



Enterprise Storage Newsletter

February 2021

Highlights

[Big Memory Storage](#)

[Micron Predictions for 2021](#)

[Flash or Cache? Deep Dive on SSD versus HDD \(plus cache\)
for Arrays \(Infinidat in the Hot Seat\)](#)

[Jeff Bezos/Amazon Timeline](#)

**Big Memory, Bigger
Application Performance**

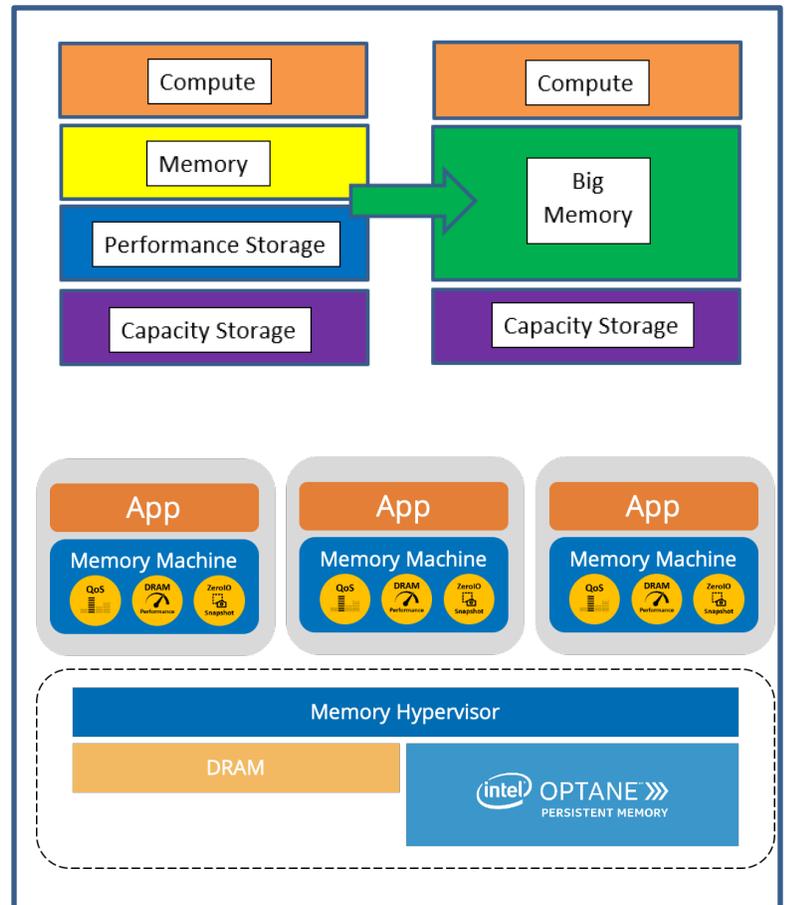


The fastest way to run any application has always been to run the application in dynamic random-access memory (DRAM) as much as possible. Several classes of applications have been built to take advantage of the growth in DRAM capacity of modern servers. These include in-memory databases such as [SAP HANA](#), [Aerospike](#), [Redis](#), and [Apache Ignite](#), as well as applications that run on top of these databases – think artificial intelligence, business intelligence, and similar workloads. Supported in-memory applications include 3D animation and rendering, in-memory big data analytics framework, machine learning, persistent key value database, in-memory key-value database, NoSQL database, and environment for statistical

computing. The importance of this market is underscored by the fact that the in-memory computing market is expected to nearly double between 2020 and 2025 from \$11.4 billion (US) to \$24.5B.

From the application perspective, the great thing about “big memory” servers is that applications can be written to run in a single virtual machine that has a huge number of cores and a gargantuan memory footprint. However, one of the inconvenient truths of big data has been that DRAM is one of the most expensive components in servers (especially and unsurprisingly in “big memory” servers). So how can you have big memory without the big memory price tag?

Turns out that there are a couple of reasonably affordable options to solve this problem. The obvious way is to use lower-cost storage in conjunction with a software program that presents the cheaper storage as memory. Today, that “cheaper storage” in most cases means [Intel’s Optane persistent memory](#) technology, first introduced in 2017. Optane combines the attributes of persistent storage (which DRAM doesn’t have) along with speeds that are roughly in-between NAND flash solid-state drives (SSDs) and DRAM. Intel provides two native modes to access the Optane memory: **App Direct Mode**, that allows developers to access the memory through an API, **and Memory Mode** that makes the Optane look just like (volatile) DRAM.



