Enabling Breakthroughs in Life Sciences

Lenovo TruScale High-Performance Computing as a Service (HPCaaS)

Introduction

More than ever, Life Sciences organizations must shorten discovery pipelines, enhance scientific collaboration, deal with data explosion, integrate data and applications faster, and respond to rapidly changing requirements. With a shortage of qualified lab techs in some regions compounding the problem, biotech groups are looking to leverage high-performance infrastructure to address these challenges. "High-performance" infrastructure refers to compute resources dedicated to running computationally intensive workloads like bioinformatics.

This whitepaper discusses these infrastructure challenges and the limitations of traditional on-premises high-performance computing (HPC) infrastructure and public clouds. It then outlines how new, more agile service-based models can give life sciences organizations the flexibility to leverage new opportunities and accelerate their discovery and research journey to the next level.

The remarkable impact of computing and genomics

In 2003, mapping the human genome cost about $3 billion. By 2019, it was less than $1,000. Within a decade or even sooner, the cost could be less than **$100**[1]. Consequently, human genome data analysis for large populations is feasible but, in the aggregate, will require tens of exabytes of storage and trillions of core hours.

Front and center to this "Bio-Revolution" are proven approaches to genomic analysis, such as next-generation sequencing (NGS). Figure 1, based on an NHGRI (a division of the NIH) framework, shows how NGS is transforming healthcare by facilitating the development of targeted drugs and helping to deliver personalized healthcare.

Figure 1: NHGRI Framework

NGS is vital to get better insights into the root causes of diseases, find new biomarkers of particular diseases, explore new drug targets, and individualize treatments based on a person's genotype. However, deploying NGS's information technology (IT) infrastructure is challenging.

IT infrastructure challenges and requirements for NGS

There are many challenges (Figure 2) in deploying an NGS high-performance infrastructure:

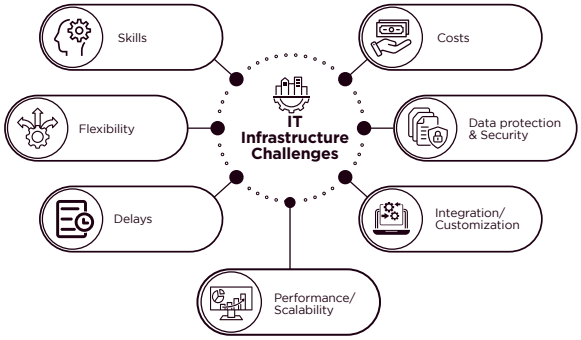

Figure 2: IT infrastructure challenges for life sciences

Costs: The capital cost of acquiring a traditional on-premises high-performance infrastructure can be prohibitively high, especially for smaller organizations or academic institutions. Additionally, these systems may have lower and sub-optimal utilization under normal operations since they are sized to handle peak loads. These systems also require considerable electricity, backup power, data center space, and cooling, which can add even more costs.

Skills: Operating an on-site high-performance infrastructure is complex and time-consuming. It needs a highly skilled team of technical professionals, which may be challenging to recruit and retain.

Performance/Scalability: Scaling out population genomics productions quickly depends on high-performance clusters, hardware accelerators, faster storage, and proximity to data. In addition, each step of the NGS workflow must be optimized for performance.

Flexibility: Applications developed for one high-performance cluster will likely not work optimally on another system. Migrating the entire NGS environment onto a new high-performance infrastructure can take weeks or months.

Delays: High-performance infrastructure is in high demand across industries, not just biotech, and is complex to install. The procurement, delivery, and deployment time can slow down the pace of implementation.

Integration/Customization: As a rapidly evolving field, life sciences customers must customize and integrate newer solutions, such as artificial intelligence and deep learning, into their existing workflows.

Data Protection and Security: Life sciences organizations deal with sensitive personal information in highly regulated environments and must ensure this data is protected and secure.

In the future, these infrastructure challenges will increase in complexity as life sciences organizations grapple with several thorny scientific challenges (Figure 3). The shortening of discovery pipelines, greater emphasis on collaboration, and the need to conduct more research with fewer scientists means life sciences organizations must accelerate time to results while managing, processing, and analyzing higher volumes of data with greater efficiency.

