in from silicon upwards

Processor-Centric Computing → Memory-Driven Computing

Memory + Fabric
Open Architecture

Performance demonstration – similarity search

From offline to decision time

- Millions of images
- Searches per minute

Disk-based	In-memory	Simulated MDC
4 / 0.3	80 / 500	**80** / **1200**

Use cases:
- Content-based image/video retrieval
- Near-duplicate web page detection
- Similar document retrieval
- Outlier detection for e-commerce fraud mitigation
- Fingerprint matching
- Scalable object recognition
- Nearest-neighbor classification
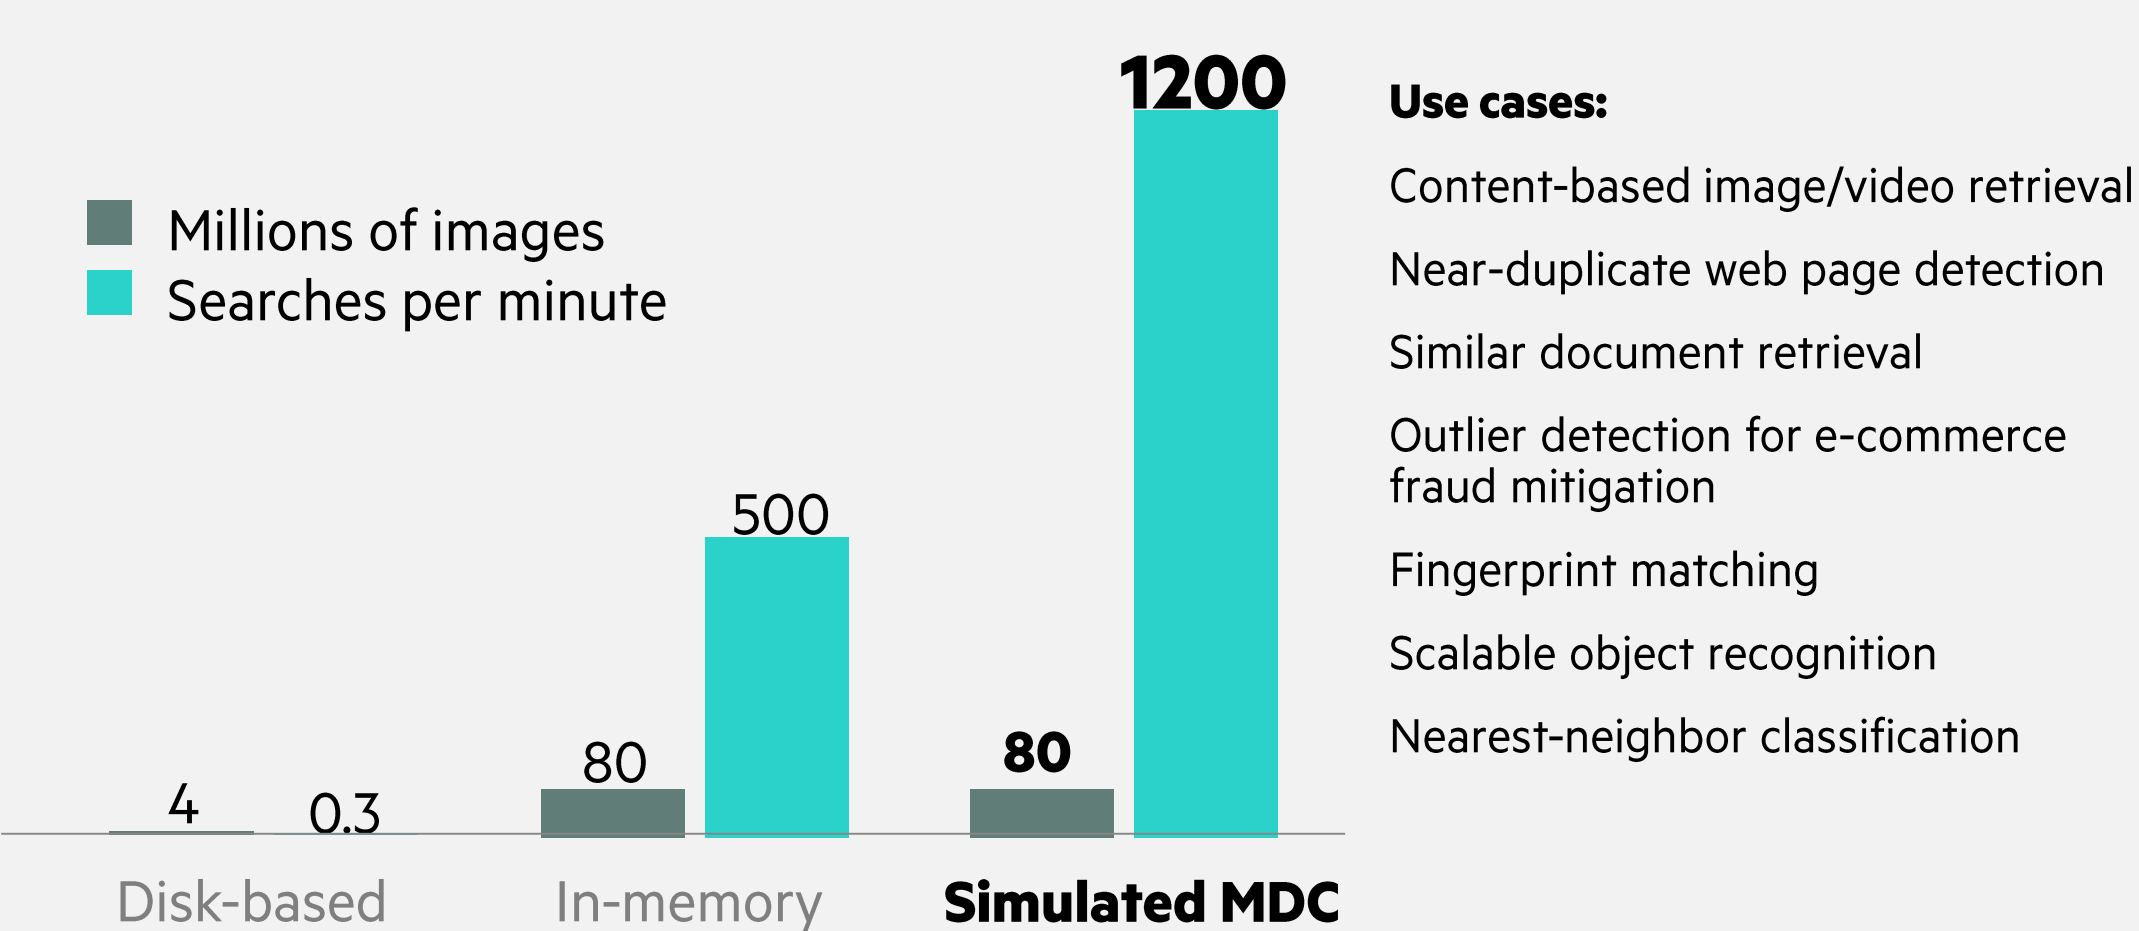
Performance demonstration – similarity search

