

Building  
the Future,  
Together

H L R I S

High-Performance Computing Center Stuttgart



2025/26  
Annual Report



# 2025/26

## Annual Report

The High-Performance Computing Center Stuttgart (HLRS) was established in 1996 as Germany's first national high-performance computing (HPC) center. As a research institution affiliated with the University of Stuttgart and a founding member of the Gauss Centre for Supercomputing, HLRS provides comprehensive HPC services to academic users, industry, and the public sector. HLRS operates one of Europe's most powerful supercomputers, provides advanced training in HPC programming and simulation, and conducts research to address key problems facing the future of supercomputing. Among HLRS's areas of expertise are parallel programming, numerical methods for HPC, visualization, grid and cloud computing concepts, data analytics, and artificial intelligence. Users of HLRS computing systems are active across a wide range of disciplines, with an emphasis on computational engineering and applied science.

In 2025, HLRS added solar panels to the roof of its existing building. The facade of the center's new building, HLRS III, will also include photovoltaic capabilities.

# Director's Welcome

## Grußwort



Prof. Dr.-Ing. Michael M. Resch, Director, HLRS

Looking back at another successful year at the High-Performance Computing Center Stuttgart, one day in particular stands out as a highlight. It was mid-December 2025, and in a season that is typically cold and gray, the temperature outside was pleasant and the sun shone brightly in a cloudless sky. For the many friends and honored guests of HLRS who had gathered, the conditions were a spectacular backdrop for celebrating the cornerstone laying for HLRS III, a building that will soon house HLRS's next-generation high-performance computing (HPC) systems. The ceremony was in many ways not just about a building, but rather a transformative vision that in the past 12 months has become much more concrete, both figuratively and

Wenn ich auf das vergangene Jahr am Höchstleistungsrechenzentrum Stuttgart zurückblicke, bleibt mir ein Tag besonders im Gedächtnis. Mitte Dezember 2025: ungewöhnlich mild, heller Sonnenschein, ein wolkenloser Himmel. Für unsere Gäste bot dieses Wetter eine spektakuläre Kulisse für die feierliche Grundsteinlegung für HLRS III – das Gebäude, in dem bald unsere nächste Generation von Höchstleistungsrechnern stehen wird. Doch an diesem Tag ging es um mehr als Beton und Technik. Es ging um eine Vision, die im Laufe des Jahres greifbar geworden ist. In diesem Bericht zeigen wir Ihnen, welche Schritte wir 2025 gegangen sind und wie sie unsere Arbeit im kommenden Jahrzehnt prägen werden.

literally. In this annual report, we invite you to learn about some important steps that HLRS took in 2025 that will define its future for the coming decade.

In late 2025, for example, we announced the details of Herder, our next supercomputer. When it is installed in HLRS III in 2027 it will offer a significant leap in computational performance, opening exciting new possibilities for research and technology. To prepare for its arrival, one of our most important tasks this year has been to support users of our current Hunter system to port their codes to GPUs. This will make it possible to take advantage of Herder's formidable capabilities when it arrives.

Another key activity this year involved serving as coordinator of HammerHAI, an important project selected by the EuroHPC Joint Undertaking (JU) as Germany's first AI Factory. We worked closely with our consortium partners throughout 2025 to design a service portfolio that will accelerate AI-powered innovation, and to connect with AI-focused stakeholders and initiatives from across Baden-Württemberg. In addition, the JU has now announced the results of the procurement for a new, AI-optimized HammerHAI supercomputer, which will be installed at HLRS later in 2026. Already we are seeing great anticipation for these developments from across the German research, industrial, SME, and start-up communities, and we look forward to seeing how they will use these powerful, sovereign AI capabilities.

The rise of GPU-accelerated systems for simulation and artificial intelligence exemplified in Herder and HammerHAI has changed HPC dramatically in recent years. But what will come next? Quantum computing, neuromorphic computing, and other emerging emerging technologies hold enormous potential to improve the performance and energy efficiency of simulation applications. To prepare, HLRS this year founded a Future Computing Group, which will collaborate with developers of such technologies to test and evaluate their capabilities, particularly in light of our users' computing requirements. Regardless of what the future holds, this

Ein wichtiger Meilenstein war die Vorstellung der Details von Herder, unserem kommenden Supercomputer. Wenn wir ihn 2027 im HLRS III in Betrieb nehmen, wird er unsere Rechenleistung deutlich steigern und neue Chancen für Forschung und Industrie eröffnen. Damit wir diese voll ausschöpfen können, haben wir im Jahr 2025 unsere Nutzer:innen intensiv dabei unterstützt, ihre Codes auf GPUs zu portieren. So werden sie in der Lage sein, Herders Fähigkeiten vollumfassend zu nutzen, sobald der Supercomputer verfügbar ist.

Eine weitere Neuigkeit im Jahr 2025 war der Beginn der Koordination von HammerHAI – einem bedeutenden Projekt, das von der EuroHPC Joint Undertaking (JU) als Deutschlands erste AI Factory ausgewählt wurde. Gemeinsam mit unseren Partnern haben wir ein Serviceangebot entwickelt, das Innovationen mit künstlicher Intelligenz beschleunigt. Gleichzeitig knüpfen wir Kontakte zu Unternehmen, Start-ups und Forschungseinrichtungen in Baden-Württemberg und bauen bestehende Beziehungen weiterhin aus. Inzwischen steht auch fest: Ein neuer, KI-optimierter HammerHAI-Supercomputer wird 2026 am HLRS installiert. Die Erwartungen daran sind hoch – in der Forschung ebenso wie in der Industrie – vor allem vonseiten Start-ups sowie KMUs.

Das Aufkommen GPU-beschleunigter Systeme für Simulationen und künstliche Intelligenz, beispielsweise bei HammerHAI und in Zukunft bei Herder, hat das Höchstleistungsrechnen in den letzten Jahren grundlegend verändert. Doch die Entwicklung geht weiter: Quantencomputing, neuromorphes Computing und weitere neue Technologien versprechen zusätzliche Leistung und mehr Energieeffizienz. Um diese Technologien früh zu prüfen, haben wir eine Future Computing Group gegründet, die eng mit Entwicklern zusammenarbeitet und neue Ansätze im Hinblick auf die Anforderungen unserer Nutzer:innen testet. So stellen wir sicher, dass das HLRS auch künftig Zugang zu den leistungsfähigsten Werkzeugen bietet.

Technik allein reicht allerdings für Fortschritte im Höchstleistungsrechnen nicht aus – sie entstehen nur

effort places us in an optimal position to continue providing access to the most advanced computing tools and resources.

Success in high-performance computing is not simply a matter of technology, though. As the motto of this annual report suggests – “Building the Future, Together” – great things are only really possible through cooperation. At HLRS, the impulse to collaborate with others at the state, national, European, and international levels has for decades been deeply embedded in everything we do. In this annual report we describe why such partnerships are so important, and the kinds of impacts they can have. I am also delighted that we were able to include an interview with Steve Conway, a widely respected HPC industry analyst and longtime friend of HLRS. Providing insights into issues that are currently shaping the evolving HPC landscape, Steve also speaks to the importance of cooperation in Europe’s impressive expansion of HPC capabilities, remarking that this achievement is the most exciting development that he has witnessed in his long and accomplished career. I am proud of HLRS’s contributions to these efforts through its ongoing leadership within EuroCC, CASTIEL, and other European projects.

Whether it be managing the effects of climate change, finding new strategies for fighting disease, or developing the high-tech products that will enable Germany to remain competitive globally, simulation and data science methods offer absolutely essential tools for addressing major challenges. One outstanding example is the work of Everllence, one of HLRS’s more than 75 industrial HPC users. A manufacturer of large-scale industrial machinery like engines for supertankers and facilities for gas and oil refineries, Everllence aims to make major contributions to global decarbonization efforts by enabling the transition to alternative fuels, and making carbon capture and storage practical. Elsewhere in this annual report we also highlight interesting scientific accomplishments that benefited from access to HLRS’s computing systems. These include exciting new approaches that combine machine learning and simulation to gain insights into human physiology and

mittels Zusammenarbeit. Unter dem Motto „Gemeinsam die Zukunft bauen“ zeigen wir in diesem Jahresbericht, wie wichtig Partnerschaften für unsere Arbeit sind – in Baden-Württemberg, Deutschland, Europa und weltweit. Besonders freue ich mich über ein Interview mit Steve Conway, einem anerkannten Analysten der HPC-Branche und langjährigen Freund des HLRS. Als Begleiter der HPC-Entwicklung seit vielen Jahren und Trendspotter in der Industrie ist seine Einschätzung eindeutig: Der Ausbau der europäischen HPC-Infrastruktur gehört zu den spannendsten Entwicklungen seiner Karriere. Ich bin stolz darauf, dass das HLRS viel dazu beiträgt – etwa dank seiner führenden Rolle in EuroCC, CASTIEL und anderen europäischen Projekten.