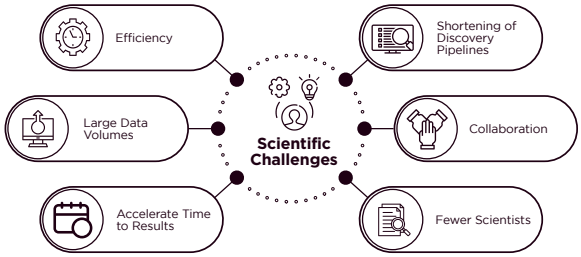

Figure 3: Scientific challenges for life sciences organizations

Researchers need flexible, scalable environments that rapidly process tremendous amounts of data to avoid these obstacles and speed up innovation. They must also collaborate and share large data sets with upstream and downstream partners, and they need an affordable infrastructure that supports automation to simplify data aggregation, assimilation, and management. High-Performance Computing as a Service (HPCaaS) provides this infrastructure to overcome many IT and scientific challenges for life sciences organizations.

Overcoming life sciences infrastructure challenges with "as a Service" models

New lab equipment, like next-generation sequencers, cryogenic electron microscopy (cryo-EM), and single-cell sequencing technologies, produce multiple terabytes of data per run that must be processed, stored, analyzed, and compared against large genomic databases.

Increasingly, research and clinical applications involve rich data sets, including MRIs, genome data, and ultrasound imaging. In addition, data from medical devices are increasingly attached to patient records or clinical experiments, driving the need to store and process ever-larger data sets with better efficiency while maintaining patient privacy. So, the scale (and costs) of computing in all life sciences disciplines is so colossal that even the largest life sciences organization may be unable to install and operate it without using some form of "as a Service" or cloud-like model.

With HPC as a Service (HPCaaS), users can acquire and access a range of bare-metal HPC clusters (dedicated single-tenant physical servers with no operating system overheads) from a central pool in a hosted data center (Figure 4 – left). They can build resource-intensive applications with the option to expand their footprint in the future if required. In addition to raw computing power, HPCaaS (Figure 4 – right) typically includes software for cluster management, workload optimization, security, and highly skilled people resources to operate the HPC clusters and ensure they are fully optimized to deliver maximum performance for life sciences applications.

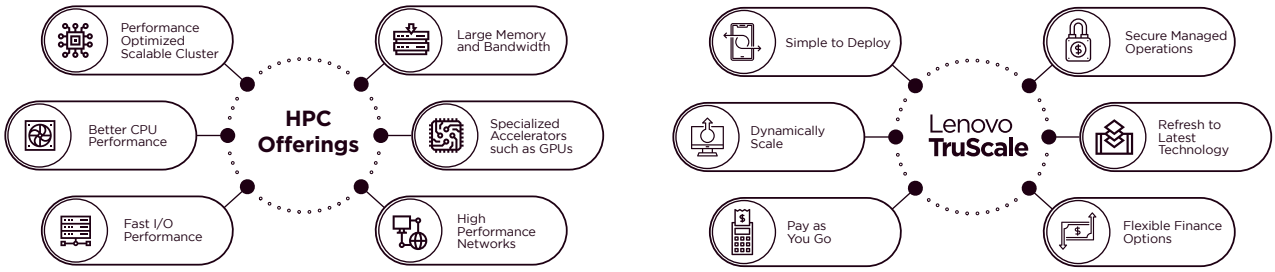

Figure 4: Characteristics of HPC cluster offerings and benefits of "as a Service" models

HPCaaS helps users save on capital spending and still realize their business objectives. Since the service provider manages all the complexities of the computing environment, customers can get the best value from their investment and focus all their precious resources on improving research and discovery. However, HPC public clouds have many limitations.

Overcoming public cloud limitations

Cloud computing has become an increasingly popular tool for genomics research due to its ability to store, manage, and analyze large amounts of data. However, there are several limitations of cloud computing for genomics (which multiply with the number of users and the size of the data files), including:

