


- Subscribe to EnterpriseTech Weekly Updates: @



# HP Puts Memristors At The Heart Of A New Machine

June 12, 2014 by [Timothy Prickett Morgan](#)



It is not every day that a major IT company throws down the gauntlet to its peers and declares a new computing architecture that will wipe away current system designs, but that is precisely what Hewlett-Packard did in a surprise move at its Discover 2014 conference in Las Vegas this week.

Meg Whitman, the company's CEO, and Martin Fink, its CTO and the head of HP Labs, took the center stage to unveil a massive project called simply the Machine, one that harkens back to a big bet that the company made back in the mid-1980s (along with Sun Microsystems) on something called Reduced Instruction Set Computing, which of course transformed the nature of processing on all manner of devices. Fink

joined HP back in May 1985 as the first fruits of “Project Spectrum” were coming to market, and he was among the first people to install the “Indigo” HP 9000 systems based on that processor into customer accounts. Fink wants HP to shake things up again like that with the Machine.

With the Machine, HP is going to tightly couple custom-tuned processors with its memristor non-volatile memory, which has been in development for the past six years, to create a system that eliminates the nine to eleven stages of memory and storage hierarchy that are in a server these days.

“HP has been talking about the individual component technologies for some time and now we are bringing them together into a single project to make a revolutionary new computer architecture that will be available by the end of the decade,” said Whitman. “This changes everything.”

HP is keen on using moonshot analogies these days, harkening back to the dawning of the Apollo missions to the Moon during the Kennedy Administration, and the Machine seems like more of a moonshot than the company’s actual Moonshot hyperscale servers, which are just starting to make their way into commercial datacenters.



With the Machine, HP is tackling a number of problems that plague hyperscale systems today and make the exascale systems of tomorrow difficult to bring into being. Current computer architectures are more energy efficient than their predecessors, to be sure, but up to 90 percent of the energy and time spent in a modern system is dedicated to moving data up and down the memory and storage hierarchy, not actually manipulating or processing that information. Systems need to be able to store more data and process it more cleanly, and of course the density of the compute and storage has to increase along with the energy efficiency gains.

This means, among other things, shifting away from copper interconnects between compute and storage components, and Fink showed off a mockup of a Machine processing and storage node that used silicon photonics to link a multicore system-on-chip to a bank of memristor memory cards.

The base element in the Machine is that memristor non-volatile memory,

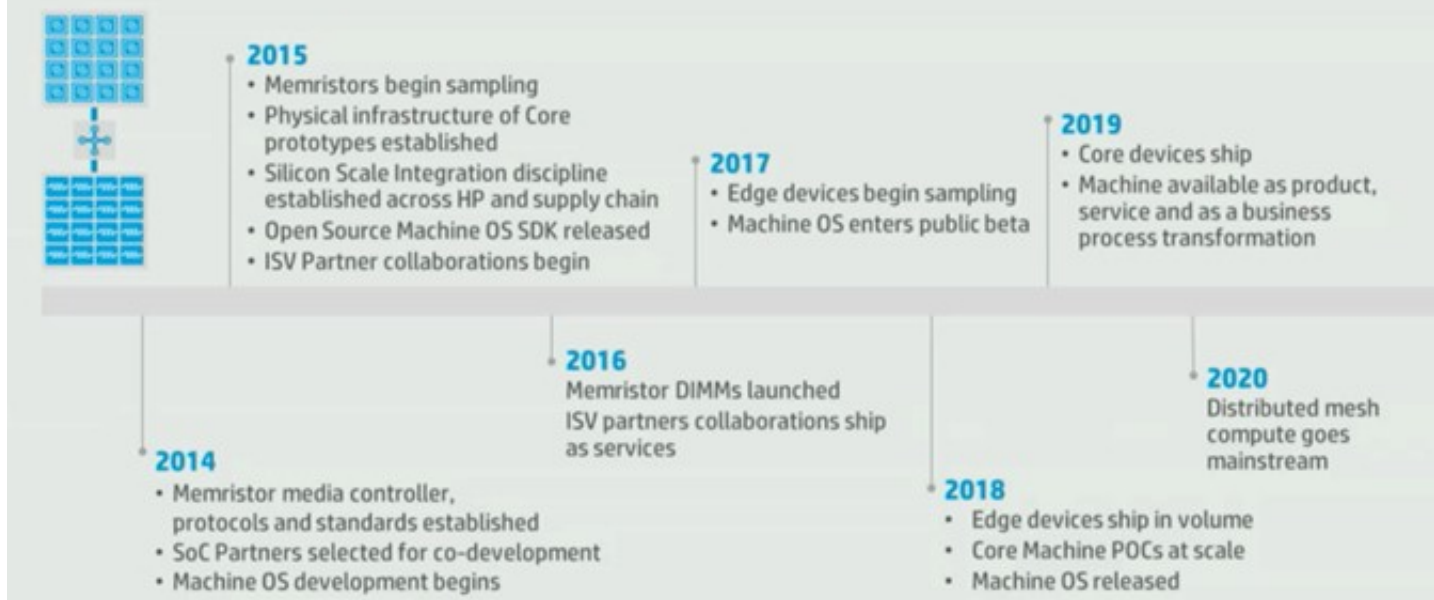
which HP has struggled to bring to market over the past several years. The memristor, short for memory resistor, has been an oddity examined by researchers since 1971. It is a special kind of resistor circuit that remembers the last voltage that has been applied to it even after the power is turned off, and as Fink explained, this is done by removing electrons from oxides in the memristor. Depending on the ionic state of the oxygen in the compound (he did not identify it, but titanium dioxide has been used in the past with memristors under development), it can be set as a one or a zero and thus be used for storage.