Ob es darum geht, die Auswirkungen des Klimawandels zu mindern, Krankheiten zu bekämpfen oder die globale Wettbewerbsfähigkeit der deutschen Industrie zu stärken – Simulations- und datenwissenschaftliche Methoden bieten unverzichtbare Werkzeuge zur Bewältigung großer Herausforderungen. Ein herausragendes Beispiel hierfür ist die Arbeit von Everllence, einem aus mehr als 75 Industriekunden des HLRS. Als Hersteller von großen Industriemaschinen wie Motoren für Supertanker und Anlagen für Raffinerien möchte Everllence einen wesentlichen Beitrag zur globalen Dekarbonisierung leisten. Mithilfe von Rechenressourcen des HLRS möchte das Unternehmen den Übergang zu alternativen Kraftstoffen ermöglichen und die Speicherung von Kohlenstoffdioxid in die Praxis umsetzen. Auch in der Wissenschaft sehen wir beeindruckende Fortschritte: Forschende kombinieren maschinelles Lernen mit Simulation, unter anderem um den menschlichen Körper besser zu verstehen. Andere Projekte zeigen großes Potenzial, den Betrieb von Satelliten in niedrigen Umlaufbahnen zu verbessern oder liefern neue Einblicke in Turbulenzen.

Auch Nachhaltigkeit bleibt für uns weiterhin ein zentrales Thema. Im Zuge des Neubaus HLRS III entsteht derzeit eine Anlage, die die Abwärme des kommenden Rechners Herder nutzen wird. Sie wird den Campus Vaihingen der Universität Stuttgart mit Wärme versorgen und so zur Dekarbonisierung beitragen. Zudem haben

knowledge that will make the operation of satellites in very low earth orbits more economically sustainable, as well as a large-scale simulation that offered the first full-resolution dataset for studying turbulence as it moves from low to high Reynolds numbers.

At HLRS, too, an important part of our vision for the future is to maximize our sustainability. As part of the HLRS III project, construction has begun on our waste heat management facility, which will supply heat produced by Herder to the University of Stuttgart’s Vaihingen campus, supporting its decarbonization efforts. We also successfully completed our third EMAS certification, and as this annual report explains, are increasingly becoming a resource for other institutions in search of advice on how to operate more sustainably.

In his public remarks at the HLRS III cornerstone ceremony, Baden-Württemberg Minister of Finance Dr. Danyal Bayaz called the event a “sign of hope that the best is still ahead of us, that we don’t sit satisfied in a corner, but continually create the conditions to reinvent ourselves technologically.” At a time when the world we thought we knew seems to be shifting under our feet, this optimism about the potential of high-performance computing to help us navigate these difficult times is something that motivates all of us in the HPC community.

Which brings me back to our theme: Building the Future, Together. It is clear to me every day that HLRS would not be what it is without our large and engaged community of funders, partners, and system users, as well as our skillful and dedicated staff. Thank you all for your ongoing support and commitment in our efforts. Let’s keep doing great things together.

With best regards,



Prof. Dr.-Ing. Dr. h.c. Prof. E.h. Michael M. Resch  
Director, HLRS

wir unsere dritte EMAS-Zertifizierung erfolgreich abgeschlossen. Aufgrund unserer langjährigen Erfahrung mit umweltschonendem Rechenzentrumsbetrieb wenden sich immer häufiger andere Einrichtungen für fachliche Beratung an uns.

Bei der Grundsteinlegung bezeichnete der baden-württembergische Finanzminister Danyal Bayaz das HLRS III als „ein Zeichen auf Hoffnung, dass das Beste noch vor uns liegt und wir nicht satt und zufrieden in der Ecke sitzen, sondern immer wieder den Anspruch und die Ambition haben, uns technologisch neu zu erfinden.“ Dieser Gedanke begleitet mich. Wir leben in einer Zeit des Umbruchs. Vieles verändert sich schneller, als wir es gewohnt sind. Umso wichtiger ist es, dass wir die Chancen neuer Technologien nutzen. Höchstleistungsrechnen kann uns dabei helfen, komplexe Probleme zu verstehen und Lösungen zu entwickeln.

Damit schließt sich der Kreis zu unserem Motto: Gemeinsam die Zukunft bauen. Das HLRS lebt von seinen Partnern, seinen Nutzer:innen und vor allem von seinen Mitarbeitenden. Ihnen allen danke ich herzlich für Ihr Engagement und Ihre Unterstützung. Lassen Sie uns diesen Weg gemeinsam weitergehen.

Mit freundlichen Grüßen

Prof. Dr.-Ing. Dr. h.c. Prof. E.h. Michael M. Resch  
Direktor des HLRS

2	Director's Welcome
<b>8</b>	<b>News Highlights</b>
9	Cornerstone Laid for New Building at HLRS
12	AI Factory HammerHAI Off to a Strong Start
15	Anticipation Builds for HLRS's Next Supercomputer, Herder
17	Future Computing Group Will Evaluate New Technologies for HPC
20	Everllence Uses HLRS Supercomputers to Accelerate Industrial-Scale Decarbonization
23	Supercomputing Meets the Arts
24	Code Porting Workshops Support Transition to GPUs
25	Good Advice from Experts in Green IT
27	Sustainability Highlights 2025
<b>28</b>	<b>Building the Future, Together</b>
29	Collaboration Drives Progress in Supercomputing
34	International Partnerships
36	Strength through Cooperation: An Interview with Steve Conway
<b>40</b>	<b>News Briefs</b>
<b>48</b>	<b>User Research</b>
49	Particle Scattering Model Could Improve Low-Orbit Spaceflight
51	Large-Scale Simulation Reveals Novel Insights about Turbulence
54	Evolutionary Molecular Dynamics Delivers Insights into Cholesterol Attraction
57	Selected Publications by Our Users in 2025
<b>64</b>	<b>HLRS by Numbers</b>
<b>66</b>	<b>About Us</b>
67	Inside Our Computing Room
69	User Profile
70	Third-Party Funded Research Projects
74	HPC Training Courses
76	Workshops and Conferences
77	Organization Chart
78	Departments



# News Highlights



Architectural rendering of HLRS III.

## Cornerstone Laid for New Building at HLRS


**HLRS III is the foundation for the future of HPC in Stuttgart. It will enable significant advances in computing power and energy efficiency.**

In March 2025, construction began on a major new addition to HLRS. The new data center, called HLRS III, will house two new large-scale computers: HLRS's next flagship supercomputer, called Herder (see page 15), and an AI-optimized supercomputer to be installed in conjunction with the AI Factory HammerHAI (see page 12). Scheduled to open in 2027, HLRS III is being built to accommodate the demanding technical requirements of future high-performance computing (HPC) systems, while also using energy as efficiently as possible. The new facility will enable Stuttgart to continue playing a leading role in European high-performance computing.

By December 15, 2025, HLRS III's foundation and parts of its basement were already in place, providing an inspiring backdrop for a festive ceremony that celebrated the laying of its cornerstone. Among the speakers at the event was Baden-Württemberg Finance Minister Danyal Bayaz, who in a press release said, "We are constructing a state-of-the-art building, promoting innovation and advanced research in the state, and investing in our future. This will create space for a new generation of computers. HLRS III in Stuttgart will be a top-class facility in Europe and a major European cluster for high-performance computing and artificial intelligence."



A processing facility located next to HLRS III will capture and distribute waste heat produced by HLRS supercomputers to the University of Stuttgart's district heating network.



The "Canyon" will offer views into the HLRS III computer room, while enabling passive cooling in the warmer months.

HLRS III will expand on existing HLRS facilities to provide an additional 7,000 m<sup>2</sup> of new space. A hybrid construction using reinforced recycled concrete, steel, and wood, the building will contain a state-of-the-art computer room with a redundant power supply, a highly energy-efficient cooling infrastructure, and a building structure capable of supporting the higher weight of future supercomputers.

"The construction of HLRS III marks the beginning of a new stage in HLRS's history, and affirms our HPC center's important role in powering research and innovation in Germany and across Europe," said Prof. Dr. Michael Resch, Director of HLRS. "The new building is a critical piece of infrastructure that is necessary for us to be able to provide our scientific and industrial user communities with access to the next level of supercomputing power. At the same time, it will allow us to operate our systems sustainably."

### Energy efficiency and waste heat reuse

The increased computational capabilities of Herder and other future HPC systems will require more energy than has been available to HLRS in the past. For this reason, maximizing energy efficiency is an essential element of the HLRS III building concept. All of the electricity consumed by supercomputers is transformed into heat, which must be removed from the machine to prevent it from overheating. To take advantage of this valuable resource, a new waste heat processing facility also being constructed next to HLRS III will capture all waste heat generated by the center's supercomputers and distribute it to the district heating network of the University of Stuttgart's Vaihingen campus. Highly efficient, large-scale industrial heat pumps using a climate-neutral refrigerant will be used for this purpose. This approach will cover up to one-third of the campus's heating needs during winter months and all of its process heat requirements in the summer.

University of Stuttgart Rector Prof. Peter Middendorf highlighted the significance of HLRS III for the entire university community saying, "The visionary new construction of HLRS III is a milestone for the digital future of the University of Stuttgart, and is an excellent example of how we bring cutting-edge technology and sustainability together. The supercomputer Herder will not only make high-performance science possible, but will also actively contribute to the decarbonization of our campus."

When used to its full capacity, it is estimated that the HLRS III waste heat processing facility could reduce CO<sub>2</sub> emissions on campus by nearly 50%.

