Revolutionizing life sciences and interdisciplinary research through easy access to ultra-modern HPC cloud platform.

How **CLIP at Vienna BioCenter** uses Lenovo ThinkSystem servers, powered by NVIDIA® GPUs, to accelerate breakthrough research in bioscience and physics.

Lenovo Infrastructure Solutions for The Data-Centered

NVIDIA

Lenovo

1

Background

Vienna BioCenter is a leading life sciences campus in Europe with 2,685 staff and 141 research groups, combining fundamental research and applied sciences with cutting-edge facilities offering the latest technologies to advance the understanding of biomedical processes. Vienna BioCenter was founded by five research institutions to innovate and accelerate interdisciplinary collaboration. The effort is driven by the Austrian Academy of Sciences (OeAW), the largest Austrian non-university institution for science and research with 25 research institutes, two of which are founding partners of the Vienna BioCenter. Additional partners include the University of Vienna, the Medical University of Vienna, and the global, research-driven pharmaceutical company Boehringer Ingelheim.

To help tackle the toughest research questions, Vienna BioCenter established the Cloud Infrastructure Platform (CLIP), a high-performance computing (HPC) service currently used by 14 research institutes and open to researchers around the world. One of Austria's largest supercomputer clusters, CLIP offers an advanced and easy-to-use cloud service with over 400 standard tools for high energy physics, subatomic physics, mathematics, sound mechanics, space exploration, and the full array of life sciences.

CLIP is a partner of CERN and supports the operations of the world's largest and most powerful particle accelerator, the Large Hadron Collider (LHC). Its HPC resources form part of the Worldwide LHC Computing Grid (WLCG), hosting the only Tier-2 Grid computing site in Austria to run analysis tasks for some of the world's top research scientists.

2

Challenge

To help 350 leading researchers uncover new insights into the physical world, plants, and the human body, CLIP needed an extensible and adjustable HPC platform that could grow over time and adapt to meet the evolving needs of the research community.

Ronny Zimmermann, Head of IT at Cloud Infrastructure Platform (CLIP) at Vienna BioCenter, says: “We provide researchers with an easy-to-use and scalable scientific cloud computing environment. We enable them to run experiments, without having to set up tools or buy and configure complex workstations individually. With this approach we give curiosity-driven research much more space. That in turn allows for initial explorations or large-scale experiments on up to 500 TB of data—saving all sides time and money.”

CLIP realized that more and more scientific tools now require support for GPUs. With cutting-edge technologies and algorithms in the area of artificial intelligence and machine learning becoming easier to access for a larger group of scientists, more groups have started using neural networks and deep learning to push the boundaries in their fields.

Erich Birngruber, HPC Specialist at CLIP at Vienna BioCenter, comments: “When working from home became the default for many researchers during the COVID-19 pandemic, we noticed a spike in demand for our JupyterHub service where scientists run interactive sessions using widely used tools such as Jupyter Notebooks to develop complex machine learning models. We now keep 150 interactive sessions running in parallel, on top of the standard batch workload that already constantly runs in the background.”

Tasked with helping a specific group of scientists with groundbreaking research, CLIP also had to act quickly to avoid costly delays to important projects that required more computing power and new, faster GPU processing capabilities than the organization could offer. “The service had worked so brilliantly that other institutes of the Austrian Academy of Sciences wanted to join. We quickly had to adapt our capacity to ensure the high standard of service they are used to getting from us,” says Ronny Zimmermann. “As a result, the demand for cost-efficient GPU performance exploded. Together with NVIDIA and Lenovo, we were able to provide additional resource in the shortest time possible.”

