Powering Breakthroughs in Life Sciences

High Performance Computing on AWS Powered by Intel® Xeon® Processors Accelerates Research in Genomics, Clinical Research, and More

3 The Importance of HPC for Life Sciences Research

4 How HPC on AWS Works

5 Genomics Data Processing

- How HPC on AWS benefits genomics data processing
- Rady Children's Hospital accelerates diagnoses with HPC on AWS

6 Clinical Trial Simulation

- How HPC on AWS benefits clinical trial simulation
- Bristol-Myers Squibb runs *in silico* simulations using HPC on AWS

7 Computational Chemistry

- How HPC on AWS benefits computational chemistry
- OpenEye adopts HPC on AWS

8 Deep Learning

- How HPC on AWS benefits deep learning
- GE Healthcare and UCSF use HPC on AWS for AI research

9 The Nuts and Bolts of AWS

10 How HPC on AWS Benefits Research

11 Getting Started with HPC on AWS

The Importance of HPC for Life Sciences Research

The need for high performance computing (HPC) is growing across a wide range of use cases in life sciences, such as genomics data processing, clinical trial simulation, computational chemistry, and deep learning.

As HPC workloads increase in number and complexity, on-premises infrastructures are coming under considerable strain. To allow scientists to continue their research work at full speed, organizations are under constant pressure to make sizable investments in on-premises hardware to ensure a high level of performance and sufficient capacity—an exercise that has to be repeated all too often to keep up with increasing demand. Cloud-based HPC on Amazon Web Services (AWS) offers an alternative by offering researchers access to the latest technologies and virtually unlimited HPC infrastructure capacity. Harnessing the power of Intel® Xeon® processor technology, AWS helps researchers get meaningful results faster, collaborate securely with colleagues anywhere in the world, and ask questions they couldn't ask before.

How HPC on AWS Works

Think of AWS as an extension of your current data center—without the capital expense of adding new hardware. Now you can have essentially unlimited amounts of computing power and data storage when you need them.

Ready, set, go! Forget waiting for procurement cycles—AWS is ready when you are. Launch new workloads right away and easily scale resources to meet the needs of your research workloads, not the other way around.

Stay current. AWS leverages the power of the latest Intel processors. You always get the benefit of technology advances as soon as they hit the market.

Watch your budget. AWS offers a range of pricing options. Pay only for the resources you use, with little or no upfront investment.

As computing technology advances, the performance and value of HPC solutions on AWS advances in tandem. For example, in November 2017, AWS introduced the EC2 C5 instance type based on the latest Intel® Xeon® Platinum 8000-series processors. C5 instances are 24 percent faster and up to 15 percent cheaper than C4 instances. This computing advancement helps life sciences companies achieve faster time to results and conduct more experiments, while reducing research costs.

Life Sciences Compliance in the Cloud

Run your own HPC applications or choose from a range of commonly used and verified applications from our trusted AWS ecosystem partners.

Q: What are Good Laboratory, Clinical, and Manufacturing Practices (GxP)?
A: GxP refers to the regulations and guidelines applicable to life sciences organizations that make food and medical products such as drugs, medical devices, and medical software applications. The overall intent of GxP requirements is to ensure that food and medical products are safe for consumers and to ensure the integrity of data used to make product-related safety decisions.

Q: (GxP)? How can AWS help us build a GxP environment?
A: By leveraging the AWS cloud to enable your GxP environment, you can create templates that allow you to use your infrastructure throughout your organization with a high degree of consistency. AWS also gives you deep control over who can affect elements of your infrastructure software—including when, where, and how they do it.

Q: What are the advantages of using AWS when it comes to testing?
A: When you build and validate applications using AWS, you are using infrastructure software products instead of physical hardware. This means you can repeatedly test and monitor your environment for compliance on a continuous basis, rather than rely on manual, point-in-time activities you might be using for services built with on-premises infrastructure.

Q: How does AWS help with audits?
A: When you deploy the AWS software infrastructure, you can also use AWS tools to automatically log a wide range of activities in your environment, including how the infrastructure is deployed, and how it is accessed and configured. This improves traceability in your environment, making it easier to support audit requests.