1. **Operating Cost:** While cloud computing offers scalability and flexibility, it can be expensive, particularly for large-scale genomics projects that generate massive amounts of data. In addition, cloud providers charge additional fees for accessing, storing, and using the stored data. These budget constraints can limit access to cloud computing resources for smaller research teams or institutions with limited funding.
2. **Data privacy and security:** Storing large amounts of sensitive genomic data in the cloud raises concerns about privacy and security. The data must be protected from unauthorized access, theft, or breaches, no matter where it resides. What are the costs associated with protecting the data in the cloud? Is data stored redundantly to ensure recoverability? How much control does the user have over remote storage?
3. **Regulatory Compliance:** Most life sciences IT solutions must be validated for compliance with FDA and other agency regulations. The applications and the installed systems environment must be validated on the cloud. With localized IT infrastructure, this is assured by the internal IT group. However, the cloud provider must also do this in the cloud. This typically adds additional costs, risks, and complexity, especially during some compliance audits.
4. **Usability:** Genomics involves a pipeline of scripts and a command line interface. Setting up nodes in the cloud requires users to understand the system's intricacies to set up these command lines and scripts manually. In addition, sourcing and optimizing specific applications needed for the entire life sciences workflow is cumbersome and sometimes not feasible. All this is expensive and time-consuming and may have to be repeated for each user type in the life sciences organization. Life sciences R&D and innovation require close collaborations between scientists to visualize and interpret analysis results. Effective remote visualization is needed to mitigate considerable data movement challenges from the cloud to remote users when working with large data sets on a cloud. Unfortunately, many public clouds do not support these capabilities effectively, thus hampering user productivity and innovation.
5. **Internet connectivity/Data transfer times:** Cloud computing requires high-speed internet connectivity to access and analyze large datasets. Internet connectivity can be a limitation for researchers working in remote or low-resource settings, where internet connectivity may be unreliable or unavailable. In addition, uploading and downloading large datasets to and from the cloud can take significant time and resources, particularly for researchers with limited bandwidth or slow internet connections.
6. **Technical expertise/Dependence on cloud providers:** Cloud computing requires technical expertise to set up and manage the infrastructure and tools necessary for genomics research.

Lack of technical expertise can be a limitation for researchers who may not have the skills or resource to use cloud computing resources effectively. In addition, cloud computing relies on third-party service providers, which can limit researchers' control over their data and analysis tools. This can raise concerns about vendor lock-in and the long-term sustainability of cloud computing resources.

However, life sciences users can still benefit from public cloud environments and bridge the gap between the freedom of the public cloud and the security and control of an on-premises solution. Indeed, factors such as application availability, performance requirements, and data governance and sovereignty regulations mean organizations should consider on-premises technologies that can deliver the benefits of a cloud operating model while keeping the business entirely in control of the infrastructure supporting these applications. These solutions can provide the best of both worlds: cloud-based economics with on-premises resources, combining their performance and security requirements with the flexibility of embracing a hybrid approach.

High-level architecture of Lenovo TruScale HPCaaS

The Lenovo TruScale HPCaaS offering for life sciences is built with this hybrid approach. It uses a foundation of leadership and expertly engineered on-premises high-performance infrastructure ranging from workstations to the edge to the data center.

Expertly engineered Lenovo solutions for life sciences

Lenovo offers a broad portfolio of workstations, servers, storage, software, and genomics services, many of which are part of Lenovo TruScale (Figure 5). Key components of the portfolio include:

- **Lenovo GOAST Architecture :** Lenovo's Genomics Optimization and Scalability Toolkit (GOAST) leverages an optimized variant-calling workflow and a concise, simple, non-specialty hardware recipe to deliver an affordable solution with peak performance on Lenovo hardware. GOAST Architecture delivers a 27X to 40X speed-up in analyzing whole genome sequencing (WGS).

  GOAST Scaler: This is a tool for sizing and scaling HPC for life sciences workloads. GOAST Scaler calculates the projected HPC usage for an expected workload. For example, it outputs the compute nodes, active and archive storage needed to meet a workload quota (e.g., 50K genomes/yr.). GOAST Scaler can also be used to size the current production capabilities of an existing cluster e.g to answer the question of how many genomes can process with my current cluster or how many genomes/yr. can this year's budget afford me?

- **Lenovo Workstations:** **ThinkStation P** Series workstations deliver powerful performance: from the **ThinkStation P620** with its AMD Ryzen™ Threadripper PRO to the latest generation of Intel® Xeon® processors combined with up to 2x NVIDIA RTX™ A6000 graphics cards in dual and single-processor systems. These are ISV-certified, energy-efficient, and highly versatile.

- **Lenovo Edge Servers:** Deliver purpose-built and secure platforms suitable for compute-intensive and latency-sensitive applications deployed outside traditional data centers. They range from purpose-built, compact, and secure **ThinkEdge SE-350** to the latest AI server for the edge – **ThinkEdge SE-450**.