Sustainable technologies on the facades and roofs of the two buildings will complement the high-performance computers contained within. The south and east sides of HLRS III will be covered with photovoltaic panels to generate electricity.

Although HLRS III's computer room will be securely separated from public areas of the building, an area called the "Canyon" will offer observation decks overlooking the computer room. Ventilation in the ceiling of the canyon will permit passive cooling of the building during warm summer months.

HLRS III has been identified as a component within the State of Baden-Württemberg's Green IT strategy. HLRS received the German Environmental Management Award 2025 (see page 27) in recognition of the particularly energy-efficient and sustainable operation of its computers. In 2024, the center also won a Datacenter Strategy Award for "Transformation," highlighting the sustainable measures being implemented in HLRS III.

### Funding and planning

The total cost of constructing HLRS III, including the waste heat processing center and underground piping needed to deliver heat to the campus district heating network, is €178.6 million. Funding of €138.5 million has been provided by the State of Baden-Württemberg, with the balance of €40.1 million coming from the German Federal Ministry for Research, Technology, and Space.

HLRS III was designed by the architecture firm Benthem Crowell International in Diemen, Netherlands, with building planning being coordinated by Birk Heilmeyer und Frenzel Architekten from Stuttgart. The University Construction Office of Stuttgart and Hohenheim is responsible for the project and the building's future maintenance.



Baden-Württemberg Finance Minister Danyal Bayaz placed a time capsule inside the cornerstone to HLRS, which was then lowered by crane into the construction site. Also speaking at the event were (l-r): Prof. Michael Resch (Director, HLRS), Prof. Peter Middendorf (Rector, University of Stuttgart), and Andreas Hölting (Director, Baden-Württemberg Agency for Property and Construction).

# AI Factory HammerHAI Off to a Strong Start

**HLRS and its partners in the HammerHAI consortium established a service portfolio, developed partnerships, and completed procurement of a new AI-optimized supercomputer, laying a foundation that will significantly improve Germany's AI capabilities in industry and research.**

In December 2024, the EuroHPC Joint Undertaking (JU) announced the establishment of HammerHAI, Germany's first AI factory. Coordinated by the High-Performance Computing Center Stuttgart, HammerHAI was one of the first seven AI factories in this landmark effort to establish a public, coordinated, large-scale European AI infrastructure. In 2025 the initiative grew to include 19 AI Factories across Europe, a network that holds great potential to promote European digital independence.

Officially launched in 2025, HammerHAI has already begun providing computing infrastructure, expert support, and training to facilitate the development of artificial intelligence-powered applications, business concepts, and research. The AI Factory is concentrating in particular on supporting SMEs and startups in the fields of engineering, manufacturing, automotive and

mobility, as well as academic science. Operated in compliance with the European Union's AI Act, the AI Factories will provide a secure, reliable platform for European industry and research.

Throughout 2025 HLRS worked with its partners in the HammerHAI consortium – the Leibniz Supercomputing Centre, Karlsruhe Institute of Technology, Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen, and SICOS BW – to begin the transformative process of launching the AI factory. This included beginning development of a service portfolio to address the needs of its future user community, and defining operational procedures that will enable it to meet its users' demands efficiently. In parallel, procurement was completed for a new HammerHAI supercomputer that will be installed at HLRS, and the AI Factory conducted extensive outreach to connect with potential users and



HammerHAI is coordinated by HLRS in cooperation with the Leibniz Supercomputing Centre, Karlsruhe Institute of Technology, Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen, and SICOS BW.

partners, building a network that will propel the new system into service as soon as it arrives.

## AI-optimized supercomputer will be installed in Stuttgart

The EuroHPC Joint Undertaking signed a contract for the HammerHAI supercomputer, which will be installed at HLRS in the second half of 2026. The system will be manufactured by HPE, based on the liquid-cooled NVIDIA GB200 NVL4 architecture. Combining NVIDIA Grace CPUs with NVIDIA Blackwell GPUs and scaled with NVIDIA Quantum-X800 InfiniBand networking, the NVIDIA GB200 NVL4 by HPE will offer more than 15 Exaflops of peak AI inference performance. It will integrate the VAST Data DASE storage architecture, which provides a unified data platform for AI and HPC workloads, as well as a partition based on AI-optimized inference engines and hardware accelerators from Netherlands-based Axelera AI. The HPE Morpheus Enterprise software will be used as a unified AI control plane, enabling automated provisioning, governance, and workload lifecycle management.

This configuration will make the HammerHAI supercomputer a powerful tool for handling medium- to large-scale AI workloads for machine learning and artificial intelligence. It will incorporate a cloud-native software stack familiar to the AI community, making it straightforward to migrate or scale applications from local systems or commercial cloud environments. The system will be configured to support research and technology development in disciplines prioritized with the HammerHAI consortium, with an emphasis on engineering, manufacturing, automotive and mobility.

## Service portfolio will accelerate uptake of AI capabilities

In parallel with the procurement, HammerHAI began developing a service portfolio aimed at supporting research and industry. The portfolio includes resources developed and hosted by partners within the HammerHAI consortium, including a selection of software developed at the AI Service Center for Sensitive and Critical Infrastructures (KISSKI) that offers do-

mestic, privacy-protecting alternatives to offshore commercial tools. These include a privacy-oriented chatbot (ChatAI), applications for image generation (ImageAI), voice transcription and translation (VoiceAI). KIT provided access to edge.ml, a tool for deploying and monitoring machine learning models on edge devices. HammerHAI will also offer access to brokered services maintained by partners of the AI Factory, such as a facility for digital twin object recognition and movement detection located at ARENA2036, a laboratory for advanced automotive research.

In addition, HammerHAI began offering services for skill building and networking that are designed to lower barriers to entry for companies interested in developing new AI products and services. These included workshops to help companies identify potential in their data sets, develop roadmaps toward the implementation of AI-enabled applications and workflows, understand the requirements and implications of the European AI Act, and mitigate the ethical, legal, organizational, and technical risks inherent in AI projects within organizations.

## Building networks across the AI community

Representatives of HammerHAI were also very active in 2025 in conducting outreach to promote awareness of the AI Factory's capabilities across German industry, and to begin building contacts with potential partners and users. Throughout the year, HammerHAI maintained a strong presence at leading industry conferences and conventions, such as the European Big Data Value Forum 2025, the Automotive Testing Expo 2025, the EuroHPC User Days, ISC High Performance, SC25, and the 4<sup>th</sup> Smart Data Innovation Day.

HLRS and HammerHAI have also been actively integrating themselves into the German AI landscape, developing partnerships with other AI-focused organizations in ways that will build expertise and increase impact. For example, HLRS joined the German AI Association, the country's largest networking and advocacy organization for the artificial intelligence community. HammerHAI joined the Baden-Württemberg

“AI Landbox,” and will make its supercomputer available within a collaboration involving ARENA2036, Cyber Valley GmbH, and other stakeholders that plans to develop a collaborative infrastructure for accelerating AI innovation across the state. The AI Factory is also participating in the Manufacturing and Robotics Data Lab, an initiative that will pool data securely and make it available for the training of foundation models for automating manufacturing. And finally, HLRS joined the European Connected and Autonomous Vehicle Alliance (ECAVA), an EU-wide initiative that is bringing together stakeholders from across the automotive industry to promote digital innovation. These and other activities will ensure that HammerHAI becomes a key resource for AI capabilities across Germany.

#### Partnership with University of Edinburgh establishes UK AI Factory Antenna

In a further extension of the European AI Factories network, the EuroHPC Joint Undertaking in October 2025 announced the creation of 13 “AI Factory Antennas.”

This new category of competence centers for artificial intelligence will complement the activities of AI Factories across Europe.

Among the winning proposals was the UK AI Factory Antenna (UKAIFA), led by EPCC, a UK national high-performance computing center based at the University of Edinburgh. UKAIFA will work in partnership with HammerHAI, which will serve as a bridge between the United Kingdom and the European AI Factory network. The partners will also collaborate on a variety of efforts to accelerate the adoption of artificial intelligence in SMEs, startups, industry, and the public sector. Under the AI Antenna agreement, HammerHAI will provide access for the UKAIFA user community to the new HammerHAI supercomputer and will work with EPCC to enable secure, federated data sharing and cross-platform compatibility.

**More information:**  
To learn more about HammerHAI, visit [www.hammerhai.eu](http://www.hammerhai.eu).

## Anticipation Builds for HLRS’s Next Supercomputer, Herder

**Based on next-generation technologies from HPE and AMD, HLRS’s upcoming flagship supercomputer will be more than seven times faster than its current system and offer improved capabilities for simulation, artificial intelligence, and converged computing.**

In December 2025, HLRS announced the architecture of its next supercomputer. Developed with its technology partners, HPE and AMD, the new system – called Herder – will offer both a major increase in performance and an optimized architecture for today’s most advanced computational applications. When it arrives in 2027, Herder will give German and European research and industry powerful new capabilities for scientific discovery and innovation.

Prof. Michael Resch, Director of HLRS, celebrated the announcement, saying, “Herder will offer HLRS’s scientific and industrial user communities state-of-the-art technologies for both high-performance simulation and artificial intelligence. Combining these capabilities on a single system will also support some of today’s most exciting computational research, which is bringing numerical simulation and AI together in extremely interesting ways. Herder will enable our users to simulate

systems at a degree of precision and at a scale that is unprecedented in Stuttgart, while also offering local, sovereign capacity for AI model training and generative AI. We are excited about the kinds of results that it will make possible.”