Third-party companies have also jumped into the deep end of the memory pool as well. [MemVerge](#) expands the amount of memory presented by including SSD storage pools. The MemVerge Memory Machine virtualizes DRAM and persistent memory to provide a transparent “tiered memory” structure. Memory Machine also addresses a few other issues that have been problematic for big memory applications, such as cloning databases over RDMA and providing snapshot functionality. These capabilities significantly simplify the fielding of big memory applications. Of course, MemVerge is not the only player in this market – there are others such as [OmniTier](#) and [Violin Systems](#). Each has a different approach to addressing the same concept – how to present more memory to applications (and improve their performance) without breaking the bank.

Micron Predictions – 2021



[Rajeeb Hazra](#), SVP of Emerging Products and Corporate Strategy, at [Micron](#) give his thoughts on predictions for this year, from [Storage Newsletter](#).

In 2021, the prevalence of remote work – even post-pandemic – will continue accelerating capabilities in the cloud. Companies will look to create preparedness for a new normal whether it be more IT solutions for a flexible workforce, larger data stores to fuel continued growth of online commerce, or resilient IT systems to address any future health care crises. This will drive unprecedented demand for agile IT infrastructure, multi-cloud solutions and pervasive connectivity to power edge-to-cloud use



cases. While we see great opportunity for memory and storage to fuel increasingly data-centric cloud services, we will also see a rise in data center operators evaluating disaggregated, composable systems to better scale for coming enterprise demands and data growth.

Boundaries between memory and storage will blur: 2021 is going to see AI-as-a-service become mainstream, intelligence migrate to the edge, and 5G come to life. This is going to propel fundamental changes in the way server systems are architected. Memory will extend into multiple infrastructure pools – and will become a shared resource. And the lines between storage and memory will blur. You'll no longer think "*DRAM for memory and NAND for storage*". Instead, faster NAND will create the ability to use it as memory, and applications will grow in their sophistication to utilize resourcing in innovative ways. In 2021, we'll also see enterprises seeking new kinds of solutions such as storage-class memory and memory virtualization to further unlock the value of AI and exploding volumes of data.

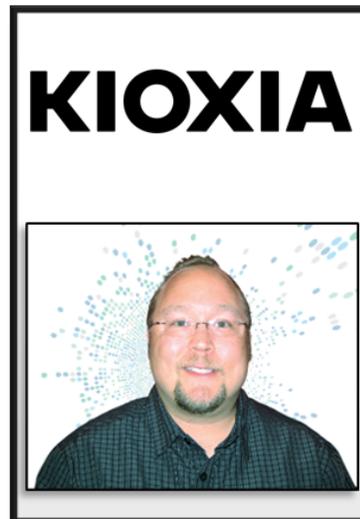
More pressure for an energy-efficient cloud: The move toward composable infrastructure will be critical in reducing over-provisioned resources, and thus, mitigating the rising environmental impact of IT. Information and communication technology is already predicted to use 20% of the world's electricity by 2030. As companies look to incorporate sustainability into business strategy and reduce Opex for compute-intensive workloads such as AI and high-performance computing,

we'll see escalating demand for energy-efficient architectures, enabled by composable infrastructure.

AI will become more accurate and more ubiquitous, and, we'll start to see it filling in more gaps for simple tasks where people would traditionally say *"Why would I ever use an AI algorithm for that?"* For example, grocery stores might tap AI-enabled cameras that periodically check to see if a shelf is empty and if so, alert a clerk to restock. In a post-Covid world, we'll see more businesses adopting AI for use cases like these to create these contactless experiences. We'll also see AI moving into infrastructure such as data centers and telecom base stations as neural network algorithms become more adept at workload and system error correction and recovery.

The rise of edge data centers: There are lots of startups that are focused on building edge data centers that look like transport containers that sit in metro areas to enable content – like your Hulu videos – to be closer to the consumption. We'll see the adoption of these edge data centers in the next few years, as enterprises and consumers look to tap massive amounts of data for insight and faster services closer to the source.

High-bandwidth solutions for high-compute at the edge are becoming a requirement. With fully autonomous solutions, the amount of compute performance needed for cars is reaching data center levels; in ADAS and autonomous driving, cars need hundreds of tera operations per second. This is some of the highest levels of performance in the industry today, rivaling what you will find in data centers.