Here is how Fink outlined what the Machine would do. “We are going to specialize the compute to the actual workload that we are running,” he said. “We are going to connect that to a large, single pool of what we call universal memory. And then we are going to connect the two with a very high speed, low latency fabric based on photonics where we use light for communications. This will enable us to deal with massive, massive datasets, and not to just be able to take those massive datasets, but ingest them, store them, and manipulate them – and do this at orders of magnitude less energy per bit or unit of compute.”

The concept behind the machine could be summarized in six words, Fink said: *electrons compute, photons communicate, ions store*. You will be hearing that mantra a lot in the next several years from HP, we think.

HP has been working with SK Hynix to manufacture memristor memory since 2010, and had expected it to come to market first in 2013 and then sometime in 2014. Now the HP roadmaps relating to the Machine show it coming to market in stages over the next two years:

# Future History



The protocols for the memristor media controllers will be set this year, and HP will establish partnerships for the SoCs that will be used in the Machine as well in 2014. Memristor chips will start sampling in 2015 and DIMMs based on the technology will appear in 2016. Initial “edge devices” based on the Machine architecture will come out in 2017.

Both Fink and Whitman harped on the fact that what most of an operating system is doing these days is shuffling data in and out of various layers of storage, and HP intends for the Machine to have one layer as a universal memory pool, eliminating the need for various layers of disk and flash memory in the hierarchy. Presumably the processors in the SoCs in the Machine will have their own main memory – but perhaps not. Fink did not elaborate on that. But what he did say is that a rack of Machine modules would have 1 PB of memory, and that HP could envision up to 160 racks all linked together using the photonics fabrics to have any compute node in the system to be able to address any byte anywhere in those 160 racks in under 250 nanoseconds. He added that HP has “line of sight” to photonics that can deliver 6 Tb/sec of bandwidth on a single piece of fiber optics, and that if you wanted to do that with copper, you would need a bundle of wires about the size of a

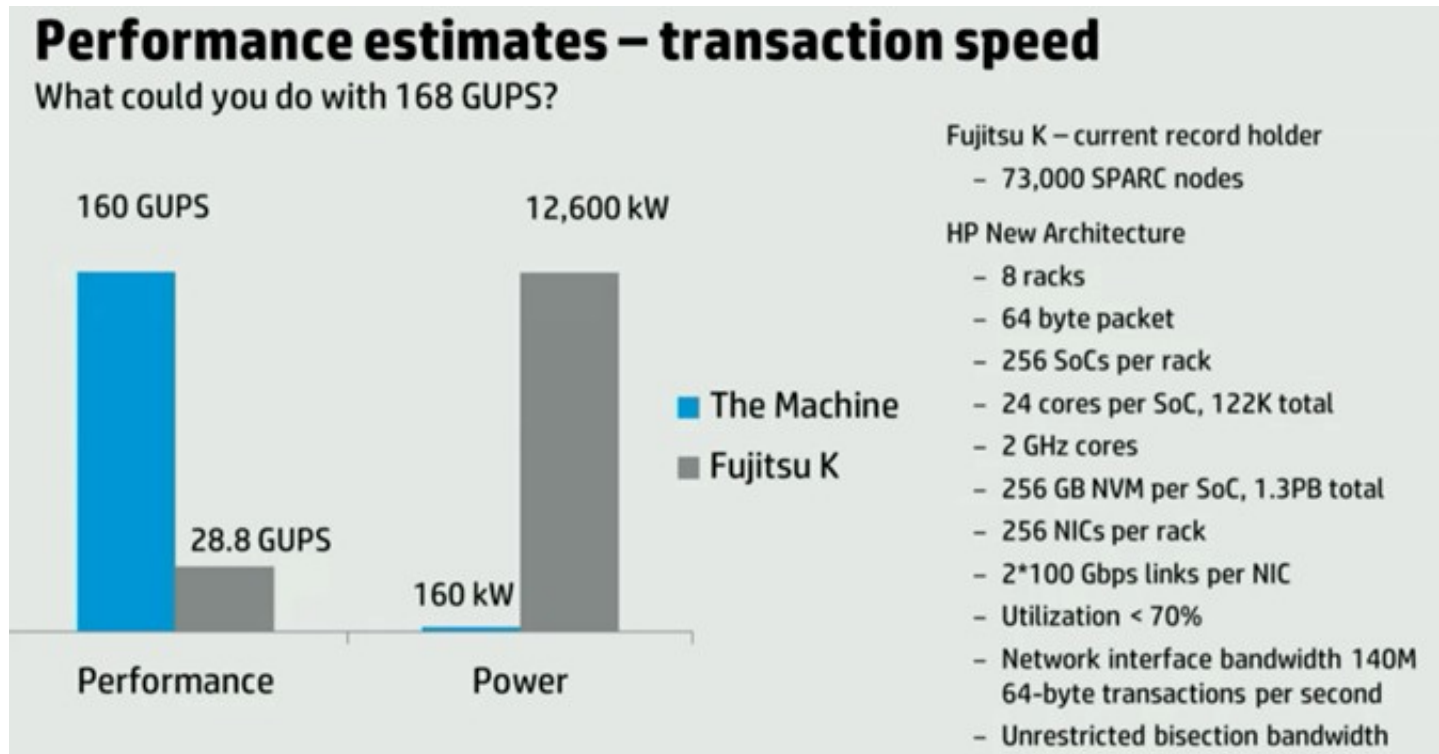
large dinner plate and it would require thousands of times more energy to pump the signals through it compared to the fiber. The memristors have a switching speed on the order of picoseconds, by the way.

