Goethe University Frankfurt delivers broad range of scientific research with AMD

2nd Gen AMD EPYC™ CPUs and AMD Radeon Instinct™ MI50 GPUs power discoveries in particle physics, climate research, digital medicine and more

AMD
EPYC | INSTINCT

CUSTOMER

FIAS Frankfurt Institute for Advanced Studies

INDUSTRY

Theoretical science research

CHALLENGES

Delivering cost-effective HPC modeling and analysis capability to nearly 50 different research groups

SOLUTION

Deploy Servers with AMD EPYC™ CPU and Radeon Instinct™ MI50 GPU

RESULTS

Achieved cost-performance objectives to grow capacity to support more research

AMD TECHNOLOGY AT A GLANCE

AMD Radeon Instinct MI50 GPUs

Theoretical science pushes boundaries, furthering our understanding of fundamental natural phenomena. The computing platforms used to model and analyze such complexity in detail must be capable of handling everything imaginable, and more. For the Center for Scientific Computing (CSC) at Frankfurt Institute for Advanced Studies (FIAS), located at Goethe University, that demand is amplified by the challenge of supporting high-performance computing needs across nearly 50 independent research groups exploring life sciences, theoretical physics, neuroscience, computer science and systemic risk.

"More than 98 percent of processing for the 600 gigabytes per second data stream runs entirely on servers with eight AMD MI50 GPUs that deliver 90% of the total compute performance we need using just the GPUs."

Professor Volker Lindenstruth, chairman of the Board of Directors, FIAS

Since the progress of such critical work is constrained by the available compute capabilities, CSC is hard at work deploying a new HPC platform; its first to take advantage of the performance and scalability afforded by combining CPUs with multiple high-performance GPUs. "If you launch a lot of jobs on a multi-core CPU alone, you quickly begin to watch the clock rate deteriorate," says Professor Volker Lindenstruth, chairman of the Board of Directors, FIAS. "You lose a good fraction of the performance you would expect to gain by adding parallelism, and, unfortunately, defeat the purpose of a multi-core processor."

Flexibility supports incredible variety

To overcome that limitation, Lindenstruth and his team selected the combination of 2nd Gen AMD EPYC™ processors and AMD Radeon Instinct™ MI50 GPU computing accelerators to obtain the combination of performance, advanced I/O capabilities, and operating efficiency needed to accommodate their clients' many varied applications.

And the diversity of research is impressive. There are groups conducting ab initio calculations called lattice quantum chromodynamics (QCD) that model nuclear reactions without free parameters. Another group uses the ultra-relativistic molecular dynamics (UrQMD) simulation code, developed at Goethe University, for applications across particle physics, high energy experimental physics and engineering, shielding, detector design, cosmic ray studies, and medical physics. There are yet other groups exploring quantum chemistry, AI, deep learning, atmospheric and climate research, and, increasingly, digital medicine supporting the Institute's world-renowned cancer research efforts.

Solving today and preparing for tomorrow with AMD Radeon Instinct MI50

To carry this breadth of experimentation, FIAS CSC is deploying a new server cluster featuring 880 AMD Radeon Instinct MI50 GPUs and supported by 100 gigabits per second InfiniBand networking. Incredibly, even as this new system just comes online, "We are in the process of buying an additional system with 2,000 AMD MI50s," says Lindenstruth.

That third system will support the work of the GSI Helmholtz Centre for Heavy Ion Research, which analyzes data generated by the Large Hadron Collider at the CERN Particle Physics Laboratory, in Geneva. GSI is preparing to support CERN's third accelerator, expected to come online within a few years. Of course, the GSI team isn't waiting. Lindenstruth explains, "We have already achieved a quantum leap with the online software that runs as the LHC experiment is taking data. More than 98 percent of processing for the 600 gigabytes per second data stream now runs entirely on SMC servers with eight AMD MI50 GPUs that deliver 90% of the total compute performance we need using just the GPUs."

Eight turns out to be a sort of magic number for the entire FIAS MI50 deployment. "We learned from our older systems that the cost-effectiveness of just one GPU on a server wasn't good when you factor in the expense of building and maintaining the server," says Lindenstruth. "With some analysis, we determined that a server built with AMD EPYC CPUs and eight AMD MI50 GPUs delivered ideal cost-performance. And, as our GPU code is better optimized, we achieve even greater efficiency."

"We determined that a server built with AMD EPYC CPUs and eight AMD MI50 GPUs delivered ideal cost-performance. And, as our GPU code is better optimized, we achieve even greater efficiency."

Professor Volker Lindenstruth, chairman of the Board of Directors, FIAS

The need for speed

Looking back, Lindenstruth explains, "One of the limiting factors in the past was memory bandwidth. In CPU-only systems, we did our best to hide transfer latency, but that is only possible to a limited extent." He continues, "You have to get the data in and out at some point and this is a factor which affects overall performance."