#### The foundation: the HPE Cray Supercomputing GX5000

Herder will be based on the HPE Cray Supercomputing GX5000 system. Announced by HPE in November 2025, it is designed to support the kinds of large-scale numerical simulations that are essential within HLRS’s traditional user communities, as well as data-science approaches for AI model training and generative AI. Bringing these capabilities together within a single system will enable the development of new kinds of converged computing workflows that seamlessly integrate the two approaches.



#### EU Digital Commissioner Henna Virkkunen Provides Outlook on AI Strategy

On July 7, 2025, HLRS welcomed Henna Virkkunen, European Commission Executive Vice-President for Tech Sovereignty, Security, and Democracy. The visit took place as part of her tour of Europe’s new AI Factories. Commenting on this new initiative, Virkkunen said, “AI Factories like HammerHAI will be critical in fostering a German and European ecosystem for generative AI and a key asset for developing European competitiveness in AI and strengthening Europe’s digital sovereignty.” During her visit, Virkkunen also provided an outlook on the European Union’s plans to expand on the AI factory model and create future EU “Gigafactories” for artificial intelligence. AI Gigafactories will provide extreme-scale capabilities on par with the world’s largest AI-optimized supercomputers, enabling the development, training, and deployment of next-generation AI models and applications.



Rendering of the HPE Cray Supercomputing GX5000. HPE’s direct liquid cooling technology will efficiently maintain safe operating conditions for high-powered processors from AMD. Image courtesy of HPE.



Representatives of the University of Stuttgart and HPE signed the contract for Herder in Barcelona. (l-r): Antonio Neri (President and CEO, HPE), Anna Steiger (Chancellor, University of Stuttgart), Prof. Dr. Michael Resch (Director, HLRS), and Marc Fischer (Managing Director and Vice President Global Sales Germany, HPE).

Herder will incorporate the HPE Supercomputing Management Software, which supports the containerized methods needed for HPC / AI workflows. The new supercomputer will use the Lustre-based HPE Cray Supercomputing Storage Systems E2000 to alleviate bottlenecks in the management of large datasets and improve overall system performance. HPE Slingshot 400 switches will move data across the system at high speeds and with low latency.

When it goes into service in late 2027, Herder will replace HLRS's current flagship supercomputer, called Hunter. It will also be in good company among the world's most advanced high-performance computing systems. The HPE Cray Supercomputing GX5000 is the same platform being used in the upcoming Discovery, Mission, and Vision supercomputers in the United

States and the Blue Lion system at Leibniz Supercomputing Centre near Munich.

#### Faster GPUs, denser CPUs from AMD

Herder will contain next-generation processors from AMD, including the AMD Instinct™ MI430X GPU and AMD EPYC™ “Venice” CPU. Each MI430X supports 432GB of HBM4 memory and 19.6TB/s of memory bandwidth, offering powerful capabilities for data-intensive operations in both HPC and AI. The Venice CPU is the world's first processor to use TSMC's (Taiwan Semiconductor Manufacturing Company's) 2-nanometer fabrication methods, which make it possible to pack more transistors onto each computer chip than was possible in the past. By bringing the transistors closer together, information can be moved through the system more quickly.

The increased memory and communications density that these processors offer in comparison to current technologies will accelerate performance to new heights. HLRS anticipates that Herder will achieve a peak performance of more than seven times the capabilities of Hunter.

#### Energy efficiency and sustainability

This jump in performance means that Herder will use substantially more power than Hunter, and so energy efficiency and sustainability have been key considerations in its planning. The HPE Cray Supercomputing GX5000 integrates several technologies to maximize energy efficiency. It uses direct liquid cooling, which efficiently removes heat from the system and is critical for the reliable operation of the dense AMD MI430X and Venice processors. Working together with HPE, HLRS also plans to implement a dynamic power management approach called Powersched on the new system. HLRS received an Environmental Management Prize in November 2025 (see page 27) for its collaboration with HPE on the development and testing of this approach. Installed in the new datacenter HLRS III (see page 9), Herder's waste heat will also be captured and delivered to the University of Stuttgart's district heating network, contributing to its decarbonization efforts.

## Future Computing Group Will Evaluate New Technologies for HPC

A new HLRS initiative is planning collaborations that will investigate the potential of quantum, neuromorphic, and other emerging computing architectures to accelerate its users' workflows.

In recent decades, advances in hardware design have enabled today's leading supercomputers to achieve incredible gains in speed and efficiency, reaching the exascale in 2022. Historically this progress has followed Moore's Law, which in the 1970s predicted that the number of transistors on a microchip would double approximately every two years, leading to consistent increases in computing power and efficiency. Recent advances have been achieved using graphics processing units (GPUs) as accelerators, which can be found in most of the fastest supercomputers on the current Top500 List.

In the coming years, however, experts anticipate that the steady performance gains we have come to expect under Moore's Law will stagnate. Microchip fabrication can now be done on the astonishing scale of 2 nanometers, meaning that miniaturization is approaching its

physical limits. And although it would be possible to build faster GPUs, this will only be commercially viable by pricing them more expensively, making large-scale systems for high-performance computing (HPC) increasingly unaffordable.

As a result, the computer engineering community continues to innovate. In parallel with the rise of GPUs, hardware companies large and small have been developing promising new paradigms that are either fundamentally different than those behind traditional CPU-based and GPU-accelerated systems, or that build on the basis of more established architectures in new ways. Although many of these approaches have not yet fully matured, they hold the potential to solve at least certain kinds of computational problems faster and using much less energy than even today's best supercomputers.



Led by Dr. Johannes Gebert (second from left), the Future Computing Group brings together a dozen specialists in HPC hardware and software from across multiple HLRS departments.

As HPC centers like the High-Performance Computing Center Stuttgart look towards the next generation of supercomputing, determining which of these new technologies to adopt, and when, is extremely important. This is because investing in a new technology is not purely a financial decision, but can lock a center into a specific technological trajectory that will determine its relevance as HPC continues to evolve.

In 2025, HLRS launched a new initiative called the Future Computing Group to explore which emerging technologies for high-performance computing hold the greatest potential to address the needs of HLRS's scientific and industrial user communities. Working closely with technology companies, the group will also provide HPC expertise that will support industrial research on next-generation computing hardware.

According to Future Computing Group leader Dr. Johannes Gebert, "We are open to collaborating with companies and to evaluate their computer paradigms. For example, we need to understand whether the hardware can solve the kinds of problems our user community has. This means answering many questions. Is the new hardware relevant for certain algorithms and software packages? And do the advantages that it might offer outweigh those of other, more established technologies and companies?" As computing architectures become more heterogeneous, the Future Computing Group will help evaluate which technology combination best meets HLRS's needs.

#### Keeping up with the race for improved HPC technologies

HLRS is scheduled to install its next supercomputer, called Herder, in 2027. Like other world-class supercomputers, Herder will use GPU accelerators. What comes after Herder, however, is still an open question, and will in many ways depend on how high-performance computing hardware and user demands evolve in the coming years.

As the supercomputing community anticipates a future in which performance gains result not exclusively from

the use of GPUs, one potential addition to the portfolio of HPC technologies is quantum computing. Modest quantum systems based on a variety of architectures are already in operation, and as research continues, it is possible that more than a single architecture could prove to be reliable, versatile, and easy to use in everyday practice. Another emerging technology is neuromorphic computing, whose architecture is inspired by the human brain. This approach promises to calculate at similar speeds to those of today's fastest systems while using a small fraction of the energy. Practical implementation of neuro-inspired computing in HPC settings is still some time away, though. Meanwhile, a variety of other interesting, if less radical, technologies are under development with the goal of achieving performance and efficiency breakthroughs using more traditional architectures.

In the coming years, HLRS's Future Computing Group will evaluate the maturity of such architectures, research the performance of algorithms on these technologies, and develop software. By publishing the results of this work, the group will support scientific and industrial communities in benefiting from post-Moore hardware.

#### Collaboration will investigate a new microchip concept for HPC

In November 2025 the Future Computing Group launched the first of what Gebert anticipates will be a number of technology evaluation partnerships. HLRS has begun collaborating with Openchip, an international company founded and headquartered in Barcelona that focuses on the design and development of high-performance accelerator chips, platforms, and software stacks for advanced computing applications, based on the global RISC-V architecture standard. Using an approach to acceleration that leverages vector processing in a new way, Openchip's system-on-a-chip solutions promise to offer a "designed in Europe" technology that is optimized for high-performance computing, delivering greater sustainability and productivity without the complexities of porting codes to GPUs. Considered an Important Project of Common European



Dr. Kateryna Bondar (third from left) and members of the Openchip team are working with HLRS to assess the compatibility of their RISC-V chip architecture in typical large-scale HPC workflows. Photo courtesy of Openchip.

Interest (IPCEI), Openchip seeks to advance Europe's overall digital sovereignty.

