

Scientists Tap Cloud-Based Supercomputers in Search of Physics, Medical Breakthroughs



BY AARON RICADELA

November 29th, 2018

At Cardiff University in Wales, scientists are solving Einstein's equations that describe how black holes circle one another before slamming together, part of groundbreaking work that led to the first few detections of faint waves from massive collisions a billion years ago. The researchers are also creating models of what weak gravitational waves may look like when they reach miles-long detectors in the U.S. and Italy depending on the super-dense monsters' long-ago mass, spin and orientation in the sky.

Instead of running all of its intensive calculations on U.K. and European national supercomputers, Cardiff Professor Mark Hannam's group has been moving some programs to Oracle high-performance computers in the cloud, tapping computational resources that are starting to outstrip what's available at national and university facilities.

"Ten to twenty years ago, if you wanted to do cutting-edge computing you had to build your own cluster. The idea that a company was going to provide you with high-performance computing—none of these words made sense," says Hannam, a professor in the university's school of physics and astronomy. "I assumed before I started working with Oracle that the cloud wasn't useful for these types of calculations," which run across hundreds of processing cores and require extremely fast data transfers.

Researchers have a long tradition of building and wiring their own supercomputers, installing customized software to wring out more performance. “The applications where that wins out are becoming a smaller and smaller portion of scientific computing,” Hannam says. “The cloud is the next step, where you just have insane levels of resources. You could not and would not want to purchase those yourself.”

As businesses have moved their financial, HR, sales and other applications from computers in their own data centers to those leased as cloud services, scientific workloads have been more resistant to change. Among the factors that have throttled use of the cloud for high-performance scientific computing: requirements for fast connections among processors; performance slowdowns caused by some vendors’ use of virtual machines in their clouds; and university funding conventions that let researchers freely access computing time before applying for grants instead of having to pay for cloud computing credits.

That’s starting to change. Global spending on high-performance cloud computing rose 44 percent last year, to \$1.1 billion, according to market researcher Intersect360. That’s still only about 3 percent of the \$35.4 billion high-performance computing market, but cloud growth outpaced an overall 1.6 percent spending rise last year. Spending on high-performance cloud computing is forecast to reach \$3 billion by 2022.

“Just as enterprise computing is undergoing an on-premises-to-cloud shift, the same thing is happening in the research community, though they’re in some ways behind where industry is,” notes Phil Bates, a cloud architect at Oracle.

Public clouds are more likely to have the newest, most powerful CPU and graphics processing chips, as vendors continually add new hardware to their data centers. University supercomputers are bought and installed on longer procurement cycles and may have older technology. Demand for specialized hardware, including GPUs for running simulations and artificial intelligence applications, and support for low-latency networking are also pushing more high-performance work to cloud infrastructure.

Researchers with computationally intensive workloads already are getting comparable performance with Oracle Cloud Infrastructure as with the fastest on-premises supercomputers, which run batches of jobs that often wait in queues. Less time between running experiments and seeing results can let scientists change parameters to inquire further, says David Glowacki, a Royal Society research fellow at England’s Bristol University and founder of startup Simulitix Research.

Oracle Advances

Oracle, like other cloud computing companies, offers researchers high-performance computing, storage and graphics chips as cloud services, as well as fast network connectivity between CPU and GPU cores.

At the SC18 supercomputing conference in Dallas in November, Oracle said it's equipping its cloud with remote memory direct access (RDMA) technology, so that computers in a network can exchange data in memory faster without having to access each other's operating systems. That innovation improves the performance of software code whose work is distributed among many computers working in tandem.

In October, Oracle became the first cloud provider to announce it would offer 16 of Nvidia's powerful HGX-2 GPUs on a single machine via its service, at the same performance level as hardware Nvidia sells. Each GPU offers 2 petaflops of compute power (each performing 1 quadrillion floating point operations per second) and half a terabyte of graphics memory. Unlike some of its cloud competitors, Oracle doesn't insist that customers run their applications on a virtual machine, which can slow performance.

Such work is changing the way high-performance computing is bought and deployed, much as powerful clusters of low-cost PCs replaced supercomputers from Cray, IBM and Silicon Graphics in the late 1990s.

