The Machine

The Problem

- We are in the middle of a data explosion. The rate at which data is created is growing at an exponential rate, driven ever higher as the volume of machine-to-machine interactions begins to surpass human-to-machine interactions.
- This massive data explosion will be so big, it will flood legacy systems. Even the most efficient conventional data centers today will fail to keep up with this exponential growth.
- Big Data promises to positively transform every aspect of human science and industry, but without fundamental changes to the physics, economy and security of conventional information technology, that transformation will always be out of reach.
- A centralized solution to the data explosion is not workable. It will take too much time and energy to transport all the data, and a central location could leave it vulnerable to focused attacks.
- Current security technology isn't nearly good enough to inspect and classify every bit of data entering and leaving an enterprise, without bringing today's computing infrastructure to its knees. At the same time, the volume and sophistication of threats is rapidly increasing.
- What's more, we're going to run out of energy to facilitate all of these data needs. There's no point to building endless data centers if you can't get the electricity to run them.
- There are simple questions you can't even ask today because the ability to process the answers at this scale hasn't been invented. There are answers in your data but conventional tools have made them impossible to find.
- Bottom line: current systems can't handle where we are headed and we need a new solution. HP has that solution in The Machine.

The Solution

- The Machine will reinvent the fundamental architecture of computers to enable a quantum leap in performance and efficiency, lower costs over the long term, and improve security.
- It will be scalable from handheld to laptop to workstation to data center to supercomputer.
- Every computation, communication or memory action is designed to consume the least energy possible.
- The Machine will collapse the current memory hierarchy, combining memory and storage using Memristors, to enable the manipulation of massive data sets that are impossible using today's technology
- The Machine will communicate using light – photonics – instead of electrons for massive bandwidth and power-efficiency increases.
- To harness these new technologies, we are building the first all-new operating system in decades. Programming will be radically simplified and enable new applications that we can't start to build today.
- Many of the new component technologies of The Machine will enter the market over the next 5 years to augment the capabilities of HP's current products. This will allow HP to provide a smooth transition to the future capabilities provided by The Machine.

Recommended viewing:

Martin Fink announces The Machine at HP Discover 2014

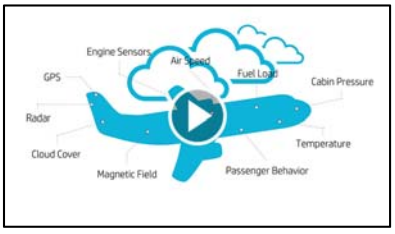

The Machine: Three-minute overview

Machine Components

By 2020, 30 billion connected devices will generate unprecedented amounts of data. The infrastructure required to collect, process, store, and analyze this data requires transformational changes in the foundations of computing. By discarding a computing model that has stood unchallenged for sixty years, we are poised to leave sixty years of compromises and inefficiencies behind. We're pushing the boundaries of the physics behind IT, using electrons for computation, photons for communication, and ions for storage.

The Machine will fuse memory and storage, flatten complex data hierarchies, bring processing closer to the data, embed security control points throughout the hardware and software stacks, and enable management and assurance of the system at scale. We're also not waiting for new ultra-efficient hardware, we've already begun to characterize workloads and develop novel analytics platforms and techniques for The Machine.

The industry is at a technology inflection point that HP is uniquely positioned to take advantage of going forward. The Machine demonstrates the innovation agenda that will drive our company, and the world, forward.

Energy-efficient processing

The processor remains a major component of total system power. By moving from general purpose processors to task-specific, system-on-chip (SoC) ecosystems we can simultaneously achieve significant gains in both compute performance and power-efficiency. Building on HP's groundbreaking Moonshot concept, we will automatically match computing tasks with the most appropriate processor type, including general purpose, graphics and digital signal processing cores.

Photonics and System Fabrics

Photonics research is focused on replacing the copper-based electrical connections used in today's IT systems with optical communication links, using micrometer-scale lasers on microchips to convert electrical signals to light, and back again. We are enabling future computing systems to communicate at the extreme scale and low latency that data-centric computing will require to transition from today's storage and networking stacks to a radically simplified universal memory model with extreme energy-efficiency.

Memristors and Universal Memory

Memristor is a new type of memory that offers the speed of DRAM combined with the low cost of flash memory and hard drives. It allows you to permanently store large amounts of data – similar to hard disks – but up to 100,000 times faster and at much lower energy.

Memristors are:

- Non-volatile, which means that the data stored can be maintained using no energy and will not be compromised in the event of a power loss
- High density, meaning they have a high storage capability in an extremely small area
- Very fast, ultimately capable of providing the speed and performance of system memory
- Low power, consuming up to 1,000,000 times less energy than flash or hard drive memory

Memristors will allow us to collapse the memory/storage hierarchy in today's computing architectures to create what we call universal memory. Imagine millions of lines of code that won't have to be maintained or executed and reducing the majority of energy consumption in today's computing architecture. Universal memory fuses high performance and easy to program main memory and persistent, reliable storage in one hyper-efficient package, leading to increased system data bandwidth, and decreased latency and power consumption.

System Architecture

As we build towards The Machine, we're also investigating the IT architecture, thermo-mechanical design and power infrastructure it will require. Fundamental shifts in technology - such as system-on-chip (SoC), hybrid memory systems and in-memory data storage - will enable us to create systems that uniquely meet the capacity, performance and TCO (total cost of ownership) demands of future data-driven workloads.

Embedded Security

A computing architecture that will enable processors to directly access large amounts of persistent memory creates new challenges in enforcing system security. There is no longer active data buried in volatile caches or memories and segregated data at rest. Every store must be assumed to be durable in the face of physical or cyber-attack. We're designing The Machine's architecture and operating systems to support much higher levels of security and assurance than today's commercial systems.

The Machine Operating System

As we fundamentally re-think computing hardware architecture in The Machine, we must also consider a complete re-evaluation of operating system principles, architecture and services to enable systems that will operate at unprecedented scale. We're developing system software technologies for a new class of applications that will allow for effective utilization of The Machine's massive, distributed non-volatile universal memory.

The Machine OS has two phases: first we will use familiar programming constructs to allow legacy applications to run with improved performance. The second phase will build on this to deliver the first all-new OS in decades, enabling developers to achieve orders-of-magnitude improvements in speed and scale from a whole new class of applications.

Management at Scale

Massively scalable systems will require a shift from today's systems-administrator centric approach to management. Recognizing that future computing will require distributed, complex systems that interact with data across multiple corporate, national and legal boundaries, while protecting customer privacy, we're working on techniques to configure and manage million-node systems.

Exascale Analytics

The unprecedented scale of future big data is also requiring us to develop new analytics algorithms, techniques, and tools that are ready to take full advantage of the multi-petabyte data sets that The Machine will enable us to manipulate. Today, Big Data means bringing all the data into one place. Tomorrow, some of data will be too big and too expensive to move. Tomorrow's analytics will work where the data is created, transforming data locally into intelligence which is then sent to a centralized learning engine powered by The Machine.

Learn more at hpl.hp.com

Created by HP Labs