CASE STUDY
Intel® Xeon® Processor E5 v3 Family
Data Center Efficiency
High-Performance Computing

Making the most of HPC

The High Performance Computing Center in Stuttgart can process huge amounts of data in a short time for a wide range of clients thanks to the latest Intel® Xeon® processors

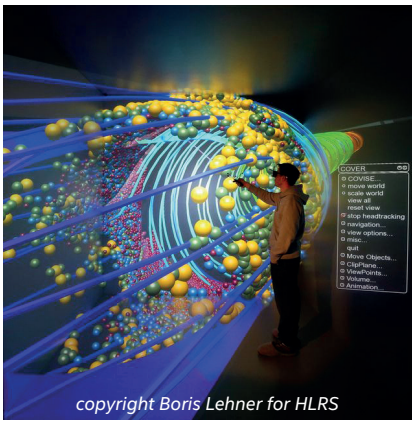

copyright Boris Lehner for HLRS

H L R S

"That such a high-level organization with so many different client needs can be served by standard Intel® solutions demonstrates the power and performance of the technology."

Bastian Koller,
Manager,
High Performance Computing Center in Stuttgart

The High Performance Computing Center (HLRS) in Stuttgart is a publically owned research institute and high-performance computing (HPC) center. It supports researchers and industry leaders throughout Germany and the rest of Europe with supercomputing technology. HLRS upgraded to the Intel® Xeon® processor E5-2680 v3 product family to achieve increased capacity and even better performance. The technology was running fully in November 2014 and, only weeks later, HLRS was ranked ninth in the 'Top500' list of supercomputers not running accelerator cards. Now HLRS can take on more projects and deliver outstanding computing services to academic clients and local companies alike.

Challenge

- **Data influx.** HLRS must process huge volumes of data for each project team, so it needed increased computing capacity to serve its users
- **Limited infrastructure.** The working maximum for HLRS is four megawatts, so the IT team needed a system that was power-efficient and still the most effective for the existing infrastructure
- **Cost efficiency.** HLRS needs to provide resources on a budget and has a fixed IT spend for the next seven years, so it relies on technological improvements to achieve better results at the same cost

Solutions

- **HPC technology.** To improve capacity and delivery time, HLRS upgraded to the Intel Xeon processor E5-2680 v3 product family
- **Performance enhancer.** HLRS uses Intel® Advanced Vector Extensions 2 (Intel® AVX2), which doubles integer packet processing ability on the same clock speeds, with significant workload performance gains

Impact

- **Reliable technology.** Intel® technology delivers reliable performance and efficient computing so HLRS is satisfied it can deliver what clients need when they need it
- **Power boost.** According to its internal tests, HLRS found the latest Intel® Xeon® processor v3 family delivered four times the peak performance of its previous system [1, 2]
- **Efficient use.** All servers can be used at all times, reducing power waste

Computing for all

Based in the University of Stuttgart, HLRS is one of three large supercomputer sites in Germany. These three centers are part of the Gauss Center for Supercomputing umbrella organization and align their IT procurements and installations so that every four to five years each center's system is updated. Each of these HPC centers has its own business model and strategy for supporting customers. HLRS makes its computing facilities available to various projects and is the only one that serves businesses. It is atypical for a university HPC center since, although it supports a great deal of research from The University of Stuttgart and further afield, it dedicates around 10 percent of its computer power to industrial projects, either directly or through a separately registered company, hww GmbH. hww GmbH is one way of handling charging from commercial clients, since academic users do not get charged due to the funding of the machines with public money.

The HLRS IT team wants to deliver top service, but having a mix of customers puts more pressure on the system, especially since commercial deals rely on it. "To achieve our goal of taking on more projects and serving more industrial users, we needed more computing capacity and processing power, which meant investing in new processors," said Bastian Koller, manager of HLRS. "We are addicted to giving best possible service, so it is vital that we run on reliable technology and are able to solve any issues immediately. For the sake of our commercial and academic customers, the system simply cannot go down. The challenge is that our scientific and commercial clients want their projects run immediately. They want a tailored service and they want it now, so we cannot have an underperforming system. Our previous work with companies like Porsche and Daimler gives us a lot of credibility in industry, which is important, since organizations need to know that our system is stable and reliable, thus they won't have an outage. We want to build on this."

Leading the way

HLRS had a long-term plan to invest in new hardware to be able to run more projects in a quicker timescale. Following an internal selection process, HLRS upgraded to the Intel Xeon processor E5-2680 v3 product family. HLRS uses a mix of clients and the team valued how Intel technology maintains good performance across various

HLRS can handle more data thanks to the Intel® Xeon® processor E5 v3 family

infrastructures. The team still runs its previous machine, but added these new processors to boost capacity and project delivery.

According to internal tests, HLRS was at 86 percent usage after the first month of running the machine, which indicates how in-demand this processing power was. "It was a smooth transfer process and we are proud of our new supercomputer," said Koller. In the same month, HLRS was rated in the global 'Top 500' and came 16th in total, but was ninth out of systems not running accelerator cards. Also, although there are 20 systems based on the Intel Xeon processor E5 v3 product family, the HLRS computing infrastructure is the highest performing of them.

The center was granted early access to this latest-generation technology, making this upgrade notably early within the European academic community. "Being able to get early access to new technology and test systems prior to official launch is valuable for us and helps us to deliver excellent service to our users," said Koller. "To make this decision, we looked at what solutions were available from our major vendor, Cray, and that hit our internal key performance indicators, focusing on peak performance, capacity and ease of integration. Processing speed and minimal latency on chipsets with other parts of the system made the latest Intel Xeon processor E5-2680 v3 product family the right technology for us. Our new system is more convenient and delivers much better performance. We needed leading technology and are happy with our choice."