Q: How can I learn more about building validated GxP systems in the cloud?
A: AWS has a whitepaper, developed with contributions from international compliance consulting firm Lachman Consultants, on building validated systems in the cloud. **Download your copy now.**

Genomics Data Processing

An increased prevalence of chronic illnesses such as cardiovascular disease, cancer, and diabetes, as well as an aging population, is straining the global healthcare system. One promising avenue lies in precision medicine methods, which incorporate patients' genomic variants into their mainstream care with the hope of providing more targeted and effective treatments.

As the use of genomic-based therapies and the number of genomes being sequenced continues to grow, the limited capacity of in-house infrastructures becomes a bottleneck in clinical research environments. As jobs take longer and longer to run, the pace of potentially life-saving research slows down. Researchers scale back their studies to avoid overloading the platform, which hampers innovation.

HPC on AWS offers a way to avoid these bottlenecks. By moving genomics workloads to the AWS cloud, researchers can process, analyze, and interpret data at ever-accelerating speeds using AWS compute instances based on powerful Intel® Xeon® processors. The ability to add virtually unlimited compute resources when they are needed—without long procurement cycles—frees scientists to test more hypotheses, even highly speculative ones. HPC on AWS gives organizations a viable, cost-effective way to pursue research at full speed and contribute to better outcomes using precision medicine.

How HPC on AWS benefits genomics data processing

- Reduces time and money spent on infrastructure and reinvests these resources to increase the pace of discovery in genomics research.
- Derives data insights faster using a genomics analysis pipeline powered by the cloud.
- Facilitates global collaboration while ensuring data security and enabling compliance with HIPAA, GDPR, and other regulations.

Rady Children's Hospital accelerates diagnoses with HPC on AWS

Rady Children's Institute Genomic Medicine

Rady Children's Institute for Genomic Medicine uses Fabric Genomics' software running on Intel Xeon processors on AWS to identify disease-causing variants in a child's genome.

Read the full case study.

"Using Fabric Genomics' interpretation software on AWS, we can interpret an entire genome's variant set within minutes, whereas other platforms might take hours. We distill millions of data points into a few distinct markers for diseases in children."

Shareef Nahas, Ph.D., Senior Director of Clinical Laboratory Operations, Rady Children's Institute for Genomic Medicine

Clinical Trial Simulation

Launching a new drug can cost up to US$4.6B and take up to 12 years. Late-stage clinical trials currently represent 40-50 percent of this time and 60-65 percent of the cost of drug development.[i]

Pharmaceutical enterprises are increasingly turning to clinical trial simulation to help decrease the time and investment required to bring a new drug to market. This *in silico* drug modeling and simulation helps provide crucial scientific insight and, in some cases, eliminates the need for clinical trials in certain populations.[ii]

However, simulations strain on-premises infrastructure, causing researchers to reduce the number of simulations and creating processing queues that delay other research work. Many pharmaceutical companies are eliminating these constraints by migrating *in silico* trial workloads to Intel® Xeon® processors running on AWS. These simulated trials help to focus human trials, in many cases shortening the study period and reducing the required number of human subjects.

How HPC on AWS benefits clinical trial simulation

- Accelerates the pace of drug discovery without incurring significant capital expense.
- Optimizes design of human clinical trials, reducing the impact on trial participants.
- Shortens or even eliminates processing queues, minimizing the impact of simulation runs on other researchers.

Bristol-Myers Squibb runs *in silico* simulations using HPC on AWS

Bristol-Myers Squibb

Scientists at Bristol-Myers Squibb perform on-demand clinical trial simulations on AWS via a secure, self-provisioning research portal.

Read the full case study.

"Because of the compute capacity of AWS, our scientists can now run thousands of clinical trials to prepare optimized designs. As a result, we're using fewer subjects in these trials, optimizing the dosing levels and safety, and fewer blood samples from pediatric patients."

Russell Towell, Senior Solutions Specialist, Bristol-Myers Squibb

Computational Chemistry

Computational chemistry is a branch of chemistry that uses computer simulation to assist in solving chemical problems. Researchers use computational chemistry methods to better understand chemical interactions and experimental data and explore reaction mechanisms that are difficult to study experimentally.

The trend in computational chemistry is to model larger and more complex molecular systems. As an example, scientists working in drug discovery rely on heavyweight applications such as Schrödinger, GOLD, and Maestro[iii], which dramatically increase the demand for compute resources. In-house HPC infrastructures are often not up to the task, delaying simulation runs and slowing the pace and scope of drug discovery.