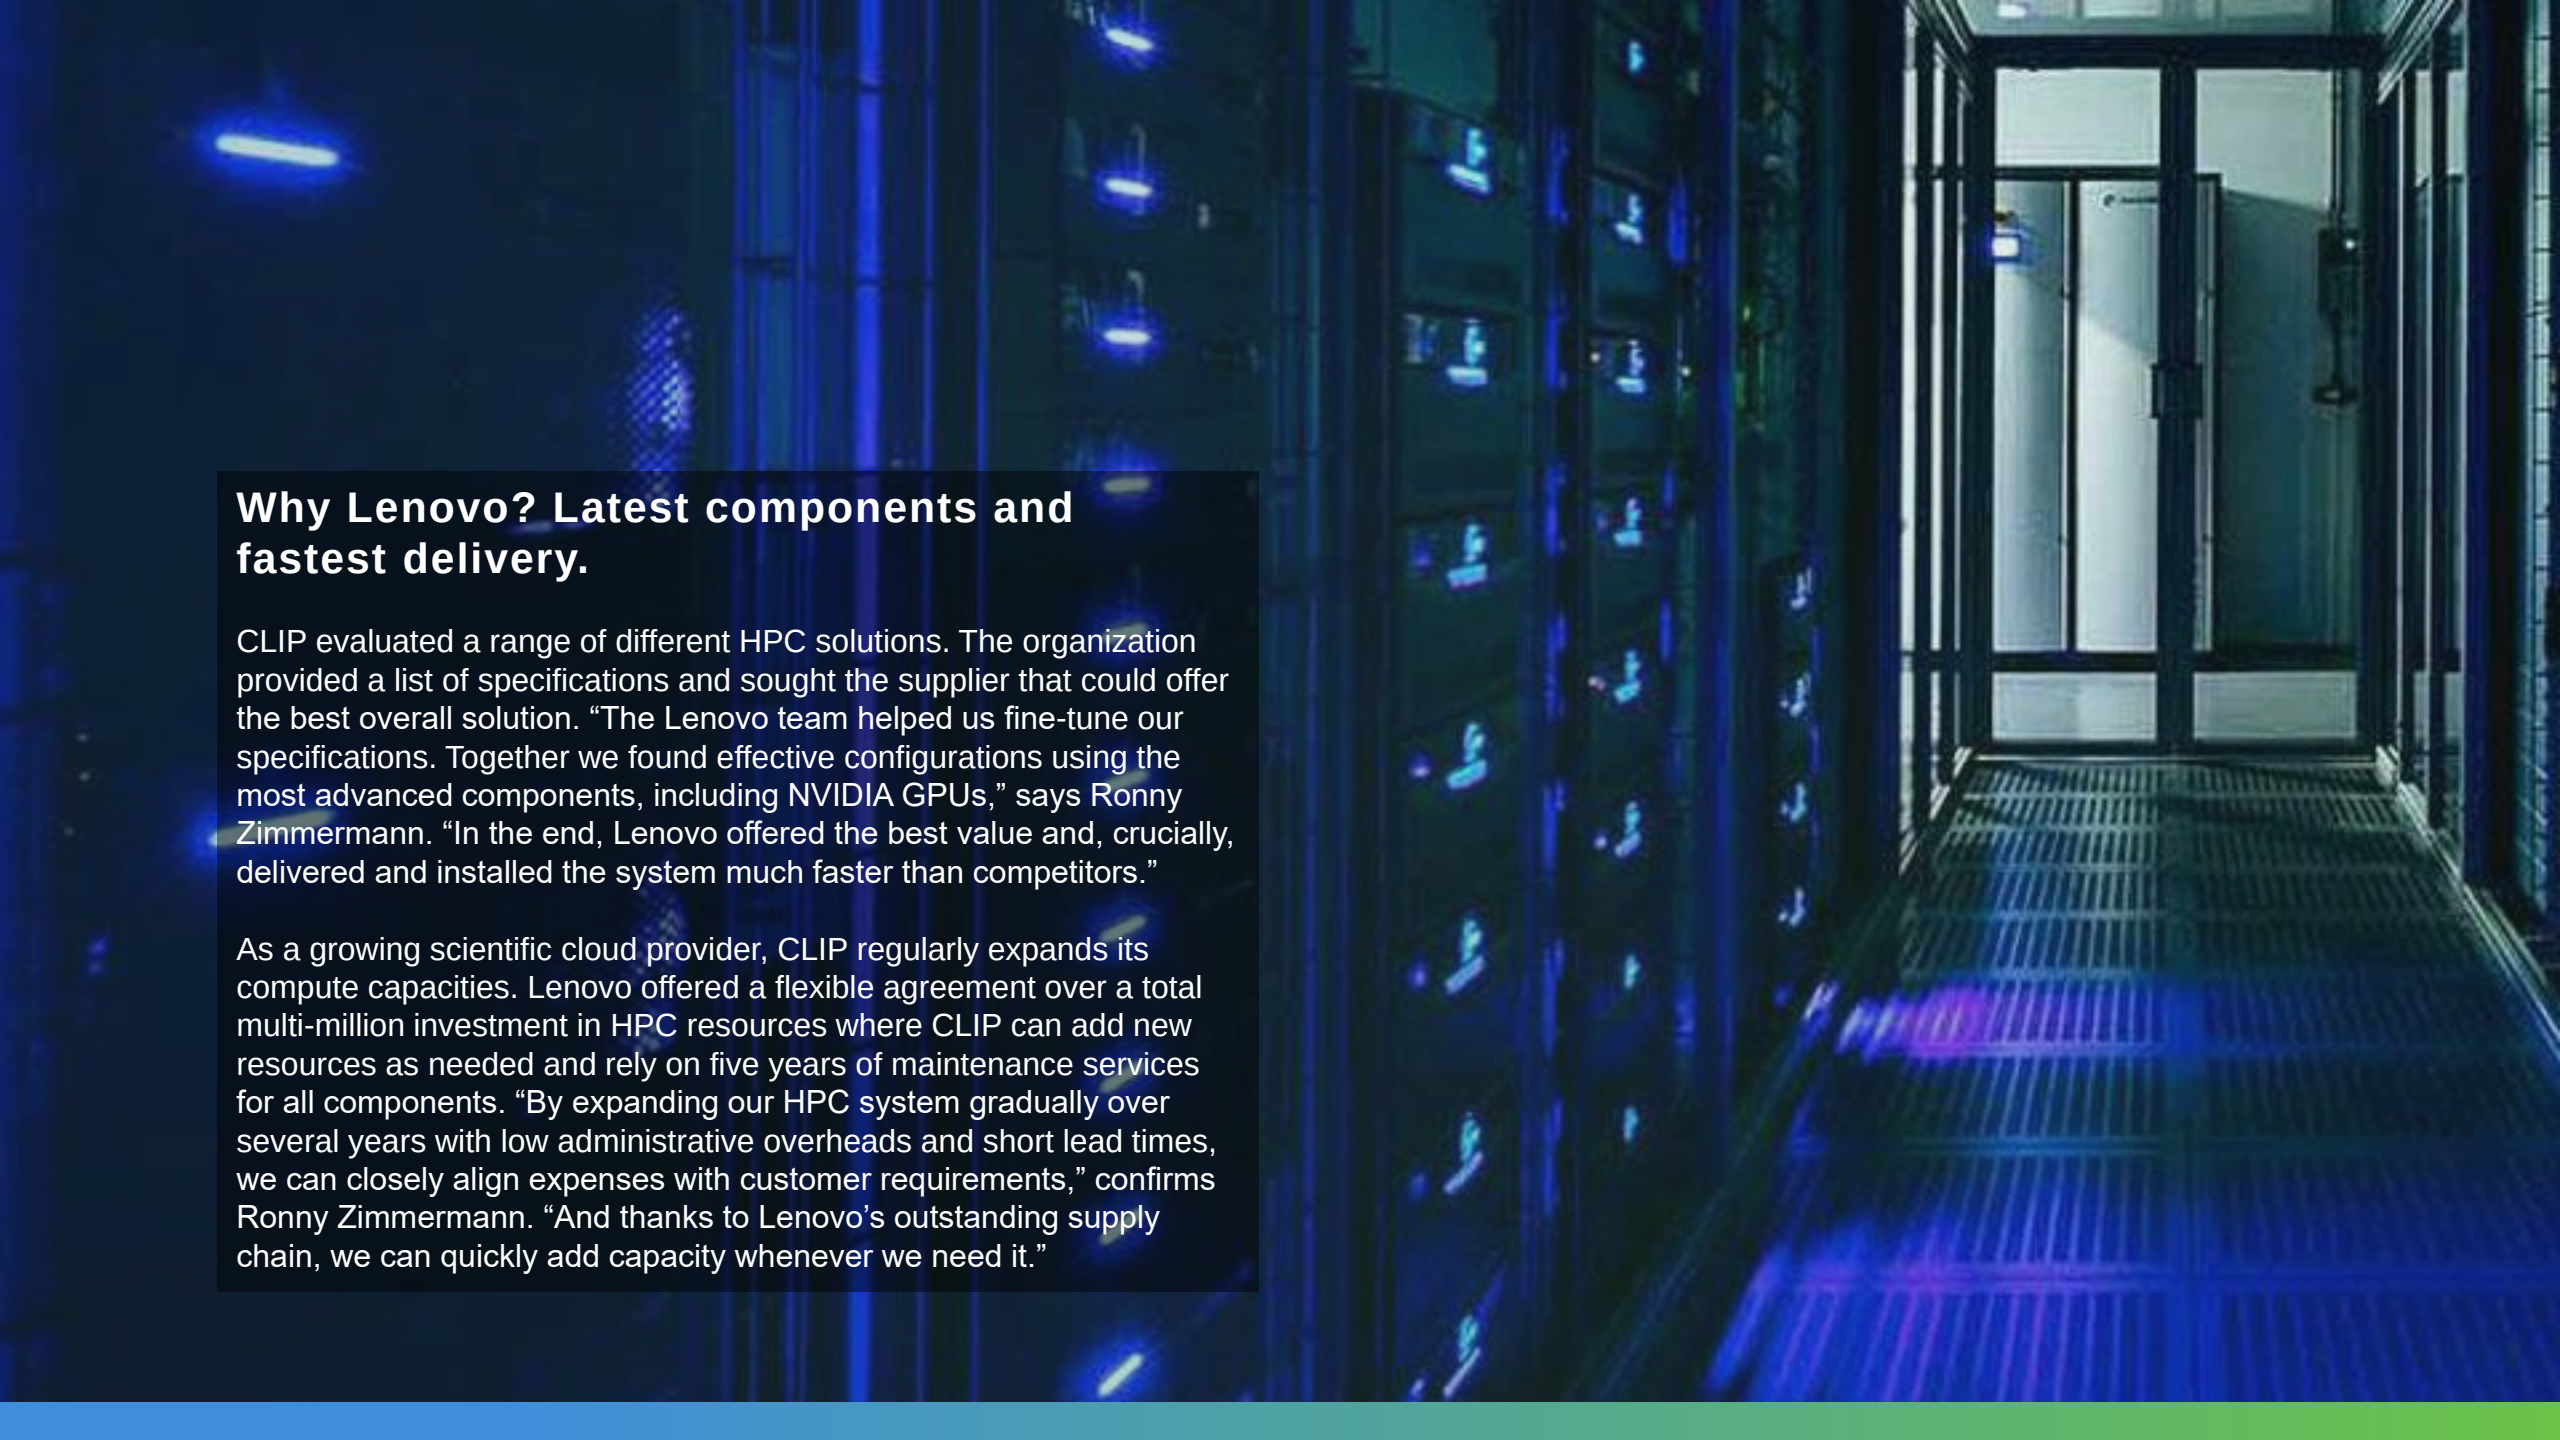

Why Lenovo? Latest components and fastest delivery.

CLIP evaluated a range of different HPC solutions. The organization provided a list of specifications and sought the supplier that could offer the best overall solution. “The Lenovo team helped us fine-tune our specifications. Together we found effective configurations using the most advanced components, including NVIDIA GPUs,” says Ronny Zimmermann. “In the end, Lenovo offered the best value and, crucially, delivered and installed the system much faster than competitors.”

As a growing scientific cloud provider, CLIP regularly expands its compute capacities. Lenovo offered a flexible agreement over a total multi-million investment in HPC resources where CLIP can add new resources as needed and rely on five years of maintenance services for all components. “By expanding our HPC system gradually over several years with low administrative overheads and short lead times, we can closely align expenses with customer requirements,” confirms Ronny Zimmermann. “And thanks to Lenovo’s outstanding supply chain, we can quickly add capacity whenever we need it.”