"Even the best technology is useless if no one wants to buy it or use it; ensuring that our product is competitive and truly satisfies users' needs is therefore critical," says Ivan Rodero, Chief Innovation Officer at Openchip. "We are working with HLRS and other centers to better understand what HPC facilities need from the hardware. By examining a wide spectrum of applications, we aim to clarify their requirements and ensure that applications can run efficiently on RISC-V. HLRS's Future Computing Group makes this collaboration much smoother and gives us the opportunity to drive real innovation – something that is not always easy."

The collaboration with Openchip is just an early step for the Future Computing Group. Gebert anticipates that several additional collaborations with technology startups will launch in the coming months.

#### Workshops promote discussion of future computing technologies

The Future Computing Group aims to support the community surrounding emerging HPC technologies by

promoting the exchange of new insights. In January 2026, it organized a "birds of a feather" session at SCAsia, the largest meeting of the Asian high-performance computing community. On March 16 – 17, 2026, HLRS also hosted the Future Computing Workshop in Stuttgart, bringing hardware vendors, research software developers, and domain-specific researchers together to discuss the current status of high-performance computing architectures, and the potential of emerging computing paradigms. With the success of the first Future Computing Workshop, HLRS plans to continue it as an annual event series.

As Gebert explains, HLRS has an important role to play in bridging the gap between technology companies and the scientists whose research their next-generation products will benefit: "HLRS has a lot of experience supporting computational research, and we know that researchers have clear ideas about what kinds of questions they could ask with access to much faster HPC systems. In the Future Computing Group, our goal is to identify the infrastructure that will best support the development of the algorithms, software, and heterogeneous computing workflows that will take scientific computing to a new stage."

# Everllence Uses HLRS Supercomputers to Accelerate Industrial-Scale Decarbonization

Access to HLRS systems gives engineers at the global company on-demand computing capabilities and has enabled standardization of its simulation pipelines.

Operation of large-scale machinery – for example engines in power plants or cargo ships, or cement manufacturing – is a major global source of the CO<sub>2</sub> emissions that are driving climate change. As the world attempts to transition to a more sustainable future, reducing the reliance of such infrastructure on fossil fuels could lead to dramatic reductions in the world's carbon footprint. For manufacturers of industrial infrastructure this presents a major opportunity, as their expertise and resources place them in an optimal position to design the new, less carbon-intensive technologies that will make this transition possible.

One prominent player in this effort is Everllence, an international mechanical engineering firm based in Germany that has grown to more than 140 offices worldwide. Until recently called MAN Energy Solutions, Everllence has its roots in MAN, a global company that is also historically significant for building both the world's first diesel engine and first diesel-fired power plant. In recent years, Everllence has turned its focus away from fossil fuels and toward developing technologies for large-scale decarbonization, including engines powered by alternative fuels, more efficient turbomachines, industrial scale heat pumps, and carbon capture and storage facilities.

“Our company's mission is to forge the path toward net zero together with our customers,” explains Dieter Schwab, Head of Business IT at Everllence. “We manufacture technologies that could be transformational, and our ambition is to set the course for sustainable change with our product portfolio.”

Simulation with high-performance computing (HPC) is an essential tool for Everllence engineers working to achieve this goal. For the past four years this has meant integrating supercomputing resources at the High-Performance Computing Center Stuttgart (HLRS) into the company's product development pipeline. Working through HWW and T-Systems, which partner with HLRS to coordinate industrial access to its computing infrastructure, Everllence engineers around the world are using these high-performance computing resources to develop innovative technologies and customized solutions more efficiently.

## Simulation for complex, customized products

Everllence is a leading provider of propulsion, decarbonization and efficiency solutions for shipping, the energy economy and industry. Its machines are also typically both very large and very complex, making research and development impossible without simula-



The Everllence B&W ME-GI dual-fuel two-stroke engine can run on marine fuel, liquified natural gas, biomethane, and synthetic natural gas.

tion. When designing a next-generation engine for a supertanker cargo ship, for example, engineers must be able to simulate fluid dynamics behavior in the casting of steel components or combustion processes in the burning of alternative fuels such as ammonia or methanol. As a global leader in a very competitive marketplace, access to supercomputing capabilities is also extremely important for accelerating speed to market and achieving the extreme precision required to satisfy demanding specifications.

In most cases Everllence does not simply sell off-the-shelf products, but delivers customized solutions that are either unique or produced in small series that are designed to meet clients' specific technical requirements. This means that simulation is required at many stages in the product design process, starting even during the proposal phase when Everllence bids for new contracts. Here, simulation can be used to evaluate the

feasibility of a client's technical requirements, and enables engineers to propose realistic solutions that will best meet those needs without the time and expense that would be required for physical prototyping.

In recent work, for instance, HLRS computing resources were used for acoustic simulations with fluid-structure interaction that wouldn't have been computationally feasible on previous in-house systems. In addition, the common web-based platform has greatly facilitated the collaboration with other off-shore simulation engineering sites.

## Flexible access to powerful computing capabilities

Before the company began computing with HLRS, simulation engineers at Everllence would face bottlenecks resulting from limited computing capabilities, leading to competition for scarce resources within the company. Large simulations also often exceeded the

practical capacity of the company's in-house systems. Such factors meant needing to wait to run a critical simulation or to purchase computing time elsewhere. Overcoming these limitations is very important, as the ability to simulate quickly and efficiently directly affects Everllence's ability to complete project proposals and push product development forward quickly.

Meanwhile, exploding demand for high-performance hardware like GPUs and data storage devices presents a supply chain problem for companies like Everllence. Because hardware manufacturers are already struggling to keep up with demand from large-scale data centers and HPC facilities, companies whose primary business is not information technology can find it difficult to purchase hardware of their own. "The prices for hardware are constantly fluctuating," Schwab says, "and in the past one might have to wait up to a year for components to be delivered after placing an order. These dependencies are painful, and it doesn't make sense for us to invest in them."

Gaining access to HPC resources at HLRS solved both of these problems, as simulation engineers now have the ability to access powerful, state-of-the-art computing resources for large simulations on an as-needed basis. For Dr. Martin Kaiser, Product Owner for Simulation Process Integration and HPC at the company's Oberhausen location, this has major benefits: "In our work it is sometimes necessary to scale up to resources on the order of 1,000 compute cores. Because of the large overhead required to purchase and operate a supercomputer, however, it doesn't make sense for us to have our own dedicated system. Through our partnership with HLRS, we have virtually unlimited, pay-per-use access to powerful high-performance computing capabilities exactly when we need them."

Reliable access to a large-scale HPC system through a partnership with HLRS has also had other benefits for Everllence. Kaiser explained that having a single,

primary resource for high-performance computing is enabling the company to standardize its simulation pipelines. In the past, individual groups working at different locations used a variety of legacy solutions. Access to HLRS's system has made it possible to implement a more standardized approach in which shared solutions based on common use cases are now used across the entire company. This has streamlined collaboration and creates new synergies that can further enhance productivity.

### Responding quickly to volatile markets

While the technologies that Everllence produces have enormous potential to support decarbonization, the company's global prominence and its direct engagement with megatrends mean that the future is anything but certain. Factors such as changing national commitments to carbon neutrality targets, political decisions that affect demand for heat pumps, new regulations in the shipbuilding industry, or conflicts in the Middle East that disrupt the fuel industry have significant effects on demand for Everllence's green technology products and services. Such insecurities also highlight why on-demand access to HLRS's systems is so important for Everllence, as it gives the company flexibility it needs to adapt quickly to changing market dynamics.

Despite such volatility, Schwab finds great motivation in what Everllence has to offer: "The need for electric vehicles is widely discussed, but even if it were possible to shift the entire passenger car market to EV's, the effect would not be nearly as large as it would be, for example, if we can make the cement industry or district heating CO<sub>2</sub>-neutral. The potential is much, much greater." For that to occur, though, research and development will need to continue in order to make carbon capture and storage practical across the world, on an industrial scale. Making progress toward that goal will mean that HPC systems like those at HLRS will be continue to be critical.

## Supercomputing Meets the Arts

As a partner within the EU's S+T+ARTS EC(H)O program, HLRS supported four artist residencies and will host theGATE Festival of art, science, and technology.

Funded by the European Union, the S+T+ARTS program promotes innovation at the intersection of science, technology in the arts. By facilitating close collaboration, S+T+ARTS gives artists access to tools that open new creative possibilities, while scientists gain new perspectives for tackling challenges facing the modern world. Throughout 2025, the initiative coordinated the artist residency program S+T+ARTS EC(H)O, which supported ten art / science collaborations focused on promoting "ecologically conscious and human-compatible digital technologies."

HLRS is a partner within the S+T+ARTS network, working in cooperation with the Media Solution Center Baden-Württemberg and prominent European arts organizations such as Ars Electronica, Sónar, and the Salzburg Festival, among others. Within S+T+ARTS EC(H)O, HLRS hosted four artist residencies. In a project called *VR Terroir*, Bernat Cuni created an algorithm that generates and visualizes a humanoid avatar in

HLRS's CAVE virtual reality facility. As the avatar walks through a data-rich environment it accumulates visual traces of its experience. In *Looping* Oliver Kruse worked with HLRS visualization scientist Thomas Obst to simulate and visualize the effortless, helicopter-like flight of the maple seed. In addition to supporting research that could be used to construct a flying sculpture, the experiment is producing new knowledge about flight that could be used in aviation engineering. Within the project *HPC Data Symphonies* Dr. Kirell Benzi developed a visualization tool that provides real-time insights into how HLRS's high-performance computing system is being used. Finally, Merve Sahin's *Merging Visions* consists of a digital twin of the Salzburg Festival site, which is built into the side of a mountain. 3D visualizations offer a novel way to experience interactions between the socio-cultural history and ecological dimensions of the location.

