REPORT REPRINT

MemVerge picks up big backers for its vision of a memory-driven universe

AUGUST 03 2020

By Tim Stammers

The startup has scored funding from Intel, NetApp, SK Hynix and other investors, and says its memory virtualization software will usher in an era of 'big memory computing' by accelerating the take-up of storage-class memories. Is this the beginning of the end of data storage as we know it?

Introduction

In 2019 MemVerge unveiled software designed to exploit the benefits of persistent or storage-class memories (SCMs). Earlier this year, the startup announced a funding round that introduced Intel Capital, Cisco Investments, NetApp and SK Hynix as new investors. MemVerge's software is now in an early-access program, ahead of full release planned for later this year.

MemVerge says the software will usher in an age of 'big memory computing,' in which many applications will run entirely in memory, and consequently enjoy boosted levels of performance and availability. Renamed this year as Memory Engine, the software currently only works with Intel's Optane Persistent Memory Modules (PMems, formerly known as DCPMMs) because those are the only mainstream devices currently presenting SCM as byte-level memory.

The company says it will support other such SCM-powered memory modules if and when they appear. The key feature of the software is that it eliminates the need to modify applications to use the Intel PMems in their fastest and most effective operating mode, while providing high-speed memory-based snapshot and replication data protection services and pooling Optane and DRAM capacity across clustered servers.

451 TAKE

MemVerge's software uses Optane as a memory in the PMem modules that Intel launched in 2019, and not as storage in the Optane NVMe drives that have been on the market since 2017. The two are very different products, and there has been a long-standing prediction that Optane PMems will enjoy stronger sales than Optane NVMe drives. Although the headline quality of MemVerge's software is that it reduces the need to modify applications to make the most of persistent memories, its other qualities as a memory virtualization platform are equally important. The size of the potential market for the software will be directly related to the sales of Intel's PMem devices, which are both a complement and an alternative to DRAM. If the PMems displaced just 15% of current spending on server DRAM, they would generate about $10bn in annual sales and create a healthy market for MemVerge's software. Although the devices are currently the only SCM modules on the market, Intel's former Optane development partner, Micron Technology, plans to develop a similar product, and other vendors have promised to bring new SCMs to market. MemVerge's predicted era of big memory will not arrive overnight, but it is likely to happen – and to significantly change the nature of data storage by moving primary data out of storage systems and into memory.

New funding and early-access program

A series B funding round announced by MemVerge in May takes the startup's total funding to date to $43.5m. Intel Capital was the lead investor for the new round. Alongside the new stakeholders, the round included previous investors Gaorong Capital, Glory Ventures, Jerusalem Venture Partners, LDV Partners, Lightspeed Venture Partners and Northern Light Venture Capital.

Intel and NetApp both contributed to a MemVerge marketing webinar held earlier this year, and the startup is pitching its software for use with data-intensive, performance-sensitive workloads such as SAP HANA, or analytics and ML applications such as RocksDB, Redis and TensorFlow. As yet, there are no production deployments, but about a dozen organizations have been trialing the software. Those include LinkedIn, JD.com and TenCent, all of which have publicly endorsed MemVerge's technology.

The 'big memory' prediction

In MemVerge's prediction, the era of big memory computing will see all applications running in memory that is abundant, persistent and highly available. The abundance and persistence will result from the higher physical density and lower cost of SCMs compared with conventional DRAM memory, as well as their persistence and nonvolatility. The high availability actually refers to the availability of applications, which will be boosted by the persistence of SCMs, which will allow for significantly faster application recoveries after server crashes or restarts. According to MemVerge, availability will also be boosted by the use of the high-speed memory-based snapshot and replication functions of its Memory Machine software.

This is a slightly different message from the one that MemVerge delivered when it first emerged from stealth in 2019. At that time, the company coined the term Memory Converged Infrastructure (MCI) to describe its software. Now MemVerge is putting less stress on the MCI label and is focusing more on its big memory forecast.

We think this makes sense, even though we endorsed MemVerge's use of the MCI label in 2019. We endorsed the label because the Memory Machine software creates a form of HCI that puts Optane to better use than other HCI products by using it as memory rather than just storage. However, the label also implies that MemVerge's software creates a platform that is an alternative to HCI, and will be used to host the same range of general workloads as HCI. Although that is a long-term prospect, in the near term the software will be used mostly (if not entirely) to host performance-sensitive, data-intensive workloads, which are rarely hosted on HCI.

Optane memory was developed jointly by Intel and its former partner Micron. This year Micron is set to launch the first of its own Optane-powered products, although Micron calls the memory 3D XPoint, and not Optane. Micron's first 3D XPoint device will be an NVMe drive, but Micron says that it will also ship a 3D XPoint-powered memory device. That would be similar to Intel's PMem device, and could be used with MemVerge's software. Separately from Micron, MemVerge says it is very confident that other suppliers will begin shipping memory modules based on new SCMs within the next two to three years.

Optane as persistent memory, not storage

451 Research, as well as Intel itself, believes that Optane PMems will eventually outsell Optane NVMe drives. This is because the memory modules exploit Optane's performance far better than the NVMe drives, and allow Optane to act as both a complement and a nonvolatile alternative to DRAM, rather than simply as a faster alternative to flash. Intel is not yet declaring Optane PMem sales numbers, but it says that it and its OEM partners have completed hundreds of PoCs, with over 80% resulting in production deployments. Some customers have skipped the PoC and moved directly to production, according to Intel.