Buying speed with persistent memory

Parameter	Disk-based	In-memory	Simulated Machine
Processing	20 commodity server nodes		20 SOCs with associated NVM
Feature vector memory	16 GB/node with "unlimited" hard drive space	As required to hold image parameters	
Index memory	None		Simulated NVM with 8x DRAM latency
Threads	32 threads/server node		32 threads/SOC
Problem	4 million images	80 million images	
Typical search time (5 images)	3 minutes	2 seconds	50 milliseconds

Response time in ms (log scale)

- Disk-based
- In-memory
- Simulated Machine (NVRAM usage)

Problem size in millions of images: 80T/B, 160T/B, 32B, 64T/B, 800

Graph analytics time machine

Massive memory and fast fabrics to ingest all data

"Are there any emerging new behaviors in my network?"

Critical node

External users

Remote users

Fast graph databases and the ability to look at things that change over time
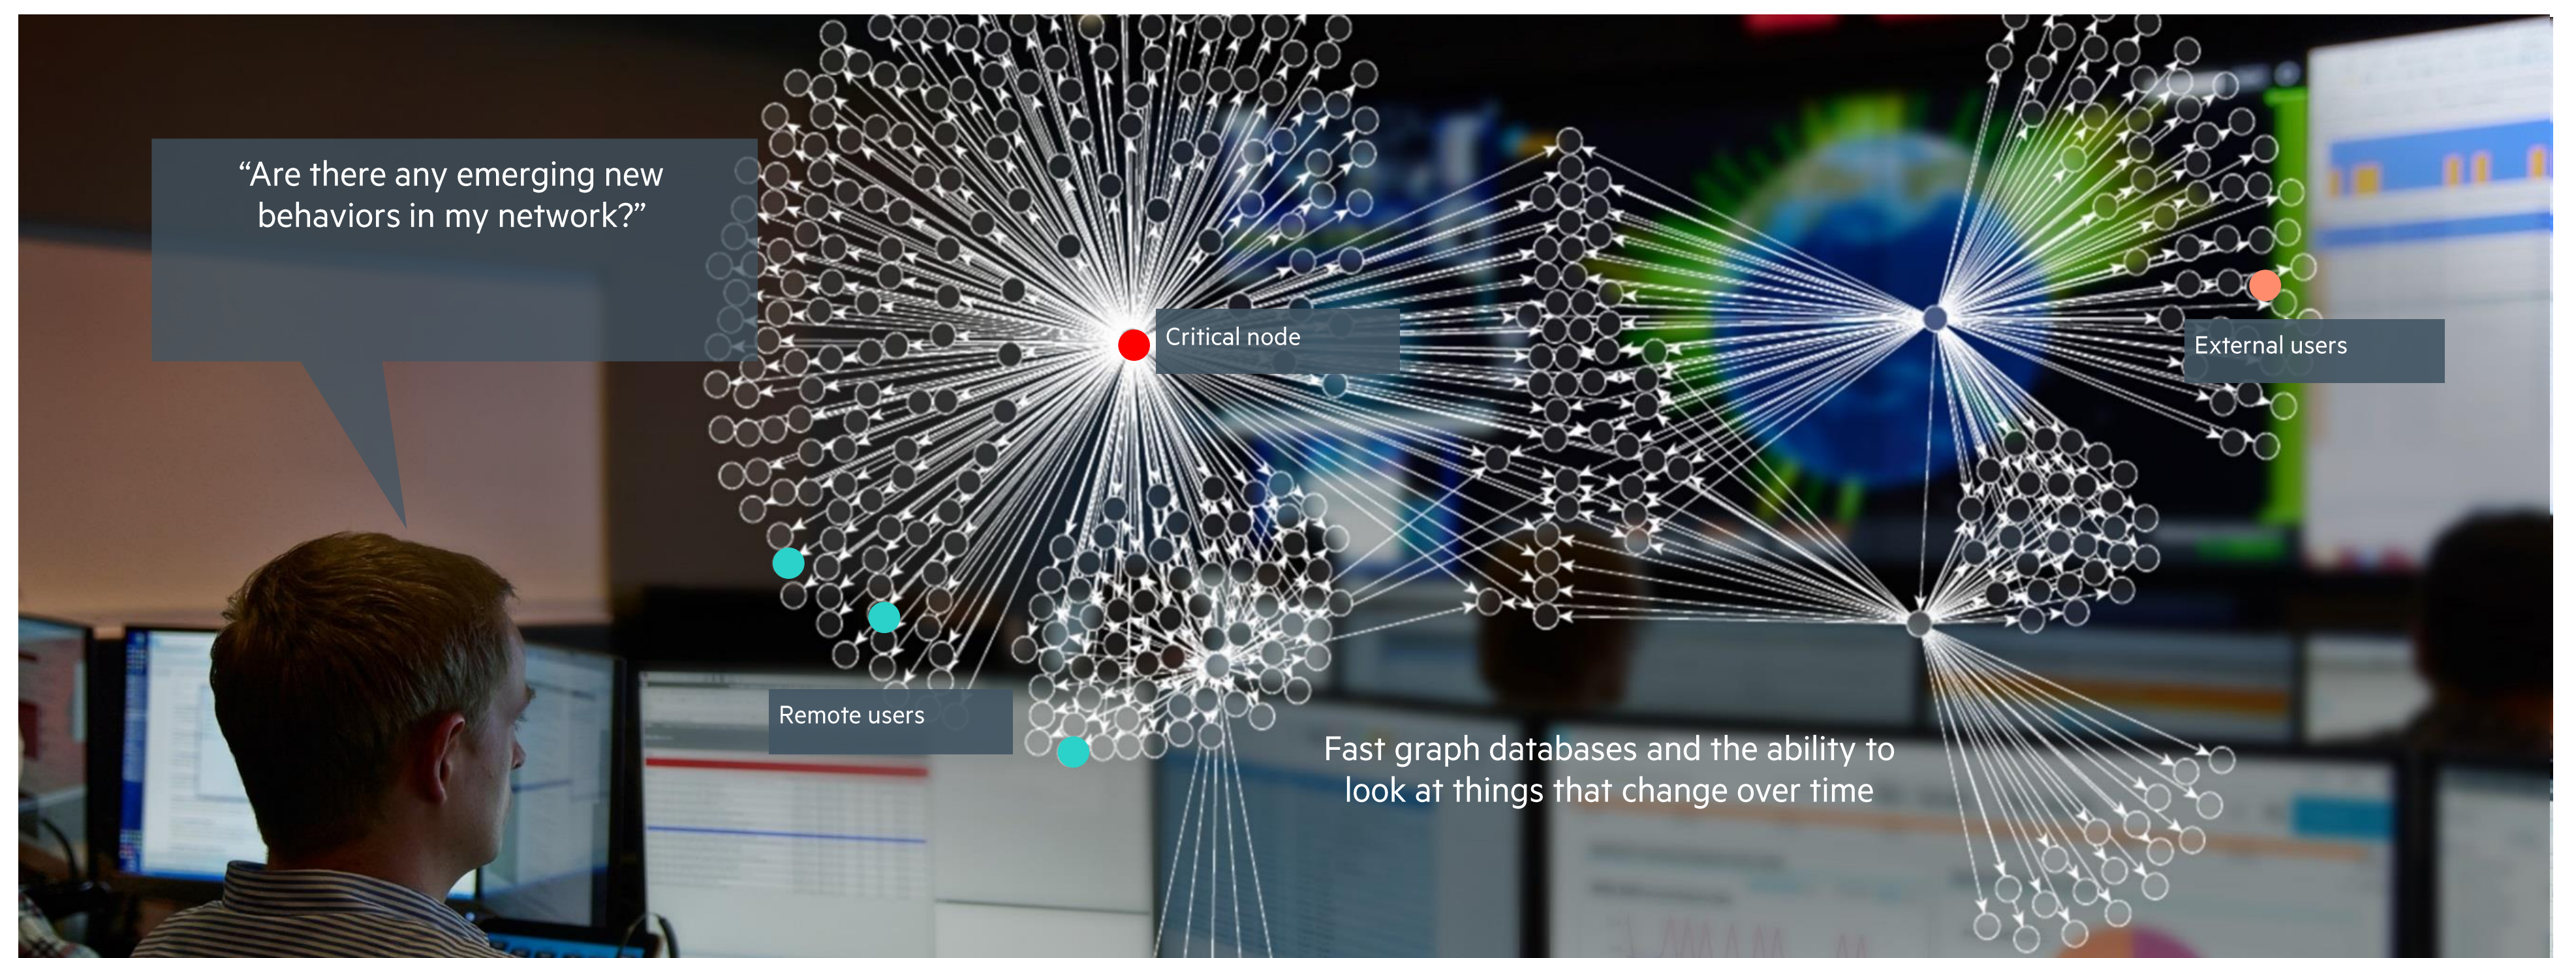
What if we could pre-compute an almost infinite set of "what ifs"?