The results of these projects – as well as other works that combine science, technology, and the arts – will be shown at HLRS on May 18-22, 2026 as part of theGATE Festival. This international event will include an art exhibit, a conference for professionals in the science / culture space, and a workshop for artists interested in learning about opportunities for integrating computational technologies into their practice.



Artist Oliver Kruse (left) has been working with the HLRS Visualization Department to investigate how a maple seed flies through the air, and how this could be reproduced in a flying sculpture.

# Code Porting Workshops Support Transition to GPUs

Users of HLRS's systems have been working closely with members of the center's user support team, as well as representatives of HPE and AMD, to adapt their simulation software for Hunter.


Across the HPC industry, demand for faster, more energy efficient supercomputers has been driving a transition from exclusively CPU-based to GPU-accelerated architectures. Most of the world's fastest supercomputers now include GPUs, including HLRS's current flagship supercomputer, Hunter. For longtime members of HLRS's user community, however, this change presents challenges. Many simulation codes that run at HLRS were written for CPUs, and will not run on GPUs. Redesigning these codes for GPUs is not straightforward, particularly for scientists and engineers who are not programming experts.

To help users navigate this transition, HLRS's user support staff organized three one-week code porting workshops in 2025. The workshops offered a general introduction to GPU porting and offloading on AMD processors, as well as plenty of time for hands-on, personalized support. Before each workshop, users could discuss their codes with advisors from HLRS's user support staff and set goals for the meetings. During the

workshops, they worked one-on-one with HLRS team members and experts from hardware manufacturers AMD and HPE, who provided invaluable advice based on their deep understanding of Hunter's architecture and relevant programming methods.

Among the workshop participants was Timo Seitz, a researcher at the German Aerospace Center (DLR). In his research he uses TASC3D, a code for simulating turbulent combustion, to gain a better understanding of dynamics in rocket combustion chambers. At the workshops he focused on code profiling, identifying computationally demanding routines that could be accelerated by being offloaded to GPUs. "The workshops were helpful because there was always an expert who was directly accessible," Seitz explains. "They could look at the code closely and offer suggestions about how to offload routines that were difficult to manage, or could help when something was missing in the compiler or there were other problems with the code. If one person couldn't help, he would call on others in the support team ... After a week, we made considerable progress." Seitz says that he has successfully ported TASC3D to Hunter, although more work is needed to adapt the numerical method so that it runs efficiently. As 2026 continues, HLRS's focus on supporting GPU porting and improving computational efficiency will remain a priority.

Each Code Porting Workshop lasted an entire week, enabling users to work in a focused way on specific elements of their codes.



# Good Advice from Experts in Green IT

Since passage of the German Energy Efficiency Law, operators of IT infrastructure have come under increased pressure to use energy more conscientiously. Based on experience it has gained over many years, HLRS is advising organizations on how to do so.



Dr. Norbert Conrad (center), leader of HLRS's environmental management team, accepted the Environmental Management Prize on behalf of the center at a ceremony in Berlin.

According to a 2023 study by Fraunhofer IZM, the increasing use of information technologies across all industries is projected to drive a continuing rise in energy consumption in Germany over the next decade. Considering limits on energy availability and the role of CO<sub>2</sub> emissions in driving climate change, the German federal government in 2023 passed the Energy Efficiency Law (EnEg), a set of regulations aimed at strategically reducing energy usage in the private and public sectors. The EnEg requires IT infrastructure operators, for example, to track energy consumption, implement energy and environmental management systems, and take steps to reuse the waste heat their computing systems generate.

For many organizations that in the past had not prioritized energy efficiency in IT operations, fulfilling the requirements of the EnEg has been a challenge. For

HLRS, however, they were nothing new. For nearly 15 years, the center has been working continually to improve both energy efficiency in supercomputing and sustainability across its organization. In 2019, it became the first high-performance computing center to achieve certification under the Eco-Management and Audit Scheme (EMAS), the world's most demanding framework for energy and environmental management in organizations. HLRS has continuously maintained its EMAS certification since then, an accomplishment that requires ongoing efforts to improve energy performance. This meant that when the EnEg went into effect, it was already solidly in compliance.

This leadership in the field of energy-efficient data center operations has made HLRS a useful resource for other organizations working to improve their energy performance. In 2025, for example, Dr. Norbert Conrad,

leader of the HLRS sustainability team, advised the Leibniz Supercomputing Centre in support of its ultimately successful application for EMAS certification, and was approached by a private cloud service provider to discuss its efficient energy management measures. He also recently joined an event organized by the Baden-Württemberg Ministry for Science, Research, and Art to speak with art and cultural institutions on the topic of green IT. By participating in conferences such as the Federal Waste Heat Conference and in an advisory board on energy management and waste heat reuse at the University of Stuttgart, he has been actively promoting HLRS's sustainability approaches.

"I always say, 'Do good things and talk about it,'" Dr. Conrad explains. "If I can find the time for it, I always enjoy giving lectures, and if I can make a small contribution to improving sustainability in other organizations, I am glad to do so."

#### Sustainability expertise gained at HLRS has wider impact

As an academic high-performance computing center, part of HLRS's mission is to increase and disseminate knowledge that can benefit society. Because of the center's longstanding emphasis on sustainability, employees in the sustainability team have come and gone over the years, taking expertise gained at HLRS that is helping other datacenters to improve their environmental management.

One former HLRS employee is Valeriia Lesnova, who worked on the HLRS project DEGREE, which focused on optimizing the use of free cooling for the center's supercomputer. She is now a project manager focusing on data center operations at Umwelttechnik BW, a state agency that supports companies across Baden-Württemberg in addressing energy-related issues. Located at a central node within the Baden-Württemberg digital community, she has connected companies to HLRS for advice, spreading awareness of the center's expertise.

Another alumnus of the HLRS sustainability team is Dr. Brigitte Lorenz, who in 2021 left the center to lead the University of Stuttgart Green Office. Carrying experience developed at HLRS, she recently oversaw certification of the University's data center (Technical Information and Communications Services, TIK) under the ISO 50001 standard for energy efficient operations. She is also working towards ISO 50001 certification for the University of Stuttgart as a whole, and intends to apply for EMAS certification for both TIK and the university in the future.

As one of Germany's three national supercomputing centers, operating world-class infrastructure for simulation and AI will always require large amounts of electricity. HLRS's achievements related to energy efficiency and sustainability in supercomputing, however, demonstrate its long-term focus on using these resources responsibly. By sharing its knowledge of best practices with other organizations, the center aims to help IT infrastructure operators across Germany achieve this goal as well – for the benefit of science, the economy, the public sector, and society at large.



In 2025 HLRS successfully completed a comprehensive environmental management audit, leading to recertification under the EMAS framework. HLRS was the first high-performance computing center to receive EMAS certification and has continuously held it since 2019.

## Sustainability Highlights 2025

#### HLRS Wins 2025 Environmental Management Prize

HLRS was awarded the 2025 Environmental Management Prize in the category "Best Measure for Improving Energy Performance." Presented in Berlin during the conference "30 Years of EMAS – a Sustainable Success", the award recognizes HLRS's use of an innovative energy management solution called PowerSched, which was developed in collaboration with Hewlett Packard Enterprise (HPE). Implemented in 2024 on HLRS's Hawk supercomputer, PowerSched distributes available power across a supercomputer's compute nodes in a targeted manner. How much energy each node receives depends on which application is running on it at a given time. The system continuously monitors and regulates the distribution of power in order to maintain an optimal balance at all times and adapt to changes in HLRS's energy and load budget. Using PowerSched, HLRS reduced the power consumption of Hawk by approximately 20 percent without negatively affecting computational performance. The Environmental Management Prize is organized by the German Federal Ministry for the Environment, Climate Action, Nature Conservation, and Nuclear Safety (BMUV) together with the Austrian Federal Ministry of Agriculture and Forestry, Climate and Environmental Protection, Regions, and Water Management (BMLUK).

#### EMAS Recertification Follows Successful Audit

In 2019 HLRS became the first high-performance computing center to receive certification under the European Union's Eco-Management and Audit Scheme (EMAS), the world's most demanding framework for environmental performance in organizations. EMAS requires continual improvement in environmental performance, as well as a comprehensive audit every three years. In 2025 HLRS successfully completed its second major renewal audit, and has now been certified under the EMAS framework for the third time. This achievement reflects HLRS's ongoing commitment to working


toward sustainability in a wide variety of ways, including maximizing energy efficiency in the operation of its supercomputers, prioritizing the use of sustainable products, promoting biodiversity on its surrounding campus, and nourishing a culture of environmental awareness and sustainable decision making among the entire HLRS staff. A central component of EMAS is the environmental statement, which reports on HLRS's environmental performance and provides comprehensive information about its activities, goals, and results in terms of the environment and sustainability.

