REPORT REPRINT

# HPE invites developers to play in its memory sandbox

DANIEL BIZO

02 JAN 2019

The company is actively seeking potential candidates with memory-bound problems to experiment with the idea of flexibly assigning memory to compute in a scale-out environment.

THIS REPORT, LICENSED TO HEWLETT PACKARD ENTERPRISE, DEVELOPED AND AS PROVIDED BY 451 RESEARCH, LLC, WAS PUBLISHED AS PART OF OUR SYNDICATED MARKET INSIGHT SUBSCRIPTION SERVICE. IT SHALL BE OWNED IN ITS ENTIRETY BY 451 RESEARCH, LLC. THIS REPORT IS SOLELY INTENDED FOR USE BY THE RECIPIENT AND MAY NOT BE REPRODUCED OR RE-POSTED, IN WHOLE OR IN PART, BY THE RECIPIENT WITHOUT EXPRESS PERMISSION FROM 451 RESEARCH.


©2019 451 Research, LLC | WWW.451RESEARCH.COM

In an ongoing effort to promote its Memory-Driven Computing (MDC) architecture (which was born from The Machine project), Hewlett Packard Enterprise has been whipping up interest for a recent initiative targeted at software developers and customers. The company is actively seeking potential candidates with memory-bound problems to experiment with the idea of flexibly assigning memory to compute. Even though many of the technologies that will underpin this capability are still in development, HPE says it will be introducing products that decouple memory capacity from processors 'soon enough.' The company is hopeful that the Memory-Driven Computing Sandbox, a custom software environment running on the company's scalable high-end servers, will help software developers nurture the idea of completely disaggregated resources, and prepare for the gradual realization of the MDC architecture.

## THE 451 TAKE

Even if the commercial introduction of technologies behind the idea of MDC is still some years out, HPE's efforts to demonstrate the underlying concepts early make perfect sense. An MDC-class system will not be an incremental improvement on existing systems, but will introduce architectural developments that will change fundamental assumptions about software engineering, server design, installation and operations. To take advantage of the expected benefits, software developers and IT buyers need to adapt their thinking – and that's the goal HPE's MDC Sandbox serves.

## CONTEXT

HPE is a major supplier of enterprise-class IT infrastructure, software and related professional services that include hybrid cloud assessment and deployment services, as well as the design and construction of datacenter facilities. HPE was formed in November 2015 after a spinoff from Hewlett-Packard (HP Inc), which retained the PC and printing units. It had revenue of $37.4bn in its 2017 financial year (ended October 31) and employs about 45,000 people worldwide.

About five years ago, HPE found itself in a bind. The enterprise and government segments – HPE's traditional strongholds – had been growing slowly, if at all, as new capacity was predominantly being added at various major cloud and web service providers, which typically don't buy much, if any, equipment via traditional channels and vendors, in their quest for rock-bottom procurement costs. Having decided it didn't want to go after low- (or even negative-) margin deals, HPE's systems business needed a way forward – or rather, a way out.

One of the areas identified as a pillar for its future is edge computing – an attempt to gain a first-mover advantage by participating in the transformation of operational technologies that use fixed-function specialty hardware into an IT-based programmable environment. HPE has made some major investments to develop IT systems for non-datacenter environments (such as the HPE Edgeline Converged System), and has entered into partnerships with industrial equipment makers such as ABB and Schneider Electric to offer micro-modular datacenters and gain exposure to an industrial channel and customer base that it had limited access to in the past.

All the while, HPE has embarked upon a long-term research and development program to develop an architecture that is markedly different from the servers and storage systems of today. Initially, it dubbed this ambitious next generation of systems The Machine, a concept it first presented in 2014. It now calls this Memory-Driven Computing to refer to the new architecture and the body of technologies behind the idea. HPE says the historical trajectory of development of current general-purpose architectures is simply not sufficient to meet the growing IT challenges in performance, power efficiency and cost-effectiveness, hence the new direction.

MDC extends the original Machine idea that was built on three key pillars: workload-optimized processors (accelerators), storage-class memory, and photonic system interconnects and networks. HPE has subsequently opened up the program to current and other future technologies in order not to be prescriptive about what customers will be able to do with the MDC systems and how. Indeed, flexibility is central to the value proposition of MDC – it comes in the form of total 'composability' where systems are dynamically born, via software controls, from pools

of components (e.g., main processors, accelerators, memory, storage, network links) sitting on a single, unified fabric. This is to eliminate the hardware imbalances between compute, memory and IO seen in today's servers that contribute to infrastructure sprawl – even the best-run cloud infrastructures suffer considerable losses of capacity because of it.

Realizing this concept requires nontrivial research and development efforts that require vast resources from the industry. Fortunately, numerous players are pursuing these goals, including heavyweights such as Intel, NVIDIA, Cisco and some hyperscale operators, and key players from across the ecosystem are collaborating on the fabric aspect through the Gen-Z Consortium.

## SOFTWARE-DEFINED MEMORY

In the almost five years since HPE's public announcement of this plan, MDC remains a goal with ongoing R&D efforts. In the meantime, Intel, which is also a proponent of server disaggregation, has productized and is commodifying two key technology components that are outwardly similar: silicon photonic networks and storage-class memory (marketed as the Optane family of products). However, HPE is not content with existing server architectures and upgrades to individual components.

