

In 2020, Cornelis Networks acquired the Omni-Path business of Intel Corporation, including the Intel® Omni-Path and Intel® True Scale products.


Cornelis Networks is now the manufacturer of [these products](#), renamed Omni-Path and True Scale.

As Cornelis Networks creates and transitions to its own written materials for the [products](#), you can reference the attached legacy materials.

All questions and product support requests should be directed to support@cornelisnetworks.com.

CASE STUDY

High Performance Computing (HPC)
2nd Generation Intel® Xeon® Scalable Processors
Intel® Omni-Path Architecture



Meeting Diverse Scientific and Research Needs

HLRN chooses 2nd Generation Intel® Xeon® Scalable processors to meet their increasingly diverse needs for high performance computing workloads

At a glance:

- HLRN was founded in 2001 by seven northern German states to support national science with computing and storage capacity.
- The new 16 PetaFLOP/s HLRN-IV system will be physically divided between the Zuse Institute Berlin, and the University of Göttingen. The sites are connected by a dedicated, redundant, 10 gigabit fiber optic cable.
- The system features the Intel® Xeon® Platinum 9200 processors for outstanding compute performance, memory bandwidth, and compute density on the nodes.
- The system features Intel® Omni-Path Architecture for high-performance data movement between nodes.



Atos

Executive Summary

HLRN supercomputers are used by over 100 universities and over 120 research institutions enabling exploration of the many frontiers of scientific research to help unlock a better future. The selection of Intel's latest processor technology to power the newest HLRN supercomputer came after detailed testing to find the best solution. Prof. Dr. Ramin Yahyapour of Göttingen University explains "the expectation for HLRN's supercomputer acquisition was to have a significant step up in computer power for new experiments."

Challenge

"Science in general is getting more compute and data intensive. This means that having larger systems available translates into an ability for the scientists to do better work. That's why HLRN is crucial for scientific research," says Prof. Dr. Ramin Yahyapour.

HLRN lays claim to being a very demanding client—HLRN has substantial expertise from their prior deployments of three supercomputer systems. Prof. Alexander Reinefeld from Zuse Institute Berlin emphasizes that "We are expecting the highest performance for all benchmark applications. Our benchmark suite was carefully chosen so that each code challenges specific parts of the system: CPU, communication network, and parallel I/O. We are not looking for peak theoretical performance—we demand real system performance which makes it more complicated for vendors to optimize their infrastructure for our applications. That meant that our selection of the right processor and the right interconnect are all crucial for the overall performance."

As with most research today, the need for more real-world computer capacity stems from the reality that simulations of many kinds are critical to the researchers. Faster computers are primarily used to increase the simulation in size and resolution—with the expectation of finding new discoveries.

"We demand real system performance... that meant that our selection of the right processor and the right interconnect are all crucial for the overall performance." —Prof. Reinefeld

Solution

HLRN procured a new supercomputer with just under a quarter of a million cores. The Intel® Xeon® Platinum 9200 processors (from the 2nd Generation Intel® Xeon® Scalable processor family) will serve as the "right processors" for HLRN. For the "right interconnect," HLRN chose Intel® Omni-Path Architecture. The system is produced by [Atos](#) (formerly Bull Computing) and will be physically split between the



Researchers at ZIB will use HLRN-IV for fluid dynamics, including developing turbulence models for aircraft wings.

Zuse-Institute Berlin (ZIB) and the Georg-August-Universität Göttingen (University of Göttingen). These sites have previously used this split system model, and already have in place a dedicated, redundant, 10 gigabit, fiber optic cable spanning the more than 170 miles between Berlin and Göttingen.

Result

HLRN has announced that the new system, HLRN-IV, will be approximately six times as fast as the prior systems—offering 16 PetaFLOP/s performance.¹ The excitement among researchers is palpable, and the list of research being done is mind-boggling. Prof. Reinefeld summed up his excitement saying, “It’s a great system. Our users will benefit right away from the more powerful system without needing to change their code. The homogeneous architecture of the 2nd Gen Intel Xeon Scalable processors will provide true performance portability which is a crucial aspect for our researchers in order to quickly benefit from the new, more powerful system.”