Oracle Executive Chairman and CTO Larry Ellison told financial analysts last month that computing and networking improvements the company is making to Oracle Cloud Infrastructure can run commercial workloads just as they power some particle-physics workloads at CERN, the European Organization for Nuclear Research.

Among those commercial applications: satellite broadband delivery, artificial intelligence in autonomous cars to detect other vehicles, and social media network modeling, says Taylor Newill, an Oracle high-performance computing product manager.

'Cultural and Social Change'

To be sure, not every high-performance workload is headed to the cloud. Some programs rely on specialized hardware developed expressly for them, making such jobs hard to move. Government funding agencies and universities may insist on data portability among vendors' clouds, a requirement that isn't always practicable. And funding agencies don't always let researchers apply grants to pay for cloud computing.

Moreover, universities have promoted their cutting-edge supercomputers to attract students and research fellows, a distinction that could dwindle as more research work moves off premises. "The cloud causes a cultural and social change in the way we do research," says Christopher Woods, a research software engineer fellow at Bristol University who is using Oracle Cloud Infrastructure on three projects.

At Bristol, a group studying nicotine receptors in the brain is trying to develop a smoking-cessation drug with fewer side effects than current treatments. The researchers ran 1,350 eight-hour molecular dynamics simulations on Oracle Cloud Infrastructure, distributed across 5,200 processing cores in 100 machines, benchmarking the results against the university's on-site supercomputer. On-demand access to the large numbers of Intel Skylake processors Oracle

offered let the group run its analysis in five days compared with an estimated three months on premises. The fast turnaround left time to quickly analyze results and perform a second computing run to test another drug, Woods says.

In another example, Bristol researchers working on the U.K.'s ADDomer project to develop vaccines against tropical viruses are using about 800,000 Oracle CPU hours and 15 terabytes of storage for high-resolution models of molecules generated by cryogenic electron microscopes. Those microscopes freeze proteins to produce three-dimensional video images that show how viruses work and genetic material is transported by RNA. "The cloud lets us construct bespoke scientific instruments to conduct our experiments," Woods says.

CERN at Cutting Edge

At CERN, physicists and engineers researching the basic components of matter using the Large Hadron Collider are pushing the industry's compute and storage capabilities. Those particles are accelerated to near light speed so scientists can study how they interact with one another.

CERN is investigating the use of Oracle cloud services to monitor data from the massive collider, which in 2012 helped physicists discover the elusive Higgs boson, a completely new particle related to a field that imbues other subatomic particles with their mass. That discovery earned two physics theorists the 2013 Nobel Prize for physics.

Six years after that discovery, CERN is laying the groundwork for discovery of new particles, as it prepares to bring online around 2026 the "High Luminosity Large Hadron Collider," capable of many more particle collisions per second. CERN's demand for processing power will be 50 to 100 times what it is today and require over 10 times more storage, says Eric Grancher, head of CERN's database group. The future collider will throw light on a range of mysteries, including theories about the existence of super-symmetric particles as well as invisible dark matter, which may hold galaxies together and constitute the bulk of the universe.

CERN now stores more than 50 petabytes (each equal to 1 million gigabytes) of data a year from its experiments, even after filtering that retains fewer than 1 out of every million events generated. The high-luminosity collider could boost total storage needs to the order of exabytes (each equal to 1 billion gigabytes) annually.

Those swelling storage requirements far outstrip the increases expected by the tech industry's normal development, and CERN is investigating the cloud as one possible way to shore up its data center and the global computer grid it underpins. "This is a major challenge we are working on," Grancher says. "It will be very interesting to see how things develop and whether cloud could become cheaper than on-premises computing for these types of workloads."

Aaron Ricadela is a director of strategic communications at Oracle.

This content was created and provided by the sponsor. Neither Dow Jones nor the news or advertising departments of the Wall Street Journal were involved in the creation or review of this

content.

Copyright 2018 Dow Jones & Company, Inc. All Rights Reserved

This copy is for your personal, non-commercial use only. Distribution and use of this material are governed by our Subscriber Agreement and by copyright law. For non-personal use or to order multiple copies, please contact Dow Jones Reprints at 1-800-843-0008 or visit www.djreprints.com.