Building a flexible supercomputer with high-performance NVIDIA GPUs.

Over the years, CLIP has relied on a range of Lenovo servers with various NVIDIA® GPU accelerators to deliver the high performance the researchers need. Today, CLIP is based on 200 Lenovo ThinkSystem SD530, SR630, SR650, SR670, SR850, and SR950 servers. As part of the latest expansion to boost machine learning performance and other GPU-optimized workloads, CLIP has deployed NVIDIA A100 Tensor Core GPUs with 40 GB HBM memory. These leading-edge GPUs are powered by the NVIDIA Ampere architecture. The team plans to take advantage of NVIDIA NVLink® to provide high-speed direct data transfer between GPUs for even faster cluster performance.

The team has also implemented fast, low-latency switches from Lenovo to maximize network performance. In addition to servers and networking, CLIP worked with Lenovo to implement Intelligent Cluster Solutions for data storage. Building on Lenovo Scalable Infrastructure (LeSI), CLIP operates several nodes with 250 TB of ultra-fast NVMe flash storage using the software-defined storage solution BeeGFS. “The nature of highly parallelized processing on GPUs requires new levels of storage performance that are difficult to achieve with traditional storage arrays,” adds Erich Birngruber. “The combination of NVIDIA GPUs and fast, scalable, software-defined storage from Lenovo accelerates innovative research across many fields with smarter technology.”

The small team at CLIP can only run such a large cluster because the infrastructure is extremely reliable and Lenovo support is always at hand if the team has any questions. “It’s great to see that we can count on Lenovo for help when we really need it,” says Erich Birngruber. “By working together, we minimize the risk of outages or failures, and ensure our users get the best performance.”

CLIP uses Lenovo XClarity to manage all 200 ThinkSystem servers cost-efficiently. Centralized monitoring and management with Lenovo XClarity Administrator and the embedded Lenovo XClarity Controller solution helps CLIP to stay on top of its infrastructure by standardizing administration tasks. “The automated health checks and ticketing integration is excellent,” notes Erich Birngruber. “When Lenovo XClarity detects any technical issues, it not only notifies us, it also opens a support ticket and notifies Lenovo—this streamlines system administration, delivering an easy-to-manage, agile, and proactive HPC solution.”

The highly automated scientific supercomputing cloud at CLIP is fully virtualized with OpenStack running optimized instances of CentOS Linux as well as Singularity containers, offering flexibility and high performance for a wide range of research tasks. Thanks to this architecture, the cloud can be dynamically and rapidly reconfigured for different workloads to offer all users an environment tailored to their specific research requirements.

3

Results

Thanks to fast Lenovo logistics, CLIP got its new HPC solution online in just two months. “Lenovo delivered the infrastructure within only 12 days and about four times faster than other vendors could commit to,” recalls Ronny Zimmermann. “This was an amazing achievement. Scientists are impatient customers. Building the infrastructure in such a short timespan meant that they could continue their work without significant interruption.”

With over 8,000 CPU cores, 32 TB memory, and more than 120 NVIDIA GPU accelerators supported by 2.4 TB GPU memory, CLIP can now provide its 350 researchers with state-of-the-art performance—enabling new discoveries in many disciplines including DNA sequencing, theoretical modeling, and scientific visualizations of structures. Using the advanced Lenovo tools, the small team at CLIP can easily manage the growing environment and focus on fast customer service, ensuring rapid onboarding of new researchers. “Our Lenovo HPC cluster powers key scientific discoveries,” confirms Erich Birngruber. “With our help, researchers at Vienna BioCenter publish 350 academic papers every year and gain international recognition.”

By deploying Lenovo ThinkSystem servers with NVIDIA GPUs, CLIP can increase performance density and deliver faster data processing with a smaller physical footprint, leaving the organization more space to grow in its existing data center. The researchers can also benefit from new technologies faster.