The MI50 features PCIe® Gen 4 technology and AMD's exclusive dual Infinity Fabric™ Link technology to provide up to 248 GB/s of aggregate GPU data transfer bandwidth. As Lindenstruth says, "The MI50 GPU combines fast memory bandwidth of a terabyte per second with very high compute capability. Even better, MI50 has a fast PCI bus. When you combine AMD MI50 GPUs, EPYC CPUs, very good memory bandwidth, fast PCI Express 4, and 100 Gb/sec Infiniband networking, you can accommodate the requirements of all our applications." He concludes, "This is the ideal system we were looking for."

Freedom through open-source HIP programming

One of the important jobs of the CSC team is to help research teams moving their legacy code to the new GPU-enabled systems. The variety is impressive. According to Lindenstruth, "The relativistic molecular dynamics code was developed in Fortran for over 25 years. We started using GPUs for CERN in 2009 using a consumer GPU and CUDA, which we later converted to OpenCL for efficiency."

Now CSC is converting to open-source HIP supported by AMD GPU accelerators. "I can't say that there is a simple red button," says Lindenstruth, "but the most important part of HIP is that the same source code will run on whatever GPUs we choose. We can make decisions based on cost performance and vendor relationship, without being constrained by software compatibility."

A vision of the future

Looking ahead, Lindestruth opines, "It doesn't make sense to focus on using only CPUs anymore, especially when you consider the huge amount of CPU cycles required by work such as CERN. There are a huge number of tier centers doing nothing but data processing for CERN, and they are all exclusively CPU based. Well, there is one exception and that is GSI! We are the pioneers making this distributed processing GPU enabled. We think that all data centers doing particle physics, and not only GSI and others supporting CERN, will be going in the direction of GPUs."

"The collaboration and support we have received from AMD are outstanding," says Lindenstruth. "We have achieved a lot. Challenges, sometimes self-made, are inevitable in large scale systems. It's good to know that we have a partner who helps us quickly iron them out."

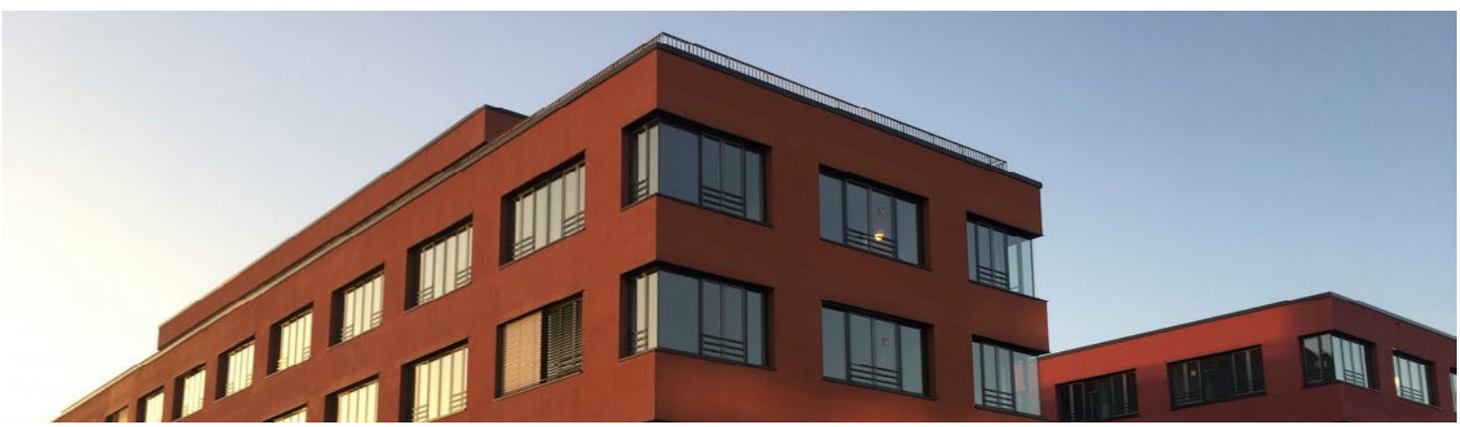

About Frankfurt Institute for Advanced Studies (FIAS)

As a platform for integrating the sciences, FIAS provides a foundation for decisive progress in research through cooperation, exchanging ideas, and overcoming structural barriers between the disciplines. About 180 scientists from more than 20 countries are doing research at FIAS. For more information visit fias.institute.

About AMD

For 50 years AMD has driven innovation in high-performance computing, graphics, and visualization technologies—the building blocks for gaming, immersive platforms, and the data center. Hundreds of millions of consumers, leading Fortune 500 businesses, and cutting-edge scientific research facilities around the world rely on AMD technology daily to improve how they live, work, and play. AMD employees around the world are focused on building great products that push the boundaries of what is possible. For more information about how AMD is enabling today and inspiring tomorrow, visit amd.com/instinct.