share:

How Cloud HPC is reducing time to insights in preclinical research

by Fierce Custom Publishing

Published by
FiercePharma

Preclinical research is an unpredictable, fast-moving field. New insights quickly open up or shut down avenues of research, forcing teams to rethink their experiments. Computing systems have historically been less adaptable, but that is changing with high performance computing (HPC) in the cloud. Using on-demand, pay-as-you-go cloud HPC, preclinical teams are achieving research goals faster than ever.

HPC refers to the aggregation of computing power. By having multiple processors working in parallel, HPC delivers more power than a desktop computer. Access to such power has underpinned advances in genomics, molecular modeling, computational chemistry and other areas of life science research. HPC enables researchers to manage the 50 terabytes of data Illumina's HiSeq X Ten can generate each week, and equips Big Pharma firms to virtually screen millions of compounds.

The downside is HPC clusters are expensive to acquire, configure and maintain. Procuring technology is a long, costly and complex process. It takes months and many millions of dollars. Configuring processors to effectively run researchers' computations takes more time and expertise. Once installed, technical problems and computing-capacity bottlenecks lead to downtime and distractions, slowing research. These characteristics undermine on-premise HPC in preclinical research.

"On-premise HPC isn't as agile as our researchers want to be. Because research collaborations can start up spontaneously, the result can be demand for IT that wasn't forecast," said Lance Smith, associate director of IT at Celgene.

The IT sector responded to these challenges with cloud HPC. This paper discusses the benefits of deploying preclinical workloads in the cloud.

Cutting time to science

Cloud HPC shifts the burden of acquiring, configuring and maintaining computing equipment from preclinical researchers to service providers. Cloud providers such as Amazon Web Services (AWS) employ compilers, schedulers, distributed network file systems and other tools familiar to users of on-premise HPC clusters.

As such, the cloud frees scientists from IT administration, allowing them to focus on research endpoints and optimal outcomes. Users can build customized HPC clusters tailored and sized to their research goals.

Cloud HPC providers have created packages optimized for different use cases. These packages, known as "instance types," have different computing, graphics, memory and storage capabilities. As molecular modeling is demanding on the visual capabilities of a HPC system, its instance has multiple, top-end graphics processing units. Genome sequencing, in contrast, needs highly parallel compute clusters.

Procuring such customized setups is simple and fast. Spinning up clusters takes minutes. This is important to IT teams and preclinical research groups alike.

"Everybody wants to … reduce the time for deploying server environments," said Russell Towell, senior solutions specialist at Bristol-Myers Squibb.

Bristol-Myers shortened deployment times by working with AWS to add to its computing capacity. The system uses a dedicated, encrypted VPN tunnel and an Amazon Virtual Private Cloud separate from AWS' public customers. Server environments are built, secured and qualified by Bristol-Myers locally before being loaded into the cloud. These precautions ensure the system meets Bristol-Myers' standards.

Part of the appeal of cloud HPC to Bristol-Myers and others is the elimination of capital expenditure. Researchers only pay for the compute and storage capacity they use. This flexible, pay-as-you-go structure has enabled research groups that could not afford on-premise HPC to access compute capacity. For larger companies, the cloud has changed the economics of expanding beyond existing on-premise HPC clusters that are at peak capacity.

Accelerating toward insights

When Novartis wanted to virtually screen 10 million compounds against a cancer target, it calculated it needed sustained access to 50,000 compute cores--capacity that would cost $40 million to install. The firm had on-premise HPC, but it was running at full capacity. Even if Novartis suspended all other uses of the HPC cluster, it would not have had enough cores.

"We were absolutely dead in the water, stuck and didn't know what do next," said Steve Litster, global lead for high performance and scientific computing at the Novartis Institutes for Biomedical Research. "We had to create a system that was fast, extremely secure, inexpensive and easy to use."

Litster contacted AWS to see if it could create a

system to handle the work. The resulting cloud system distributed 10,600 spot instances across four connected zones, giving Novartis access to approximately 87,000 compute cores.

Novartis calculated the project equated to 39 years of computational chemistry. The cloud HPC performed the screening in nine hours at a cost of $4,232. "It really was amazing," Litster said. "Most importantly, we actually identified three promising compounds."

Cloud HPC's flexibility--and the cost at which it is available--has implications for how preclinical researchers approach experiments. If Novartis had been limited to on-premise HPC, the screening program may never have happened and the three promising compounds would never have been identified. This is an example of how cloud HPC is changing the economics of drug discovery.

"When you have the capability to bring low-cost experimentation to the scientific discovery cycle, we believe there is going to be a lot more science that gets done," said Mark Johnston, director of global business development, healthcare and life sciences at AWS. "Companies won't be so cautious or conservative with how they go about resourcing experimentation because we've effectively changed the economics and the dynamics of it."

Users of cloud HPC move at whatever pace is most appropriate for their experiments. If a research team needs to complete a task very quickly, for example, because related projects cannot advance until it is done, it scales up its cloud HPC to accelerate progress. Such scalability shortens time to insight.

"By spinning up a few hundred nodes on AWS and getting results in less than a day, our scientific researchers have a lot more freedom to ask questions that weren't even possible before," Celgene's Smith said.

The cloud also optimizes resource allocation for less time-sensitive projects. Scaling down cloud HPC optimizes the balance between cost and speed.

Future-proofed computing

The flexibility of cloud HPC has long-term implications. Researchers' computing needs will change in the years to come. The amount of data generated in preclinical research is growing and will continue to do so.

Expanding on-premise HPC to manage these data is a major undertaking. Users of cloud HPC, by contrast, can automatically scale up clusters to meet their growing compute needs. The adaptability and scalability of cloud HPC ensures researchers will always have the right cluster for the job.

This promise extends to the technology used in cloud HPC. On-premise HPC systems can be dated by the time they are installed. Cloud HPC frees researchers from worries about technology obsolescence. The burden of maintaining leading-edge technology falls on service providers, which, in the competitive cloud sector, need to offer the most advanced computing components to succeed.

In this regard, as in many others, cloud HPC shifts technical pressures from researchers to service providers. The upshot is users of cloud HPC go to work each day knowing they have what they need: scalable, leading-edge compute capacity tailored to cutting the time it takes to achieve their research endpoints.

These endpoints are far from abstract goals. They are staging posts on route to therapeutics that will address major unmet medical needs and in doing so touch the lives of billions of people. Cloud HPC will be with researchers at each step in the process. Drug development will remain the toughest of tasks. But with cloud HPC researchers can strike compute capacity from their list of challenges.