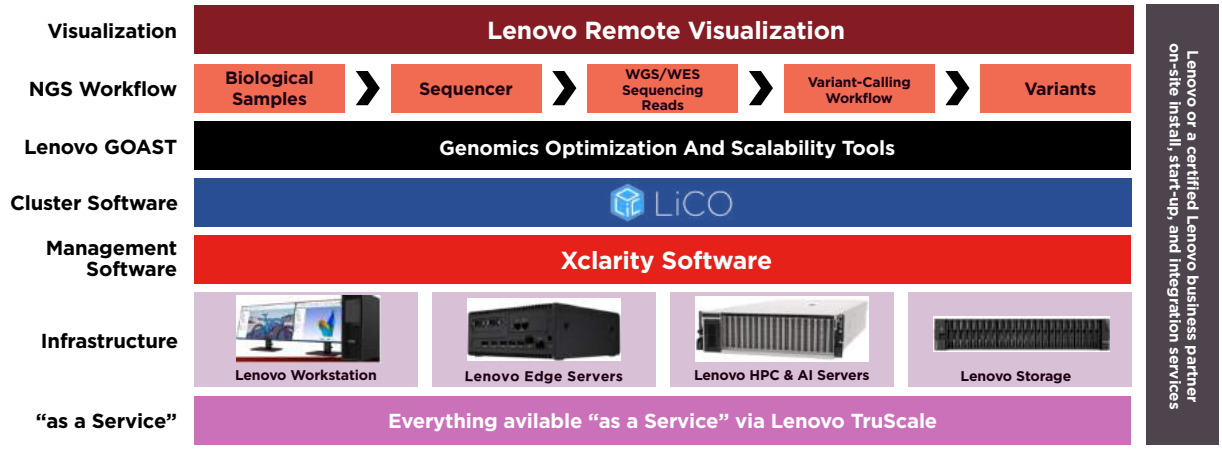

Figure 5: Lenovo life sciences solutions portfolio available on Lenovo TruScale

- **Lenovo ThinkSystem Servers:** Highly reliable, scalable, and high-performance servers to significantly accelerate Life Sciences. This Lenovo **portfolio of servers** includes **Lenovo Neptune™ liquid cooling technologies, Lenovo ThinkSystem™ SR675 V3, Lenovo ThinkSystem™ SD665 V3, Lenovo ThinkSystem™ SD665-N V3, Lenovo ThinkSystem™ SD650-N V2 rack server**; some of these servers also come with Lenovo Neptune hybrid cooling module, which quickly dissipates heat in a closed-loop liquid-to-air heat exchanger (L2A), delivering the benefits of liquid cooling without adding plumbing.
- **Lenovo Storage: Direct-Attached Storage** JBODs and expansion units provide flexible, cost-effective, high-capacity storage and are ideal for space-constrained environments and cost-sensitive customers. **ThinkSystem DE Series All-Flash Array** is designed for extreme performance with up to 1.0M IOPS and sub-100 microsecond latency. In addition, it has a full suite of APIs and application plugins for easy integration.
- **Lenovo xClarity®:** This family of software simplifies and automates the deployment and management of Lenovo infrastructure so clients can focus on their high-value projects.
- **Lenovo Intelligent Computing Orchestration (LiCO):** Reduces the complexity of using a massive HPC cluster and simplifies genomics application deployment, operation, and acceleration.
- **NGS Workflow:** Biological samples (i.e., blood, saliva, etc.) processed experimentally are input into a sequencer, generating sequencing reads (fragments of DNA strings). The sequencing reads become the input for the variant calling workflows. The output of the genomics workflow feeds into variant analyses and downstream tertiary bioinformatics work (e.g., measuring the effect of variations on function or disease).
- **Lenovo Remote Visualization:** Provides reliable and secure access to graphics-intensive applications anytime, anywhere. Instead of issuing new expensive workstations to all design staff, IT can deploy less expensive enterprise or consumer-class personal computers. In addition, IT departments can maintain security and keep costs down by using remote virtualization hosted in an internal data center or from the cloud. Remote visualization performs intensive graphics operations on a high-end graphics server and generates a 2D pixel version that users can receive quickly. In addition, server-side rendering considerably speeds up the process of using graphics in remote sessions.
- **Lenovo or Certified-Partner Services:** Lenovo and an extensive ecosystem of highly specialized HPC services partners can deliver the integrated Lenovo stack depicted above. They also provide on-site installation and start-up services to integrate this into a client's work environment, including life sciences workflows.

Lenovo TruScale includes most of the life sciences solutions portfolio (Figure 5) as an "as a Service" model and provides all the benefits of public clouds while overcoming their many limitations.