#### Hunter Places in Top 3% for Energy Efficiency among World's Fastest Supercomputers

One of the most anticipated events at each year's ISC High Performance convention is the announcement of the latest TOP500 and Green500 Lists. Released twice annually, the TOP500 and Green500 Lists rate the world's highest performing supercomputers in terms of top speed and energy efficiency, respectively. In its June 2025 inaugural appearance in the Green500 List, HLRS's Hunter supercomputer was ranked 12<sup>th</sup>. (It was also ranked 53<sup>rd</sup> for overall peak performance in the TOP500 List.) The Green500 List ranks all Top500 systems based on their HPL benchmark performance per Watt of energy consumed. Hunter achieved an energy performance of 64.653 Gigaflops per Watt.

#### Construction Begins on Waste Heat Reuse Facility

In parallel with the start of construction of HLRS III, ground was broken to install new facilities for reusing the heat that will be generated by HLRS's future Herder supercomputer. Beginning in September 2025, the Nobelstraße was closed in both directions to lay a tunnel beneath the roadway, connecting HLRS III to the district heating network of the University of Stuttgart. This approach will reduce the University's heating and process heat needs, making a significant contribution to its decarbonization efforts.

An aerial photograph of a mountainous landscape. A paved road winds through the green, forested hills. In the lower-left quadrant, there is a parking lot with several cars and a small building complex. The terrain is rugged with some rocky outcrops. The overall scene is a natural, scenic view of a mountain area.

International scientific collaboration promotes exchange of valuable expertise. In partnership with the National Center for High-Performance Computing in Taipei, Taiwan, HLRS is developing a digital twin of Yangmingshan National Park, photographed here. The knowledge gained could lead to better methods for using simulation to predict and manage forest fires and landslides, also in Germany.

## Collaboration Drives Progress in Supercomputing

**Robust partnerships with HPC centers at the state, national, European, and international levels improve HLRS's ability to offer the resources, services, and expertise its user community needs.**

Today's best science would be impossible without collaboration. At the frontiers of discovery, the ability to connect experts with complementary knowledge and skills, including specialists from different disciplines, is essential for asking new kinds of research questions. This is also true of high-performance computing (HPC), a multidisciplinary research field that brings together cutting-edge technologies on the one hand and state-of-the-art applications on the other. Motivated by the growing need for simulation and artificial intelligence capabilities in research and engineering, HPC evolves by combining specialized expertise in hardware design, systems and facilities operations, and software programming. For HPC centers like HLRS, this makes frequent dialogue among HPC technologists, HPC specialists, and HPC system users critical, as it enables them to offer the best possible infrastructure for addressing users' needs. Communication and cooperation among HPC centers across the global HPC community also ensures that the latest technological advances have the greatest possible impact.

The impulse to collaborate has been embedded in the DNA of the High-Performance Computing Center Stuttgart since its beginnings. As a national supercomputing center and member of the Gauss Centre for Supercomputing, it coordinates extensively with the Jülich Supercomputing Centre and Leibniz Supercomputing Centre to provide German and European researchers access to world-class HPC capabilities. HLRS is also involved in dozens of collaborative HPC-related

research projects, and has organized Solution Centers that offer a precompetitive forum in which players in strategic industries develop solutions for shared problems. In addition, HLRS works closely with other European HPC centers and international partners from across the global supercomputing community, cultivating relationships that facilitate the exchange of new ideas for improving performance, efficiency, and usability. Focusing on infrastructure, training, and research, these activities foster innovation.

"Ultimately, HLRS collaborates with other HPC centers because it enables us to provide the best possible portfolio of resources and services – locally, nationally, and across Europe," says Prof. Michael Resch, Director of HLRS. "From systems, to support, to training, to research, our partnerships ensure that HPC users in Europe benefit from the knowledge we collect, and that HLRS implements the best available ideas from across the international HPC community."

### **Coordinating infrastructure: an integrated approach to HPC in Baden-Württemberg**

With several major universities and a diverse landscape for higher education in technical disciplines, science in the state of Baden-Württemberg has long relied on access to HPC resources. For many years, institutions across the state have collaborated within the Working Group of the Scientific Leaders of Computing Centers in Baden-Württemberg (ALWR) to optimize the usage, accessibility, and impact of digital resources and exper-

# Building the Future, Together



Since 2020, HLRS has coordinated the EuroCC and CASTIEL projects, which through the establishment and networking of National Competence Centers have promoted the development of more consistent HPC capabilities across Europe.

tise. This effort traces its origins to the early 1990s but gained momentum in 2008 with the bwGRID initiative, which set out to construct a statewide federated computing infrastructure. At that time, HLRS raised federal funds for needed hardware, while the state of Baden-Württemberg provided financial support for staff and development.

Around 2013, this effort evolved into bwHPC, the first ever statewide concept for coordinated management of high-performance computing resources and expertise in Germany. bwHPC implemented a strategic, multi-tiered approach in which academic data centers were classified based on their capabilities for peak performance and scalability. This collaboration had the benefit of reducing competition for available financial resources. At the same time, it has guided students and researchers across Baden-Württemberg to computing resources that best meet their needs, from general purpose computing to extreme scale supercomputers located at HLRS. Over the years, bwHPC has continued to evolve, establishing an integrated data management platform for moving data between computing centers and systems, as well as specialized competence centers and coordinated programs for continuing education.

Currently, bwHPC coordinates access to five bwHPC clusters, which provide general computing capabilities for university students and researchers from specific scientific user communities. For users who require access to highly parallelized supercomputing systems for computationally intensive operations, the Karlsruhe Institute of Technology (KIT) operates HoreKA, a tier-2 system within the bwHPC framework. HLRS's Hunter supercomputer currently serves as the tier 0 / 1 HPC system within the bwHPC consortium, providing capabilities for the most demanding applications of advanced computational research. Organizing computing capabilities in this way offers a unified framework within which users can scale their knowledge and codes from small clusters up to high-performance systems as their research and computational needs grow.

Organizing things in this way also connects scientists in Baden-Württemberg to the national and international high-performance computing communities. KIT's Scientific Computing Center participates in the national high-performance computing network (NHR), while HLRS connects research in Baden-Württemberg to the formidable computing resources of the Gauss Centre for Supercomputing. This ensures that there is a constant flow of knowledge concerning the latest HPC

technologies and methods across the bwHPC network. The approach extends to skills development for HPC users, as HLRS's training program is accessible from across bwHPC, linking regional, national, and international programs for continuing professional education.

### Skills development and knowledge sharing: building expertise across Europe

Since its founding in 2018 the EuroHPC Joint Undertaking (JU) has implemented a coordinated strategy aimed at improving Europe's global competitiveness in simulation, artificial intelligence, and quantum computing. Operating on the principle that Europe becomes stronger when its member states cooperate rather than compete with one another, the JU has overseen installation of powerful computing systems across the EU, established a network of AI Factories and AI Factory Antennae, and supports a broad portfolio of training initiatives to promote HPC skills development. Such efforts have helped drive Europe's ongoing efforts to reduce reliance on offshore, private service providers for HPC and AI.

Two projects coordinated by HLRS, called EuroCC and CASTIEL, have played important roles in promoting these goals. Now in their third phase, these projects led the organization of a network of JU National Competence Centers (NCCs) and Centers of Excellence across


33 European countries. By gathering and facilitating the sharing of HPC competencies contained across the NCCs, the overall European level of expertise has risen collectively. Because each NCC now serves as a key contact point for academic and industrial scientists looking for HPC expertise and computing capabilities, the EuroCC network also ensures that regardless of where HPC users in Europe are located, they can access the same kinds of high-quality resources.

"Applying to coordinate EuroCC and CASTIEL was a strategic decision," says Dr. Bastian Koller, Managing Director at HLRS and leader of the EuroCC and CASTIEL initiatives. "It was clear that there would always be other European HPC centers beside HLRS with larger supercomputers, and just having just the fastest system in Europe was never our primary goal. By coordinating these projects for the EuroHPC Joint Undertaking, we have instead tried to be a central partner in helping to bring together the enormous knowledge base that exists across the EU. This has helped us in Germany, as we have learned a great deal from what HPC centers elsewhere are doing, and at the same time the growing list of success stories across the NCCs shows that this effort is having a great impact across Europe."

HLRS also coordinated the creation of the HPC in Europe Portal ([hpc-portal.eu](http://hpc-portal.eu)), a website launched in 2025

HLRS has partnered with the Tohoku University Cyberscience Center in Sendai, Japan since 2004. A biannual collaborative conference focuses on research aimed at optimizing sustained simulation performance.





Created using data captured by a drone, a digital twin of Yangmingshan National Park makes it possible to explore its dramatic landscape in Stuttgart using virtual reality. As the project develops, researchers plan to integrate machine learning methods for identifying features in vegetation, landscape, and geology that could help predict risk of wildfires or landslides. The knowledge gained will help HLRS in developing future digital twins that could be used for managing similar environmental hazards in Germany.

that serves as a one-stop access point for HPC users in science and industry who are looking for HPC service providers, training opportunities, and other resources. Moreover, the Stuttgart center is bringing its long-standing expertise in HPC training into projects like EVITA and HPC SPECTRA, which are improving access to up-to-date training curricula for HPC users across Europe.

#### Collaborative research: plugging into the global HPC community

Although they operate under different political, economic, and social conditions, HPC centers around the world face many of the same kinds of challenges.