“Without the NVIDIA GPUs, we simply could not offer the same level of performance,” says Ronny Zimmermann. “In recent years, more and more bioinformatics tools that combine statistics, computer sciences, mathematics, and computational statistics have closely integrated with GPUs to improve performance. In line with this trend, the number of GPUs in our Lenovo HPC cluster has grown by a factor of three, from 40 to 120. With 30 research groups, about 30% of all groups on the cluster are actively using GPUs right now. When we look at the processing power, NVIDIA GPUs already make up 40%, with traditional CPU workloads accounting for 60% of the cluster utilization.”

Vienna BioCenter

- **4x faster HPC delivery accelerates innovative research**
- **8,000 CPU cores and 120 NVIDIA GPU accelerators enable rapid insights for 350 researchers**
- **Flexible HPC agreement over EUR 10 million supports alignment of infrastructure expansion with business requirements at low administrative cost**

Supporting groundbreaking scientific research.

A core facility at Vienna BioCenter provides cutting-edge cryogenic electron microscopy services. The team relies on the Lenovo HPC cluster to analyze a constant stream of data and visualize biological samples in high resolution with 2D and 3D imaging techniques. The high-throughput sample screening solution employs the most sophisticated electron microscopes to visualize surface structures of molecules including DNA, RNA, and proteins, as well as bacteria and viruses. “For this team, the NVIDIA GPUs are extremely important,” says Ronny Zimmermann. “Processing surfaces is at least 10 times faster on GPUs. And the next generation of microscopes will generate even more data. With enough GPU performance, we will soon observe samples in real-time and quicker than ever before.”

Meanwhile, researchers at the Max Perutz Labs at Vienna BioCenter rely on CLIP to analyze terabytes of data, using polymer physics theories to build physical models of genomes. “Our teams are at the forefront of 3D genomics,” says Ronny Zimmermann. “These molecular dynamics simulations run entirely on the NVIDIA GPUs. The researchers aim to answer fundamental questions of genome biology to support breakthrough advances in human health by gaining a better understanding of biomedical processes.”

Anton Golobordko, Group Leader at Institute of Molecular Biotechnology (IMBA), confirms: “Polymer simulations of the genome require a massive amount of computational resources. One individual simulation involves up to a million particles and has to be run for up to 100 million steps. The capacities of available computational resources put a hard limit on the types of questions we can study. For example, long-range topological effects and long-term equilibration of the genome in the presence of phase separation cannot be approached without a large-scale GPU cluster.”

CLIP was also instrumental in research aimed at plant breeders to identify genetic determinants of certain characteristics such as leaf shapes and patterns. “We helped the researchers to run their ARADEEPOPSIS toolset, a deep-learning model based on convolutional neural networks, CNN for short, and the TensorFlow open source platform for machine learning with GPU support,” says Erich Birngruber. “Based on the DeepLabV3+ model the researchers label images and re-train the AI with new data. The work included automated segmentation of about 150,000 plant images.”

Additional workloads at CLIP that benefit from fast GPU performance are interactive Linux desktop sessions where users can run specialized tools with graphical user interfaces right on the HPC cluster. “We continue to focus on easy access to our HPC cloud,” states Erich Birngruber. “That’s why we offer a wide range of popular and modern tools, such as Jupyter Notebook with Python or the R programming language, integration with RStudio Server Pro, and Visual Studio Code. Our goal is to make it as easy as possible for every researcher to use our scientific cloud platform, using familiar tools from their local setup.”

Ronny Zimmermann concludes: “Our mission is to save researchers time and money, so they can be true pioneers in their fields. At the end of the day, it is thanks to science that we’re not living in caves anymore! With Lenovo and NVIDIA HPC solutions, we help advance science further every day.”

“While our focus is currently on the Austrian research community, we are more than happy to onboard scientists from around the world. Moving forward, we’re also planning to win new users by opening up our innovative cloud platform to research-intensive businesses. In supercomputing, economies of scale are essential: the more users we can attract, the more compute capacities our researchers can use for their projects to make new discoveries faster and shape the future.”

Ronny Zimmermann
Head of IT, Cloud Infrastructure Platform (CLIP) at Vienna BioCenter

What will you do with Lenovo HPC solutions?

The Data-Centered facilitate new discoveries to shape the future with Lenovo smarter infrastructure solutions, powered by NVIDIA®.

Explore Lenovo HPC Solutions