Key research areas within HLRN include:

- **Earth System Sciences**—which includes work on climate change. Subjects include the dynamics of oceans, rain forests, glaciers, Antarctic phytoplankton (microalgae), mineral dust cycles, and the stratosphere.
- **Fluid Dynamics**—which includes turbulence models for ship turbines, wind turbines, and aircraft wings. These models are notorious for needing enormous compute power—the acquisition of HLRN-IV will enable the running of more fine-grained turbulent simulations on large systems such as wind flow through a city, or across a blade on a turbine. Modeling complete cities will allow studies in how new buildings would change wind flow, and other factors that impact various microclimates within the city. This may lead to new design aspects to enhance city life. Other researchers hope to gain understanding that will pave the way for future high-lift commercial aircraft. Other researchers are hoping to save lives and ships by studying liquefaction of solid bulk cargo (such as iron ore or nickel ore). Failure to properly deal with this issue has led to the complete loss of at least seven vessels around the world in the past decade.
- **Health Care** is a broad area of research, and HLRN researchers hope to help in many ways including improving medical care at home. Gaining a better understanding of illness and treatment of diseases



At the University of Göttingen, research areas include collaborative projects on cellular and molecular machines.

stands to impact us all. Research includes simulations of drug efficacy, interactions, and side-effects. Enormous compute power allows leading researchers in these fields to start exploring the “personal medicine” aspects of these simulations, not just the average effects on a general population.

High Performance Across Diverse Research

In terms of science communities, HLRN has to support all types of workloads for their many researchers. Therefore, HLRN systems need to have the characteristics of a general-purpose system but still be of the highest performance. Their final choice had no accelerators.

“Although we looked at accelerators, including GPUs, as part of the procurement process, there was no advantage with regards to obtaining the highest performance in using GPUs or other accelerators in the system.”

— Dr. Thomas Steinke, Head of ZIB Supercomputing

HLRN’s benchmarks are open and include benchmarks that can take advantage of GPUs. HLRN found that any advantage in performance on some workloads are insufficient, when considering the reduction in general purpose compute capacity, or additional costs involved. A homogeneous system based on the 2nd Gen Intel Xeon Scalable processors proved itself to be the best choice for the diverse needs of the HLRN scientists and researchers.

Beating Back Amdahl's Law

Ever mindful of Amdahl’s Law, Dr. Thomas Steinke is fond of emphasizing the use of fast algorithms for fast computers. He shared that “The pressure of optimizing code for scaling on a node is less because of the high real-world performance of the 2nd Gen Intel Xeon Scalable processors compared to previous many-core architectures.”

The 2nd Gen Intel Xeon Scalable processor family offers an outstanding choice for HPC and helps programmers cope with Amdahl’s Law.

“Our users will benefit right away from the more powerful system without needing to change their code.”

—Prof. Reinefeld

Future of AI in HPC

AI and Machine Learning stand to impact all areas of HLRN research. A hot area of interest is the blending of machine learning and AI techniques with traditional simulation capabilities. While promising results have been reported, there is much work to be done. The exploration of algorithms is likely to take researchers in many directions, and this need for flexibility is one reason HLRN chose 2nd Gen Intel Xeon Scalable processors to support their next generation of research.

Avoid Data Movement

Prof. Yahyapour emphasized that “the CPU is quite good for artificial intelligence and machine learning. That’s an area where we see more need from our researchers. Traditionally they were not so much into data intensive work but that’s something we see as a new trend for the new system that will also be of particular interest.”

Intel® Advanced Vector Extensions 512 (Intel® AVX-512) proved to be the logical choice to help increase HLRN’s compute power, and with the addition of Intel® Deep Learning Boost (Intel® DL Boost) to augment AVX-512, offered outstanding performance for the new frontier of HPC applications.

The ability to compute data where it is, for all types of algorithms, saves data movement. That represents a boost for compute capacity, and less wasted energy. A double win!

When exploring new algorithms, and new application techniques, nothing is more important than the flexibility of a system. The 2nd Gen Intel Xeon Scalable processor delivers high performance coupled with the flexibility needed to meet future challenges.

Where to Get More Information


For more information on solutions for HPC from Intel, visit [intel.com/hpc](https://www.intel.com/hpc)

[View the HLRN Brings Advanced Performance to HPC video](#)

Solution Ingredients

[Find more information on the 2nd Generation Intel® Xeon® Scalable Processor](#)

[Learn more about Intel® Omni-Path Architecture](#)



¹The previous system, HLRN-III, consists of two complexes located at ZIB in Berlin and Leibniz Universität IT Services (LUIS) in Hannover coupled by a dedicated 10GigE fiber optic connection for HLRN to provide for the so-called single-system view. Delivered in two phases, the compute node details include: Phase one included two Cray XC30 computers each consisted of 744 compute nodes with a total of 1488 dual-socket Intel Xeon processors E5-2695v2 with a total of 93 TB main memory, which are connected via a fast Cray Aries network with Dragonfly topology. The second phase added 2064 Intel Xeon processor E5-2680 v3 compute nodes with 85248 compute cores with 1872 compute nodes in Berlin and 1680 compute nodes in Hannover totaling 2.7 PetaFlops / s Peak Performance and 222 TB main memory extended.

Intel technologies’ features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at [intel.com](https://www.intel.com).

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 complete information visit www.intel.com/benchmarks.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

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

© Intel Corporation