One promising solution is to migrate compute-intensive modeling applications to HPC on AWS. Researchers can bring up virtually unlimited numbers of EC2 instances powered by Intel® Xeon® processors and achieve faster computational speeds using Intel Advanced Vector Extensions 512 and the Intel Math Kernel Library.

How HPC on AWS benefits computational chemistry

- Allows researchers to effectively rank order the most likely drug candidates and eliminate less-promising compounds
- Brings life-saving drugs to market faster by taking advantage of the scalability of the cloud.
- Frees up on-premises resources to support other research and business needs.

OpenEye adopts HPC on AWS

OpenEye's Orion software, a drug-design platform running on Intel Xeon processors on AWS, helps chemists at pharmaceutical firms create, share, model, calculate, visualize, analyze, and organize chemical collections of different sizes and complexities.

Read the full case study.

"Our customers can deploy 100,000 cores in a few days, instead of the six months it would typically take. They can think about doing new and exciting things instead of thinking 'one day we could possibly do that.'"

Craig Bruce, Head of Infrastructure, OpenEye Scientific Software

Deep Learning

Artificial intelligence (AI) is coming to life sciences: Forty-four percent of life sciences professionals are using or experimenting with AI today.[iv] Scientists and researchers are particularly interested in the AI branch known as deep learning, which analyzes enormous datasets to find subtle patterns that elude human detection and learns how to better identify similar patterns in the future. This powerful AI technology holds enormous promise for use cases such as drug discovery, diagnosis of multifactorial diseases, and medical imaging.

However, most on-premises data centers today lack the compute and storage resources necessary to take advantage of deep learning. AI applications such as BigDL and Cortex push traditional hardware to the limit, forcing IT managers to consider costly upgrades that only solve the problem for a short time—often just months.

Early adopters can gain an advantage by turning to HPC solutions on AWS to accelerate research into deep learning and other AI disciplines. AWS and Intel have optimized deep learning engines using C5 instances powered by Intel® Xeon® scalable processors and the Intel Math Kernel Library. The result is a performance increase of more than 100x for machine learning inference.

How HPC on AWS benefits deep learning

- Improves the quality of predictions by increasing the amount of data analyzed.[v]
- Discovers promising drug candidates by finding subtle patterns in existing datasets from past testing.
- Enables visualizations of data to help physicians deliver more personalized care.

GE Healthcare and UCSF use HPC on AWS for AI research

GE Healthcare

cdhi

GE Healthcare and the Center for Digital Health Innovation at the University of California, San Francisco, have developed machine learning algorithms that can distinguish normal scans from those requiring immediate intervention. The research work made extensive use of HPC on AWS.

Learn more.

"There are many conditions we're working on with X-rays and deep learning that are incredibly compelling and very much attached to saving lives."

Keith Bigelow, General Manager of Analytics at GE Healthcare

The Nuts and Bolts of AWS

Researchers usually lack formal training in information technology, although many have developed a working understanding of IT infrastructure out of necessity. One of the key benefits of AWS is that it is simple to understand and use, allowing scientists to focus on science instead of the technology that powers the science.

At the same time, IT managers in life sciences organizations need to have a working knowledge of the AWS infrastructure to make informed decisions about moving workloads to AWS. This diagram gives a conceptual overview of the AWS architecture. For a more complete explanation of the technology inside HPC on AWS, see "HPC Capabilities on AWS."