HLRS ran its own tests with its internal innovation group. It found that, where previously, peak performance was 1.043 petaflops, with the new system it now comfortably reaches 3.8 petaflops — an increase by nearly a factor of four.[2] Also, the system's main memory capacity is now 493 Terabytes, which is quadruple the storage space the HLRS supercomputing infrastructure had previously.[2] "Users will benefit from the 5.4 Petabyte of file storage with an input/output speed in the range of 150GB/s," said Koller. "We are now able to meet the performance challenges of today's most demanding HPC users, from automotive and aerospace projects to medical, manufacturing and life science research."

Delivering customer satisfaction

HLRS has made the value of the new cluster clear for its users and will continue to do so.

RECOM SERVICES
REACTION & COMBUSTION MODELLING

Since customer satisfaction is so crucial to the team at HLRS, it regularly gathers feedback through user workshops to share experiences and apply any learnings made in one team across other projects. "A strong exchange of information is very valuable, especially in academic circles," said Koller. "Since new technologies often require users to recompile and optimize their codes to be able to take advantage of the improved technology, open communication is necessary, so we collaborate with our users to ensure that everyone feels the benefits of such an upgrade."

"Our team is very excited about using the new system at HLRS, since it gives us greater per node memory," said Ralf Schneider, senior researcher at HLRS. "Before, we had to switch to lower-resolution cluster systems due to storage limitations but now we can run all our simulations on the one supercomputer, which is much more efficient. We can get the best out of commercial off-the-shelf Intel technology thanks to how the platform is optimized."

Schneider is a permanent member of staff at The University of Stuttgart, working with structural mechanics simulations as part of an ongoing research project that is trying to achieve accurate simulation results for fractured bone, focusing on the thigh bone. His team is using the massive computing capacity at HLRS to simulate the pressures and forces a human femur endures to better understand the forces that this bone has to take in day-to-day life and then create a much better implant that is as light as possible and has the strength to handle all the same pressures as a natural bone. To develop material models for such bone remodeling simulations, these researchers need to predict bone remodeling by exploring how patients heal. For this, they need a huge amount of exact material data.

"This upgrade has sped up our entire process," said Schneider. "The higher throughput means that our turnaround times are much faster, which is critical as our development depends on it. In my world, one sample of data is two years of in-depth work. One 0.6mm cube of bone generates approximately 90,000 samples. For each of these samples, we carry out at least six Finite-Element simulations to get the field of anisotropic material data we need over the full femoral head. To carry out this large number of simulations, we definitely need the new HLRS supercomputer."

User benefits

Another satisfied user is Bendetto Risio, founder and president of RECOM Services. After completing his PhD at The University of Stuttgart, he founded RECOM, which does mathematic modeling to bridge the gap between scientific research and industrial problem-solving. Optimization is a big issue, and RECOM predicts organizational performance and generates efficient solutions. Usually, this would require a lot of testing, which would waste time and resources. However, with an accurate model, different systematic options can be demonstrated in a virtual world to gauge which delivers the best performance rate. Genetic algorithms can automatically generate different scenarios and then records results. Getting an accurate model takes thousands of simulations of different options, which calls for significant computational power. Thanks to an ongoing relationship with the university, RECOM can rent computing time to run simulations and serve significant projects much better than with conventional hardware.

"Without the HPC power we can access at HLRS, many problems we solve today would not be possible," said Risio. "Our solution strategies are strongly connected to the capabilities we can access, and we have noticed improvements since the upgrade. Since the new system is faster and has greater capacity, we can refine our models and generate even more accurate simulations with more data, increasing complexity rather than cutting run time."

One of RECOM's clients, a boiler manufacturer in Europe, needed to reduce its nitro oxide emissions to meet strict regulatory levels. Thanks to RECOM's accurate simulations, the team was able to predict how to retrofit the factory and introduce measures to limit the system and cut emissions to be within the specified limit. Without this simulation, the manufacturer would have just trusted its old burner technology and may have had to pay significant penalties. With a smart model and some powerful HPC machines, the problem was solved in just three months, which is half the time it takes to install a new boiler. "HLRS is essential for projects like this one," said Risio.

Lessons Learned

Due to this upgrade, HLRS is one of the foremost HPC sites in Germany. It has a more powerful offering and can process more data. The team is excited about the second phase of the installation, which is expected to boost performance even further, to seven petaflops. Moving forward, the organization hopes to take on more new projects to continue supporting the advancement of important scientific research, medical innovation, and industrial development.

Find the solution that's right for your organization. View **success stories from your peers**, learn more about **server products for business** and check out the **IT Center**, Intel's resource for the IT Industry.

[1] Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to http://www.intel.com/performance.

[2] HLRS Supercomputer "Hornet" Replaces "Hermit" press release from Gauss Centre for Supercomputing (July 2014), http://www.gauss-centre.eu/SharedDocs/Pressemitteilungen/GAUSS-CENTRE/EN/2014-07-Hornet.html

Intel does not control or audit the design or implementation of third party benchmark data or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmark data are reported and confirm whether the referenced benchmark data are accurate and reflect performance of systems available for purchase.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at http://www.intel.com

Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

*Other names and brands may be claimed as the property of others.

© 2015, Intel Corporation

0515/JNW/RLC/XX/PDF

332465-001EN