

Sponsored

# How all-flash storage helps DUG Technology expand high-performance computing into new markets

VB Staff

January 6, 2021 4:10 AM

Enterprise

f t in



One of 256 tiles of the Murchison Widefield Array, located 1.5km from the core of the telescope.

Image Credit: Peter Wheeler, ICRAR

*Presented by Intel*

Energy exploration and high-performance computing (HPC) are made for each other. A typical project processes a petabyte of data (1,000 terabytes), which can quickly balloon 5x, 10x, or more, and often lasts a year or longer. Providing the supercomputing and storage muscle needed for massive oil and gas searches is a bedrock business for [DUG Technology](#), an international service provider with offices in Perth, Houston, London, and Kuala Lumpur.

Yet a couple of years ago, in the continued downturn in the energy industry, the company spotted a new opportunity: providing medical and genetic researchers, astrophysicists, and others with affordable, on-demand HPC services to handle analysis of large, complex data sets. “Traditional cloud providers are not focused on HPC,” says Phil Schwan, chief technology officer at DUG. “If you want to build Netflix, they’re terrific. But if you want to build HPC, you have to do a lot of things yourself. You need a lot of your own HPC expertise.”

So in 2019 DUG set out to expand its focus and change all that. The goal: continue to provide top service for existing customers while attracting different types of new users of HPC cloud services. Eighteen months later, the company is servicing business from a diverse range of HPC disciplines. Much credit goes to what Schwan calls a “technology leapfrog event” — replacing hard disk-based legacy storage with innovative, all-flash storage service from [Vast Data](#), powered by [Intel Optane](#) storage and memory.

The modernization has enabled the international technology firm to improve availability, manageability, and reliability, with “exponential” increases in performance of 50 petabytes of storage supporting its full-stack global HPC as a Service (HPCaaS) and other offerings. It’s also a great example of how technology innovation cascades through multiple businesses, enabling new opportunities at every step.


## Hard disks limit expansion

[HPC](#) uses a network of [supercomputer](#) nodes to tackle large computational problems too big and complex for standard computers. HPCaaS delivers this high-level processing capacity and related services to customers via the cloud. Both types can run the same scientific or big data analysis workloads; the latter brings organizations the added advantage of compute-intensive processing without the cost and expertise needed to create and run their own system.

That’s especially important for smaller companies and academic users who may lack the deep resources or everyday need. Says Schwan: “We see a lot of people trying to use cloud for their HPC compute. For somebody who doesn’t have a consistent 365-day constant workload, cloud makes a lot of sense. You can burst up as you need it. You can shrink down. It’s great.”

Begun in a shed in Western Australia in 2003 as DownUnder GeoSolutions, DUG Technology today owns and operates some of the largest computer systems in the world, totaling 30 petaflops. Over the years the company has built highly sophisticated and optimized disk-based storage systems to serve its 30-Pflop global capacity. But leaders knew that continued growth and competitive advantage would require a major storage upgrade. It soon became clear, however, that scaling and expanding the existing hard disk-based system was not a viable option.





Above: The 30 fluid-cooled petaflops that power DUG's HPC-a-Service (HPCaaS) are backed by 50 petabytes of fast storage.

Image Credit: DUG Technology

## Needed: Lots of fast petabytes

One big issue was scale: A project could easily swell to 10 petabytes and 50 copies; DUG runs 100–200 projects at any time. A project can last 12–18 months. Beyond scale, new target customers brought additional challenges. Many of their HPC applications needed even faster IOPs, more bandwidth, and greater reliability.

Latency and performance were paramount considerations. Says Schwan: “In many cases, the computations require random access to data sets much too large to fit in RAM. You’re being asked to do something that’s almost impossible without something like flash. If your peak latency is seven milliseconds, like it would be on a classic hard drive, physics very quickly limits the number of inputs/outputs-per second (IOPS) you can do. You can only divide seven milliseconds into one second so many times. Some of these workloads we simply couldn’t run on spinning discs, because the software is so insanely brutal in the way it uses the file system.”

He concludes: “Many of these new markets would have been very difficult for us to serve on our classic, pure HDD platform.”

Operational smoothness was another concern. While Schwan says DUG never had a serious problem with hard drive failures, “it was a constant thorn in our side. Because of the sheer quantity of storage we have, every week we were replacing drives. It’s a constant maintenance overhead that you just don’t have with the SSDs.”

In the end, software was also a big part of the decision to switch. Like many, DUG used [Lustre](#), a popular parallel HPC file system. To Schwan, an IO expert who helped develop the open source software, it was time for a new direction. “If we had to build an all-flash system using classic enterprise-grade flash in a Lustre-style config, it would have been

totally unaffordable, impractical, and not up to the job.”

## Wanted: An affordable flash option

So in early 2019, DUG began seeking world-class SSD suppliers. Rejecting several “monstrously priced” vendors, Schwan decided to evaluate Vast Data. The New York company is builder of the world’s first all-flash data and object storage platform and services.