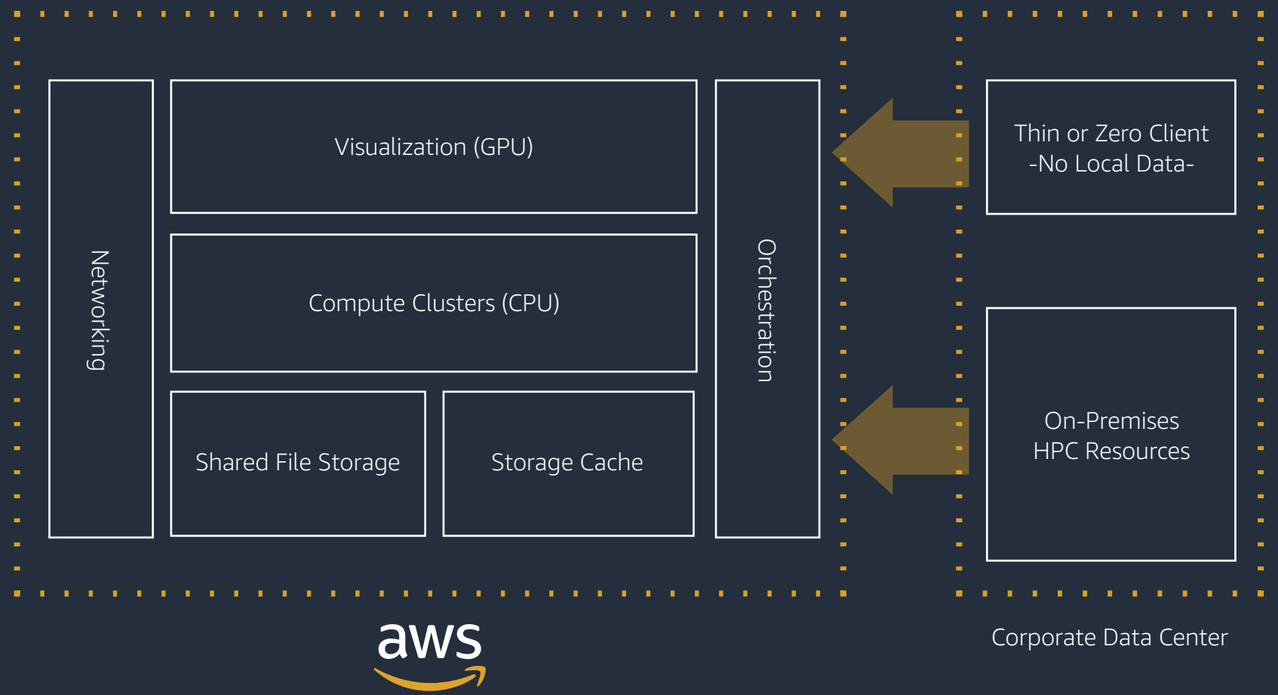

Annotations

VISUALIZATION: Accelerate graphics performance with GPU-powered instances and stream Windows graphics with Amazon AppStream 2.0, Amazon WorkSpaces, and NICE DCV.

COMPUTE: Tailor resources to match your workloads by selecting instances optimized for high-performance, memory-intensive, high I/O, or general applications.

STORAGE: Select from shared file system, block, and object storage options.

NETWORKING: Run highly parallel and high-throughput computing applications on the AWS low-latency, high-throughput network.

ORCHESTRATION: Take advantage of AWS's fully managed services or choose to configure and maintain cloud-native clusters yourself using AWS job schedulers.

How HPC on AWS Benefits Research

Transcend existing limitations. Ask more complicated questions—and more of them—without worrying about computational resources. Investigate corner cases and out-of-the-box hypotheses that exceed your on-premises capacity.

Get answers more quickly to accelerate research. Launch workloads in seconds and scale resources just as fast. Shorten time to answers, or publication, and gain a competitive advantage.

Collaborate effectively and securely. Store and share large datasets across geographies and institutions. Maintain control of your data assets and intellectual property using AWS identity, authentication, and access control services.

Share information efficiently and seamlessly. Create a single source of truth—one copy of data used by all stakeholders, internal and external. Avoid the need to move data to the researcher or maintain local copies of information that compromise data integrity.

Getting Started with HPC on AWS

Getting started with AWS is easy. You can set up an account with just a few clicks without any up-front commitment.

When you choose AWS, you gain access to Intel® Xeon® processor technologies and a full spectrum of storage, computing, and analytics.

AWS enables you to share and collaborate efficiently with team members across the globe without compromising on security.

AWS has a broad portfolio of life sciences partners who can help provide you with solutions and consulting that are specifically tailored for your research goals.

To learn more about how AWS and Intel are working together to meet the needs of life science and healthcare organizations, visit www.amazon.com/hpc or www.aws.amazon.com/intel.

Start experimenting with AWS today with a sample project or tutorial, gain deeper insight through whitepapers and videos, or find a partner to get hands-on guidance. Try it now.

i https://www.clinicalleader.com/doc/the-role-of-modeling-simulation-in-clinical-trials-0001
ii Ibid.
iii "Software based approaches for drug designing and development: A systematic review on commonly used software and its applications," Bulletin of Faculty of Pharmacy, Cairo University.
iv https://insidehpc.com/2018/01/survey-shows-ai-revolutionize-life-sciences/
v http://timdettmers.com/2017/08/31/deep-learning-research-directions/