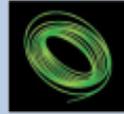


“Storage Architectures to Maximize the Performance of HPC Clusters”

Tues, Feb 23 9:00am

G2M
RESEARCH

Flash or Cache? Deep Dive on SSD versus HDD (plus cache) for Arrays



INFINIDAT

In storage systems (and in particular storage arrays), the conventional wisdom has been that flash storage in the form of solid-state drives (SSDs) are necessary for high performance, and NVMe[®] SSDs provide the highest performance of all storage media. However, a recent interview in [Blocks and Files](#) by [Infinidat](#) (started by a number of ex-[EMC](#) alumni) would like to be the start of a rewrite to this narrative. According to Kent Steinhardt (Field CTO at Infinidat), focusing on the media is looking in the wrong place. Steinhardt says the comparison of hard disk drive (HDD) vs SSD performance is like trying to “compare two different cars to determine which one will win a race based upon the horsepower of the engines. The greatest determining factor will be the skill and intelligence of the driver behind the wheel, and not the car itself.”

That Infinidat would take this position is not surprising given that they are a laggard in the move to all-flash arrays (AFAAs), and continue to focus on HDDs as their primary storage **capacity** technology which they claim they do because of the low price of HDDs vs SSDs. However, the architecture of Infinidat's arrays makes this argument somewhat hollow. The [Infinidata F6300](#), their flagship hardware product which has a usable capacity of up to 4PB and is capable of 1.3M (Read?) IOPS and 25GB/s throughput, contains over 3.0TB of memory and up to 368TB of “flash cache”. In this sense, Infinidat's products have the classical storage array architecture of the 2000s – a hierarchical storage approach where the movement of data across storage tiers is managed by software. Also interesting is the use of RDMA-based NVMe over Fabric (NVMe-oF™) for networking on both Ethernet and FC. Given that there are zero NVMe-based HDDs, that means a hefty NVMe to SCSI translation layer must sit between the network interface and the drives, which will increase latency for non-cached data.

And, as anyone who has been in the array business knows, storage performance claims should be taken with a grain of salt, especially with benchmarks as simple as IOPS and bandwidth. For instance, while Infinidat claims 25GB/s of storage throughput, the “pipes” going into their array are 25Gb/s Ethernet and/or 32Gb/s Fibre Channel (FC) connections, meaning that a compute host cannot get more than (roughly) 3GB/s of throughput unless link teaming is utilized. And, while 1.3M

IOPS (again, probably “Read” IOPS) is a big number, a cluster of Pure Storage FlashBlades has been [clocked at](#) 24M NFS IOPs. Finally, most vendors focusing on HDDs rely on the story that HDDs are significantly cheaper than SSDs on a \$/TB basis, which would be great if the storage media was the bulk of the cost of a storage array (it isn’t). Fixing “slow” (as in HDD performance) is never cheap, and Infinidat’s arguments are eerily reminiscent of those made by IBM about flash storage several years ago. That is, until IBM switched to flash, now [touting](#) NVMe-based AFAs as the right direction for enterprises to go...

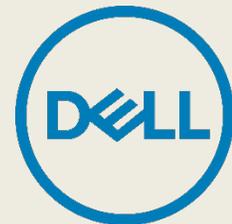
“Our Infinibox is so performant that there’s no need to use costly memory products like Optane’, said product head Yair Cohen.

That was without citing any competitors in particular but he probably meant Pure Storage and Dell EMC, who have both recently begun to offer Intel Optane in their arrays, and with whom the leadership at Infinidat seem obsessed.

Pure is a supplier that seems firmly ensconced in some enterprise sectors, such as banking, when it was so recently a startup too.

Meanwhile, Dell EMC has built its reputation on its Symmetrix/VMAX arrays, which were invented by Infinidat founder Yanai.”

Yann Serra, Computer Weekly



Upcoming 2021 Enterprise Storage Events

[International Conf on Storage Management & Data Protection](#), Feb 25-26, Japan

[Data Center Dynamics Conference](#), Feb 24-25, March 17-28, Virtual

[Developer World](#), February 17-19, Virtual

[Blueprint LDN](#), March 10-11, Virtual

[SNIA Webcast](#), The Ethics of Artificial Intelligence, March 16, Virtual

[GPU Technology Conference](#), April 12-16, Virtual

Jeff Bezos/Amazon Timeline

Jeff Bezos/Amazon Timeline

Bezos was born in NM. He loved computers, turned garage into lab, electrical Projects in the house.

High school valedictorian.

Princeton - *summa cum laude*, Computer Sci & Electrical Engineering

Amazon – Employees began developing software in his garage, 300 friends beta tested his site
 Launched July 16, 1995
 Sales of \$20k/wk in 2 mos in 46 countries.
 1997- Went public, outpacing competitors w/in 2 yrs.