Optimization over a large search space in real time becomes realistic

Solve complex problems **before** they happen
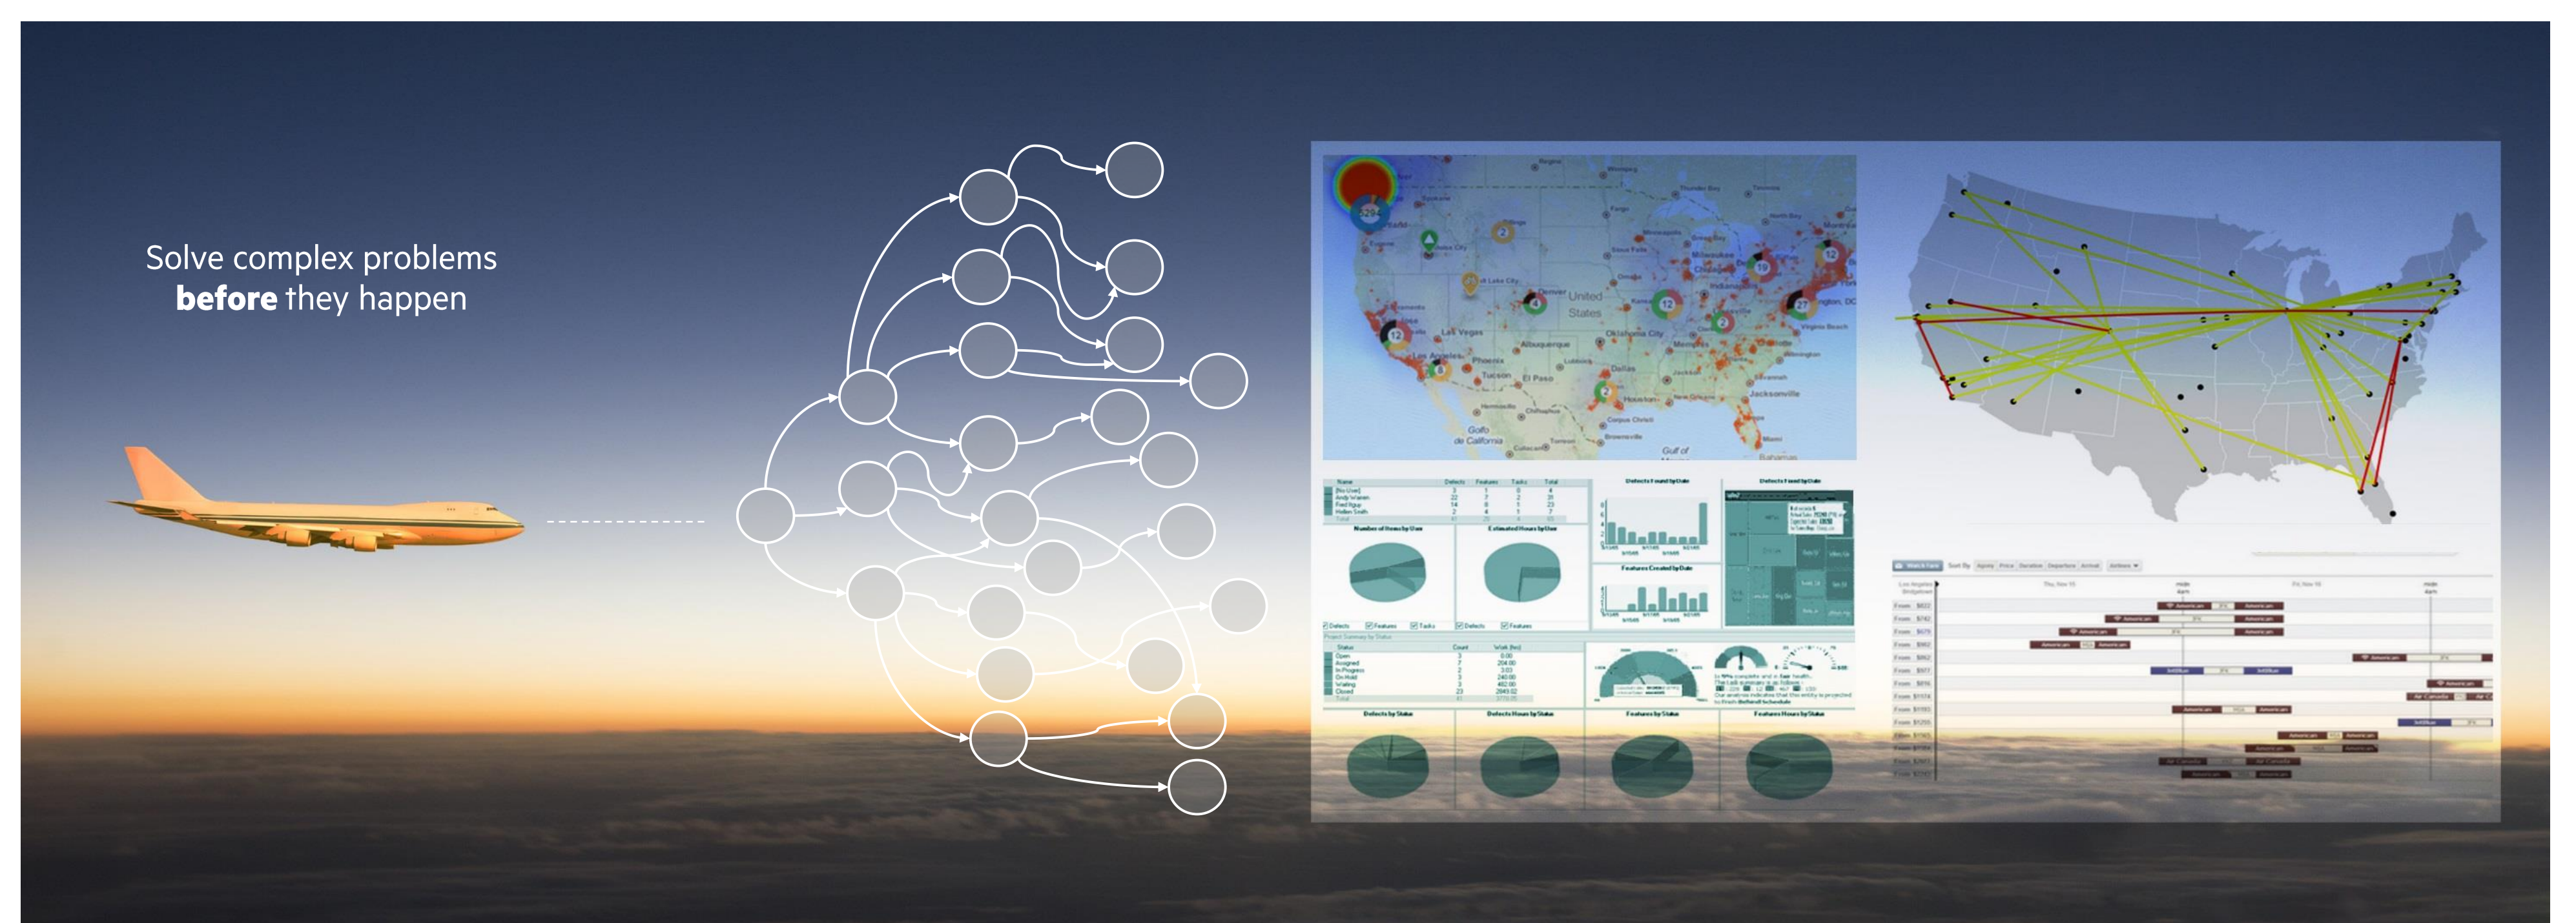
Machine Learning that can keep up with the now

Training up-to-date deep neural networks in minutes, not weeks

"Tell me the latest news on the severe weather."

Contextual ads

Relevant content

Keeping up with every click, every meme

Suggested topics

$$\lVert a, b; x, y \rVert = \frac{1}{2}\lVert h_{W,b}(x) - y \rVert^2$$

$$l_i^{(l)} = l_i^{(l)} - \alpha \frac{\partial}{\partial l_i^{(l)}} J(W, b)$$

Write code for our Fabric-Attached Memory Emulation for The Machine. The Fabric-Attached Memory Emulation is designed to get you started writing code for the Memory-Driven Computing architecture of The Machine on your laptop, before the hardware is even ready

https://github.com/FabricAttachedMemory

IEEE Rebooting Computing Responds to White House Nanotechnology Grand Challenge: "Sensible Machines" That are Smaller, Faster, and Lower Power

http://rebootingcomputing.ieee.org/archived-articles-and-videos/general/sensible-machine

To learn more about Hewlett Packard Labs, visit www.labs.hpe.com