Vast’s simplified, single-tier architecture lets companies consolidate complex storage stacks and unleash insights from reserves of data available in real time, explains Steve Paulhus, senior director, business development and alliances at Vast. Its Universal Storage architecture is powered by Intel Xeon Scalable Processors, Optane SSDs and [QLC 3D NAND SSDs](#). This innovative approach lets Vast deliver “exascale storage with NAS simplicity,” the company says. Since shipping its first system in 2018, Vast has become the fastest-selling storage company in history.

Schwan liked the economy, reliability, data reduction, and operational smoothness of Vast. He was especially impressed by how Vast layers QLC solid state drives with Optane. Doing so makes it possible to use inexpensive flash storage and affordably deliver data reduction.

Implementation began in mid-2019, and was done by year’s end. Today, a year into operation, DUG is still analyzing detailed ROI. But the benefits of the new storage platform are clear, both internally and to customers.

## “Two orders of magnitude” increase in I/O

HPC data sets and problems typically are huge. So quickly reading and writing files to physical storage is crucial to keep systems from idling or bogging down. With mechanical disks, processors must wait for the disk to rotate to the proper sector. The resulting slow input/output is a widespread bane.

“Every serious HPC center would agree: I/O is where people spend the most time, the most energy, and the most money,” says Schwan.” And it’s where failures most often occur. People build these beautiful, shiny, gleaming systems with tons of CPUs and GPUs, and they sit mostly idle because they can’t get data off the disc fast enough.”

With all-flash, DUG reports, Vast has enabled a “two orders of magnitude” increase in the number of IOPs that its systems can handle. And unlike HDD storage, operating at peak I/OPs does not degrade the experience of other users. “Now, the system can be doing half a million IOPs per second, and none of the other users even notice. It’s just totally invisible to them,” he adds.

Another benefit: All-flash helps DUG deliver a better end-user experience running varied, less optimized HPC software created by users and other third parties. Explains Schwan: “As we move into other sectors, the software is often written by scientists who are phenomenal in their domains, but they’re not — and shouldn’t have to be — computer scientists. If they don’t have the time or expertise to optimize I/O, then the system needs to soak up that abuse.”

Reliability also improved. With Lustre and hard disk storage, if any part of the file system goes down, any servers connected to the storage do too, Schwan says. “When that happens, the whole file system is offline, or at least a decent chunk of it.” In contrast, reliability at scale is much better with Vast flash. “It almost gets more reliable the bigger it is,



rather than less. That's a big draw.”

For a variety of reasons — notably cost and capabilities — DUG had never used a storage service provider. So the company was pleased with another benefit: better staff utilization. While a small DUG crew manages legacy storage as needed, Vast experts monitor and maintain the new all-flash storage service. “Our IT staff is freed from doing HDD fixes and updates. It was a very real benefit to us that somebody else would be maintaining the system going forward,” explains Schwan.

## New universe of customers

For existing oil and gas customers, Schwan says the change has been invisible. “We had already optimized our production chain to deliver high-quality service. We just arrived at it in a different way.” More noticeable is the new and expanded usage of DUG services by customers from new industries outside of energy. New and expanded engagements have produced dramatic results.

Take Curtin University in Perth, Australia. Teams there are working with the [Square Kilometre Array](#), a major multi-year, multi-national effort by governments and universities to build the largest radio telescope ever made. Scientists hope to peer back to the moments right after the Big Bang, nearly 13 billion years ago. As part of the effort, a smaller precursor telescope called the Murchison Widefield Array has been built in Western Australia by an international consortium led by Curtin. Using DUG services, researchers were able to process 450 hours of observations from the telescope in roughly a week.



Above: DUG's HPC as a Service accelerated processing of 450 hours data from the Murchison Widefield Array (MWA), a low frequency radio telescope and the first of four Square Kilometre Array (SKA) precursors to be completed. Photo supplied by ICRAR.

Image Credit: ICRAR

[Harry Perkins Institute of Medical Research](#) is using DUG technology and expertise to accelerate genomic and bioinformatic analysis of cancerous tumors and cells. Alistair Forrest, associate director of research for the Institute, says “Access to additional high-performance computing significantly increases our research capacity and allows us to expand the amount of data we can analyze.” High-performance computers and storage will help find disease mutations or analyze cell behaviors, for example. Ultimately, officials say the new capability will allow the institute to

achieve more breakthroughs sooner.