Launch of Amazon video on demand & AWS cloud computing

Bezos cameo on Star Trek Beyond
 First Prime Air delivery using drones.
 Washington Post profitable.
 Bezos invites reporters to visit Blue Origin headquarters to hear about his vision for the company

Amazon announces Bezos will step down as CEO in 3rd quarter, stay as executive chairman of the board, & remain largest shareholder

Bezo & Mackenzie divorce. She gets 25% of Amazon stake & commits to donating half to charity. Amazon - \$280B in revenue & \$11.5B net profit.

NASA - Blue Origin & 12 other companies selected to collaborate on 19 tech projects to reach moon & Mars.

1964-1982 1982-1986 1986-1994 1994-1997 2000-2002 2006 2007 2011 2013 2015 2016 2017 2018 2019 2020 2021

Worked for Wall Street firms.
 1990 - youngest VP at DE Shaw.
 1993- Met & married MacKenzie, also worked at DE Shaw.
 1994- They move to Seattle to launch Amazon, an online bookstore.

Purchases Wash Post - \$250M, hires 100s of reporters, triples tech staff
 Amazon Prime Air – drones deliver packages up to 5lbs & w/in 10 miles
 Launch of Amazon Studios

Acquires Whole Foods - \$13.7B.
 Commits \$1B Amazon stock annually to fund Blue Origin
 Wash Post ad revenue \$100M

Launch of \$10B Bezos Earth Fund to combat climate change.
 Bezos commits \$100M to Feed America.

2000 - Bezos founds Blue Origin, aerospace company to lower cost of space travel.
 2002- Blue Origin moon lander, test flights to take tourists to space for a few minutes
 2002- Launch of Amazon Web Services

Amazon, Berkshire Hathaway, & JP Morgan announce plans for healthcare for employees. Raises min wage for employees to \$15/hour. 100M+ paid subscribers for Amazon Prime
 Launches Bezos Day One Fund to help homeless families & provide preschool in low-income communities, Bezos gives \$2B to fund project

Kindle Fire

Amazon Studio - first original feature film, Spike Lee's Chi-Raq

Kindle



“Today, we employ 1.3 million talented, dedicated people, serve hundreds of millions of customers and businesses, and are widely recognized as one of the most successful companies in the world.

*How did that happen? Invention. Invention is the root of our success. We’ve done crazy things together, and then made them normal. We pioneered customer reviews, 1-Click, personalized recommendations, Prime’s insanely-fast shipping, Just Walk Out shopping, the Climate Pledge, Kindle, Alexa, marketplace, infrastructure cloud computing, Career Choice, and much more. **If you get it right, a few years after a surprising invention, the new thing has become normal. People yawn. And that yawn is the greatest compliment an inventor can receive.**”*

Jeff Bezos, CEO, Amazon

Jeff Bezos will step down as CEO of Amazon in July of this year. His replacement, Andy Jassy, currently CEO of Amazon Web Services, has been with Amazon since 1997. The above, is quote from his letter to employees regarding his exit.

G2M Research Multi-Vendor Webinar Series

Our 2021 webinar schedule is ready! Click on any of the topics to get more information about that specific webinar. Interested in Sponsoring a webinar? Contact [G2M](#) for a prospectus.

Our January webinar “Can Your Server Handle the Size of Your SSDs” was sponsored by [Kioxia](#) (Matt Hallberg), [Lightbits](#) (Josh Goldenhar), and [Intel](#) (JonMichael Hands). [View the recording](#) and/or [download a PDF of the slides](#).



- Feb 23: [Storage Architectures to Maximize the Performance of HPC Clusters](#)
- March 23: [One Year after COVID-19: How Did Storage Architectures Perform for Biotech AI Modeling & What Can We Learn From This?](#)
- April 20: [The Race to be Relevant in Autonomous Vehicle Data Storage \(both On-Vehicle and Off-Vehicle\)](#)
- May 18: [Responsive and Efficient Storage Architectures for Social Media](#)
- June 15: [It's 2021 - Where Has NVMe-oF™ Progressed To?](#)
- July 13: [Computational Storage vs Virtualized Computation/Storage in the Datacenter: “And The Winner Is”?](#)
- Aug 17: [AI/ML Storage - Distributed vs Centralized Architectures](#)
- Sept 14: [Composable Infrastructure vs Hyper-Converged Infrastructure for Business Intelligence](#)
- Oct 12: [Cloud Service Providers: Is Public Cloud, Private Datacenter, or a Hybrid Model Right for You?](#)
- Nov 9: [The Radiometry Data Explosion: Can Storage Keep Pace?](#)
- Dec 14: [2021 Enterprise Storage Wrap-up Panel Discussion](#)

G2M
COMMUNICATIONS



Effective Marketing & Communications
with Quantifiable Results