PMem is a generic industry label for any persistent memory or SCM. Intel is also using PMem as the product name for its Optane-powered memory modules, which were originally named Data Center Persistent Memory Modules (DCPMMs). This year Intel dropped the DC from the name, and began referring to the devices as PMems. The devices are DIMM memory modules powered by nonvolatile Optane rather than by the volatile DRAM used in conventional DIMMs. 451 Research has previously referred to Optane PMems as nonvolatile DIMMs (NVDIMMs) because we felt that label was simpler to understand than DCPMM.

Optane PMem modules are faster than Optane NVMe drives because they transfer data at byte level to and from server CPUs over a DDR memory bus, unlike NVMe drives that transfer data in larger blocks (over a slower PCIe bus) using a latency-inducing NVMe software driver. The PMems are currently powered by the first generation of Optane memory. Judging by Intel's statements made this year, the devices will switch to second-generation Optane during 2021. That should boost their performance and capacity, and very likely reduce their per-GB prices.

MemVerge's software and Intel PMems

451 Research outlined MemVerge's Linux-based software when it was unveiled in 2019. The company says alpha testers of the product have seen dramatic performance improvements. One customer used the software to replace 50 servers running a machine learning application with a single eight-way server, while cutting the time taken to load data for training, according to MemVerge. An eight-way server can be fitted with up to 36TB of Optane PMem capacity.

The software works with bare-metal or virtualized servers; pools DRAM and Optane PMem memory across up to 128 RDMA-clustered servers; tiers data across DRAM, Optane PMem and conventional storage; and provides fast memory-based data replication, snapshot and restore functions. As stated above, it also eliminates the need to modify applications to use Optane PMems.

PMems operate in a choice of two modes: Memory mode and Application Direct (AppDirect) mode. Memory mode requires no modifications to applications, but doesn't exploit the nonvolatility of Optane memory. For many applications, that is a significant disadvantage – nonvolatility is one of Optane's biggest virtues. In contrast, AppDirect mode requires applications to be modified, but does allow Optane's nonvolatility to be exploited to heavily boost application availability.

As its name suggests, AppDirect also allows applications to use Optane capacity directly, and potentially more efficiently than Memory mode. Intel has said that, while Memory mode is likely to be more popular at first, it expects that as applications are modified to use AppDirect, the latter will become the more commonly used mode.

In 2019 the chip giant said applications that had already been modified to support AppDirect included SAP HANA, Aerospike Enterprise Edition, Apache Spark and Cloudera. While these out-of-the-gate modifications signaled application vendors' strong interest in using AppDirect mode, we do not know the depth of the modifications. Oracle told us in 2019 that it intends to fully exploit all of the benefits of Optane, and that it expects all of its database rivals to do the same. However, Oracle also said that doing so will require major changes to database architectures.

MemVerge says it will take years to fully modify applications, and predicts that some applications will never be modified. Even when applications have been modified, they will still not match the Memory Machine's ability to provide snapshots and replication services running in memory rather than storage, according to the company.

The Memory Machine's snapshots are far faster to create than conventional snapshots because their creation only requires data to be flushed from a DRAM cache into Optane PMem, according to MemVerge. The startup says this only takes microseconds, compared with the minutes during which even in-memory applications such as SAP HANA must be quiesced while conventional snapshots are created in storage. Similarly, MemVerge's memory-to-memory replication hugely outperforms storage-to-storage replication, according to the company.

As a technical note, unmodified applications can work with PMems in AppDirect mode, but only by accessing data in blocks rather than at byte level, which is slower not just because the blocks are larger, but also because it involves the use of software-based block drivers. Intel says the applications that it listed in 2019 as having been modified to use AppDirect mode all did so at byte level.

Competition

We are aware of only two other software products that allow unmodified applications to use Optane PMems in AppDirect mode. They are both aimed at the same applications as the MemVerge software. One is fellow startup Formulus Black's FORSA software, and the other is NetApp's MAX Data software.

Formulus Black's software supports AppDirect using block- rather than byte-level data transfers via a proprietary driver, allowing Optane PMem capacity to be presented as storage. It also includes a highly unusual data reduction mechanism that boosts performance by allowing applications to store more data in either Optane or DRAM, and by accelerating writes.

The Formulus Black software also pools Optane PMem and DRAM capacity across clusters of servers, similarly to MemVerge's Memory Machine. That is not the case for NetApp's MAX Data software, which works only on single servers, where it tiers data between Optane PMem and NetApp storage systems. NetApp says that MAX Data and MemVerge's Memory Machine are complementary products.

SWOT Analysis

STRENGTHS

MemVerge's technology allows the maximum benefits of Intel's Optane PMem to be realized using unmodified applications and provides other important functions.

WEAKNESSES

The company's software is as yet little known, and introduces an extra layer of complexity and vulnerability into operating environments.

OPPORTUNITIES

The analytics and machine learning markets are growing fast. MemVerge's software was designed to meet the needs of these and other data-intensive applications for faster data access and, hence, faster overall performance.

THREATS

The fate of MemVerge's software will be tied to that of Intel's Optane PMems until other suppliers ship SCM-powered memory modules. Take-up of the PMems could be too slow to drive sufficient demand for MemVerge's software.