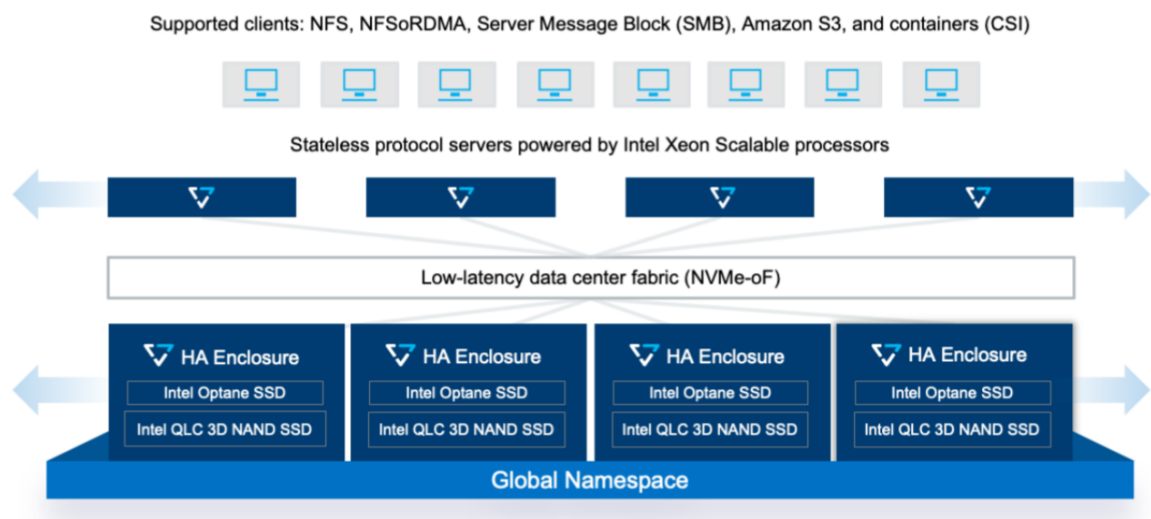
## Vast improvements, powered by Intel

Many of these performance benefits, new market opportunities, and expanded use of HPC are enabled by advances in the underlying Vast and Intel Optane technologies.

Take, for example, the challenge of providing very low-latency access to data — crucial for HPC. To solve the problem, Vast engineers took Intel Xeon Scalable Processors and a thin layer of Intel Optane, put a QLC NAND solid state drive behind it, then implemented the [NVME](#) storage communications interface over the fabrics. In essence, explains Paulhus, the Optane acts as a buffer that can write to the QLC drives in “a very friendly manner.” Reducing wear and tear in this way allows Vast to offer a rare 10-year warranty. “Server pooling” lets customers isolate users, improving privacy and security while maintaining quality of service.

Another innovation: so-called “protocol servers”. Instead of using a parallel file system like Lustre or an object store, the servers expose a Vast element store to clients via Network File System (NFS), Amazon S3, and other standard protocols. Doing away with client software has another benefit, notes DUG’s Schwan: It virtually eliminates the chance of a rogue client bringing down the file system, or otherwise negatively impacting other users. Full fault-tolerance also greatly improves reliability.

### Case Study | DUG Technology Modernizes with Exascale Flash Storage and Seizes New Opportunities Beyond Oil and Gas



### Logical diagram of the VAST Data Universal Storage solution

Going down a layer, Intel Optane and [3D XPoint memory](#) are crucial to the high performance and cost efficiency of the Vast platform. Explains Paulhus: “Because Optane is so fast, so low-latency, and — critically — non-volatile, we can use it almost like persistent RAM. Optane allows us to not only stage the data, but write optimally to the QLC (drive). We can do our erasure coding. We can do our similarity-based data reduction. Without Optane, we couldn’t do any of that.”

Schwan agrees, noting another benefit of pairing Optane and the QLC drive. “Having a layer to buffer the unorganized, random, chaotic data that’s coming from the applications and turning it into nice, large block-sized, very carefully managed I/O, and extending the life of the QLC — that’s only possible because of Optane,” he says.

# Benefits to old and new customers

Vaulting from a highly optimized but topped-out HDD storage to an all-flash service has brought DUG huge increases in performance, latency, and management reliability. It provides the modern technological foundation needed to deliver continued high levels of service to valued energy customers, while expanding into lucrative new HPC markets.

“There were customers in other verticals who we simply could not have served with the old [HDD] technology,” says Schwan. “We would have had to pick and choose our targets much more carefully, and limit ourselves to applications which had well-behaved I/O patterns.”

DUG recently finished doubling capacity in Houston, and will soon finish doing the same in Perth. Schwan says he’s looking forward to taking further advantage of Vast’s data reduction capabilities and is pleased with its all-flash services.

“I am a notorious tightwad,” he says. “It’s what’s allowed us to thrive in such a brutal oil and gas market for the last 10 years. We do relentless improvement of our cost structure. For not much more than we were already spending, we were able to make a huge leap forward. We saw how much improvement in our business it makes, and we were willing to take out the checkbook and make it happen. Anybody who knows me knows that’s my ultimate endorsement.”

---

*Sponsored articles are content produced by a company that is either paying for the post or has a business relationship with VentureBeat, and they’re always clearly marked. Content produced by our editorial team is never influenced by advertisers or sponsors in any way. For more information, contact [sales@venturebeat.com](mailto:sales@venturebeat.com).*