As you might imagine, the new architecture of the Machine will require some changes to the operating system. In fact, just as HP is throwing out a lot of the storage hierarchy in the Machine's design, researchers at HP Labs are right now gutting the Linux operating system for servers and its Android variant for clients of all the unnecessary code used to manage storage layers. This, he explained, was why it was called the Machine in the first place. It is not intended to be the basis of a server or a laptop or a smartphone, but rather used in all of these, and in fact, when he held up the memristor prototype above, he said we should consider this as the basis of a future smartphone that would have 100 TB of capacity on it and store our entire lives.

“We are announcing, as part of the Machine, our intent to build a new operating system, all open source, from the ground up optimized for non-volatile systems,” Fink said. “We want to reignite in all of the universities around the world operating system research which has been dormant or stagnant for decades.” He added that HP was happy to work with any operating system supplier to make sure its code would run on the Machines – so that means Microsoft Windows for the most part, and maybe a smattering of Unixes or maybe not. Windows and Linux dominate the enterprise datacenter these days. It is not clear what hypervisors or containers will be necessary on the Machine, but very likely a mix of the two virtualization approaches will be available as is the case with modern systems. As for systems management, HP Labs has cooked up a tool called Loom, which allows for thousands of servers and up to 40,000 virtual machines to all be managed from one console; this will be extended to span the Machine.

So where does the rubber hit the road with the Machine? Fink offered a

comparison between the K supercomputer created by Fujitsu to an eight-rack setup of the Machine:



The performance measurement that Fink used is giga-updates per second, or GUPS, which is established by [the RandomAccess benchmark](#) that is part of the HPC Challenge suite of tests and which is used by the Defense Advanced Research Projects Agency to evaluate supercomputers. The initial K system had 73,000 Sparc64-VIIIfx processors, each with eight cores, running at 2 GHz and uses a 6D torus interconnect called Tofu to link nodes together. Fink said that K burned 12.6 megawatts of juice to deliver 28.8 GUPS of performance; that machine took up about 788 cabinets of space. But eight racks of the Machine could deliver 168 GUPS in a 160 kilowatt power envelope. “We believe we can achieve a factor of six times performance increase by focusing the workload, but do it at 80 times less energy,” said Fink. And at about a factor of 10X reduction in space, too.

In the configuration above, HP is storing data and shipping it around in 64 byte sizes. Each SoC has 24 cores on it, and it could be an X86 or

ARM core or both in the initial systems. It is very likely *not* a proprietary core created by HP – particularly after Fink oversaw HP’s Itanium-based server business. HP is not interested in introducing incompatibilities in the datacenter, not after that experience. An HP spokesperson tells *EnterpriseTech* that “the Machine will use a wide range of processor types to match the right technology to the right workload, and that these may include general purpose, mobile, graphics and digital signal processing.”

There are a total of 256 SoCs per rack and at 24 cores each across eight racks, that should be 49,152 cores, not the “122K total” cited in the chart above. The cores run at 2 GHz and each SoC has 256 GB of memristor memory, for a total of 1.3 PB of non-volatile memory per rack. Each SoC has its own network interface card, and it has two 100 Gb/sec links coming off that. HP expects to be able to run the system at greater than 70 percent utilization and be able to push 140 million transactions per second out of the photonic network interfaces linking the nodes to the universal storage pool using 64 byte file sizes.

At this point, approximately 75 percent of the researchers at HP Labs are dedicated to the Machine effort, the company confirmed to *EnterpriseTech*, building on the memristor and two photonics projects that were already well under way.

---

**Share this:**



Categories: [Editor's Pick: Front Page](#), [Slider: Systems](#), [Systems](#)

Tags: [HP](#), [hyperscale](#), [Linux](#), [memristor](#), [SOC](#), [The Machine](#)