But decoupling memory from processors, a key feature of MDC, is a major R&D undertaking. Intel doesn't expect to achieve it in products for another few years at least. This means that HPE and its partners in the Gen-Z Consortium (from which Intel is notably absent) have room for differentiation as one of the development leads. In 2019, HPE plans to demonstrate memory expansion capability over fabric, after which a whole suite of memory banks, interconnects and switches will follow in 2020 and 2021, per current plans.

With the first generation of decoupled memory products a couple of years away if development goes to plan, HPE has started whipping up interest in MDC among the developer and user communities. What makes this possible (well ahead of having workable hardware) is HPE's latest mission-critical system, Superdome Flex, which offers scalability and partitioning features to create an environment that resembles a future MDC infrastructure. In order to make its high-end Superdome Flex systems accessible to developers, HPE offers a 'sandbox' as a service for those that qualify for the program via its Pointnext services arm.

The Superdome Flex is a manifestation of joint technology development (coming from HPE's acquisition of SGI in 2016) that infuses HPE's mission-critical system design with SGI's system interconnect technology, NUMAlink, which can glue multiple servers together. Running the Sandbox is a 64-processor system (composed of two interconnected Superdome Flex frames) configured with 64TB of memory a modified Linux software stack and hardware partitions – features found in high-end systems to create multiple isolated operating environments as opposed to a single-image system.

Each hardware partition acts effectively as a server in its own right. The Sandbox is different from classical partitioning in that the software gives partitions access to a shared memory pool via software on top of a dedicated 'local' memory bank. This shared pool is carved from the capacity of individual partitions. The servers share memory data via the Superdome Flex's system interconnect.

This allows developers to allocate memory dynamically as needed by the application, and share data across servers via memory. The Sandbox can create up to 16 independent server instances and pool terabytes of memory to simulate a scale-out infrastructure in which memory amount and IO resources can be flexibly matched to compute needs.

A simple example would be a virtualized environment where many servers are limited by their memory capacity or IO configuration, and processors remain underutilized. Nonetheless, maximizing memory capacity and network links across all servers is usually cost-prohibitive, and could lead to vast amounts of stranded capacity. When allocation of resources happens dynamically from shared and hot-expandable pools, capacity can closely track requirements. Another use cases involves massive many-terabyte datasets against which computationally smaller jobs are running, which don't fit existing architectures well because compute and memory capacity scale hand-in-hand. Giving one or a couple of processors access to tens of terabytes of memory is not possible for today's IT systems.

Future configurations of the Sandbox should also help developers prepare for storage-class memory. Once equipped with non-volatile memory modules (as opposed to purely DRAM), developers can not only think of allocating memory more freely to processors, but also treating memory as a lightning-fast persistent storage tier by modifying their applications. When an application uses persistent memory for writes (saving the data), it avoids the storage subsystem, which carries much higher resource overhead and is generally orders of magnitude slower. 451 Research expects HPE to announce such plans in the near future.

Applications that will likely benefit the most from large pools of persistent memory include high-volume online transaction processing systems found in e-commerce and financial services, running analytics against massive datasets, training of deep neural networks, and even 'classical' scientific and technical supercomputing workloads. HPE has demonstrated that this kind of memory access can deliver multiple order-of-magnitude gains for algorithms in genomics, graph analytics, search and Monte Carlo simulations that are widely used in algorithmic trading of currencies and securities.

## COMPETITION

HPE's Memory-Driven Computing products will compete in the broad market of server and storage systems, where it still enjoys a large business amid stiff competition. Dell Technologies, Lenovo, Inspur and Fujitsu are formidable opponents, but HPE also competes for dollar spend with cloud infrastructure providers trying to capture more enterprise and government business. High-performance computing is also a key area of interest for HPE, where it strengthened its portfolio with the acquisition of SGI in 2016.

MDC can also be seen as the future architecture for every style of compute: edge, hyperconverged, scale-up and HPC. For the next generation of hyperconvergence, the configuration of individual nodes will no longer be confined by the limitations of current server platforms. The HCI market continues to grow at a rapid pace, although it remains relatively small compared with the rest of the IT infrastructure landscape, where Nutanix is leading the charge in close partnership with Dell. HPE boosted its position with the acquisition of SimpliVity for $650m in 2017.

## SWOT ANALYSIS

### STRENGTHS

As one of the largest IT infrastructure vendors, HPE carries enough weight to affect the future direction of the IT industry. Its deep pockets mean the company can afford to invest heavily in productizing emerging technologies.

### WEAKNESSES

An arguably convoluted product portfolio and quarterly sales pressure in a difficult environment aren't ideal ideal for risky development projects that don't necessarily meet immediate market requirements. HPE gains little from the growth in spending on cloud services.

### OPPORTUNITIES

With storage-class memory, photonic system interconnects and a blossoming world of accelerators pulled by interest in deep learning, HPE has a historic opportunity to steer parts of the IT market in its favor if it manages to pull off its concept of The Machine.

### THREATS

While commercial exploitation of these technologies is looking increasingly likely or is already a reality, the speed of market acceptance remains to be seen. The market might choose to go with technology choices different from what HPE has been working on.