Lenovo TruScale is affordable and optimized for life sciences

High-performance life sciences systems are expensive, evolving, and often highly capital-intensive. In addition, life sciences users must integrate new technologies and workloads efficiently and seamlessly as technology rapidly evolves, often within the resource, budget, and capital restrictions.

Lenovo TruScale is optimized for today's life sciences workloads and designed for future data-intensive applications when NGS is integrated with analytics, AI/ML, and other disciplines throughout the digital healthcare system. Figure 6 summarizes the key benefits of TruScale for life sciences.

TruScale for Life Sciences	Benefits
Plug & Play	• Right-sized HPC infrastructure is designed, installed, and configured by Lenovo HPC experts • Deployed in client's data centers (on-premises) and col-ocation facility • Infrastructure is fully managed and operated by Lenovo HPC experts
Flexible Financial Model	• Align costs to business outcomes with the Lenovo OpEx solution • Minimize up-front capital investment • Control costs with Lenovo's predictable pricing model • Reduce TCO by helping to eliminate over-provisioning
Simplified Operations	• Easily provision HPC resources and get support through the Lenovo HPC team • Centralized management through the Lenovo HPC team • Get expert help to design, deliver, and operate from a single vendor • Incident resolution & support from Lenovo's HPC expert, securing the running of operations and uptime.
Latest Technology	• Access to the latest technology to help run client's environment • No need for capital investment: avoid long procurement approvals with Lenovo's OpEx model • Technology is refreshed as it becomes obsolete, Lenovo will manage the upgrades and replacements
Customized Solution	• Solution customized and optimized to a client's specific life sciences computational needs

Figure 6: Benefits of the Lenovo TruScale solution for life sciences

TruScale is a new way of procuring IT resources via a consumption-based subscription model. With this new model, customers never take capital ownership of the hardware or other IT assets. They pay for what they use each month only for capacity when their workloads are actively running as part of their operating expenses with no minimum capacity commitment. Lenovo TruScale includes hardware installation, deployment, management, maintenance, and removal.

Lenovo's unique metering solution remains outside the customer's data plane – providing the advantages of cloud-like economics with the security of on-premises hardware. Additionally, capacity can be scaled up or down to accommodate business needs, ensuring IT infrastructure is always right-sized. In one bill, monthly pricing structures are simple and all-inclusive of associated services (maintenance, support, remote monitoring, and system health).

Get started with Lenovo TruScale with dedicated support

Life sciences teams cannot afford performance problems, delays, or downtime. Therefore, support must be proactive, and must be carried out by technical specialists, who work closely with the customer, and have a deep understanding of their environment.

With Lenovo, life sciences researchers will have an HPC technical account manager or system admin as part of their Lenovo TruScale contract as their single point of contact. Whether onsite, working remotely, or a mixture of both, the support professionals can quickly pinpoint and resolve any issues and ensure the life sciences environment runs optimally 24/7.

However, Lenovo goes way beyond specialized technical support. TruScale's end-to-end service for life sciences includes initial consultation, analysis, and configuring the right environment through ongoing cooling assessment and maintenance services to billing and administration.

The Lenovo TruScale advantage for life sciences

Lenovo TruScale helps life sciences companies maximize their return on investment in their high-performance infrastructure solutions and accelerate time to value to drive innovation. It also gives the flexibility to rapidly add provisioned resources without being hindered by procurement delays or supply chain disruptions – something vital in today's market.

Consuming high-performance infrastructure is more transparent and affordable than traditional on-premises solutions, without the limitations of public clouds. Furthermore, as every life sciences researcher has unique needs and requirements, working with Lenovo will help achieve a solution tailored to their evolving workload, workflow, and workforce needs.

Lenovo TruScale

Start planning your cloud journey with Lenovo

Arrange a discussion with one of our experts now.
truscale@lenovo.com

Lenovo.com

[1] https://www.lenovo.com/us/en/resources/data-center-solutions/solution-brief-documents/goast-genomics-optimization-and-scalability-tool/

Lenovo reserves the right to alter product offerings and specifications, at any time, without notice. Lenovo makes every effort to ensure accuracy of information but is not liable or responsible for any editorial, photographic, or typographic errors. Images are for illustration purposes only. For full Lenovo product, service, and warranty specifications, visit www.lenovo.com. Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo. Other company, product and service names may be trademarks or service marks of others.
© Lenovo 2023. All rights reserved