Performance optimization, energy efficiency, workforce training, and the imperative to translate computing capabilities into practical applications that improve society, for example, are issues with which HPC centers everywhere must contend. Although often carried out in response to local conditions, HPC research that takes place in one region can often provide inspiration for similar approaches in other regions. International cooperation can help to accelerate the diffusion of these ideas. At the same time, maintaining scientific contacts internationally can also have political benefits, establishing informal channels that improve intercultural understanding and promote friendship and solidarity across borders.

HLRS has for many years pursued active and productive partnerships with leading academic and research institutions across the world. The center currently maintains 13 formal memoranda of understanding with HPC centers and academic institutes across Europe, Asia, and the Americas, in addition to many additional informal partnerships (see map, page 34). Focusing on HPC-related topics of mutual interest, the partnerships have supported collaborative research, staff exchange, workshops, education programs, and industrial outreach.

One of HLRS's closest and longest running international collaborations has been with the Tohoku University Cyberscience Center in Sendai, Japan. The partnership started in 2004 because Tohoku and HLRS were both operating high-performance vector computing systems manufactured by NEC. At that time, some said that vector architectures had become outdated, but HLRS and Tohoku agreed that they remained very effective in providing the sustained performance their specific user communities required. Since then, scientists from HLRS and Tohoku have typically met twice each year for the Workshop on Sustained Simulation Performance. Although the technologies in the centers' more recent supercomputers have evolved beyond vector architectures, the regular dialogue has enabled both centers to gain new insights into how to maximize performance in modern HPC systems. An annual book series documents the results of these meetings for the wider HPC community.

Growing out of a collaboration that started in 1999, scientists at HLRS have also recently been working with investigators at the Taiwan National Center for

High-Performance Computing in Taipei (NCHC). In 2025, Visualization Department head Dr. Uwe Wössner traveled with NCHC scientists to Yangmingshan National Park, a nature preserve located north of Taipei on a dramatic landscape shaped by volcanic activity. With 3D scans using multi-spectral image sensors and photogrammetry, the researchers created a digital twin of the park that can now be studied in HLRS's CAVE virtual reality facility. The high-resolution scans make it possible to distinguish among plant types, and the researchers plan to train a machine learning algorithm to recognize vegetation and landscape features. As the project progresses, they intend to integrate geologic and other data in ways that will enable park management to better predict and understand the effects of wildfires, landslides, and earthquakes. For HLRS, the expertise that it will gain will be a valuable resource for future digital twin projects back in Germany.

#### Working together to address global challenges

Ultimately, high-performance computing exists to help solve complex problems facing scientific research, industrial innovation, and society. Computing power alone, however, will never be enough to achieve these goals. Considering how exascale computing, artificial intelligence, quantum computing, and other future computing technologies are currently transforming HPC, resource coordination, collaborative research, and knowledge sharing will continue to be essential elements of HLRS's activity. Working closely with the wider HPC community will help the center to ensure that technological advantages produce the greatest societal impacts, new computing paradigms are adopted and operated in a sustainable way, and HLRS's user community derives the greatest benefits for its future research.

# International Partnerships

HLRS has signed formal collaboration agreements with numerous academic and high-performance computing centers around the world. Supporting scientific collaboration, staff exchange programs, and joint education efforts, these partnerships facilitate the exchange of expertise, improving the capabilities of high-performance computing and helping to test new technologies and applications.

- 1 Lawrence Berkeley National Laboratory
- 2 National Center for Atmospheric Research (Boulder, CO)
- 3 University of Houston
- 4 Oak Ridge National Laboratory (Knoxville, TN)
- 5 Institute for Advanced Studies, University of Sao Paulo
- 6 EPCC, University of Edinburgh
- 7 GENCI (Paris)
- 8 Barcelona Supercomputing Center
- 9 CINECA (Bologna)
- 10 PDC Center for High Performance Computing (Stockholm)
- 11 CSC – IT Center for Science Ltd. (Espoo, Finland)
- 12 Talinn Institute of Technology
- 13 St. Petersburg Institute for Informatics and Automation
- 14 Keldysh Institute of Applied Mathematics (Moscow)
- 15 Donetsk National Technical University
- 16 Institute of Computational Mathematics and Mathematical Geophysics (Novosibirsk)
- 17 Wuhan University
- 18 University of Science and Technology of China (Hefei)
- 19 Shanghai Supercomputing Center
- 20 National Supercomputing Center Guangzhou
- 21 National Center for High Performance Computing (Taipei)
- 22 Korean Institute of Science and Technology Information (KISTI) (Daejeon)
- 23 Kumoh National Institute of Technology (Gumi)
- 24 Tohoku University (Sendai, JP)
- 25 Seijo University Research Center for Textual Scholarship

- Formal partnerships
- Informal partnerships
- Inactive partnerships





## Strength through Cooperation: An Interview with Steve Conway

Steve Conway is a veteran in the field of high-performance computing and artificial intelligence. As a market analyst at IDC, Hyperion Research, and Intersect360 Research, he has spent 40 years consulting with leaders across the international HPC landscape, publishing insights into key trends in the field and recommendations for exploiting the opportunities that HPC and AI could offer.

Conway has also had a long relationship with HLRS, drawing on its expertise in his efforts to forecast HPC's future. He recently visited Stuttgart, where we spoke with him about the growth of HPC capabilities in Europe, how scientific and industrial uses of HPC complement one another, the rise of artificial intelligence, and current challenges that Europe faces in its effort to achieve digital sovereignty.

*How did you first encounter HLRS and what has been your relationship with the center over the years?*

In my early days as an analyst of the high-performance computing industry, I quickly learned that HLRS was one of only three or four HPC centers in the world that was seriously and successfully working with companies, and I wanted to learn more about it. Around 1999 I was working for the HPC industry analysis firm IDC when the United States government asked us to start a user group that was not tied to a specific hardware vendor. We wanted to include international users and needed to find places that were appropriate for holding conferences, and every couple of years we would meet in Stuttgart. In 2010, the European Commission asked us to prepare a first-ever Europe-wide HPC strategy, and I led its preparation. HLRS's director, Michael Resch, was one of six reviewers of the report. The

Commission disseminated the European HPC strategy paper in 2012 and in 2014 they asked us to measure progress, so we did another massive study. Again, HLRS was a very important source, and since then has continued to be one.

*In November 2025 you published an article in HPCWire in which you named the rise of European high-performance computing as the most exciting development you have witnessed in your career. Why do you feel this way?*

Governments have for a long time recognized HPC as something that is important for scientific research. At some point they also saw its value for industrial research. They hadn't made the leap to the natural next conclusion, however, which is that high-performance computing is also important for economic competitiveness and GDP growth. After we worked on the second study for the EU, I read the proposal that ultimately went to the European Parliament for funding. What I found remarkable was that while they hardly talked about science, and just a little bit about industry, they spoke a lot about economic competitiveness. I thought that was really smart. For the first time, they were talking the funders' language, and it made a gigantic difference.

*You also wrote about the impact of PRACE and, more recently, the EuroHPC Joint Undertaking in implementing this pan-European HPC strategy.*

PRACE started as a collaboration among four countries, and very early they did something that nobody realized would turn out to be so powerful: They characterized European HPC centers according to their supercomputing capabilities. There were national Tier 0 centers – like HLRS – and other centers were categorized as Tier 1 and Tier 2. This made it possible to start thinking on a European scale. This scheme has persisted, and has even been adopted in other countries like Australia. In our 2014 study of European progress in HPC, we said that if Europe wanted to be globally important in this field, it would have to be prepared to buy a couple of exascale computers. At the time there was no way to do that, though. We recommended increasing how much

the European Commission could contribute to purchasing a large supercomputer from 20% to 50%, and changing the rules to allow member states to collaborate economically. All of a sudden the EuroHPC Joint Undertaking had the tools it needed. This decision enabled multiple member states to work closely with each other, and to propose large supercomputers as a team. This was extremely important for Europe, because traditionally the big six economies had controlled the supercomputer scene, creating a division between wealthy and less wealthy countries. The changes implemented by the JU went a long way towards solving the rich / poor, north / south, east / west problem that had plagued European HPC for years.

*As you know, HLRS has been managing the projects EuroCC and CASTIEL, which established a Europe-wide network of national competency centers for HPC and AI. The projects have promoted collaboration and the adoption of best practices across all NCCs. How do you see the impact of these initiatives?*

Once the National Competence Centers were selected, it was suddenly clear that competencies were extremely different in different places. By bringing countries across Europe together to coordinate and share expertise, EucoCC and CASTIEL have been working to address these discrepancies. Plus, even though many consider English the *lingua franca* for high-performance computing, that might not be the case in some countries. The ability to draw on other regions' expertise in HPC and AI, while also being able to translate this knowledge within your own cultural setting has become an important component for advancing Europe's HPC strategy.

*Reading your 2012 European strategy paper more than 10 years later, many of the needs you identified have since been addressed in one way or another. In what areas is improvement still needed?*

Thinking about a European strategy leads to the discussion of sovereignty. HPC is increasingly considered to be a strategic resource, which means that you can't afford to be too dependent on foreign sources for it,



*“Industrial problems are often just as challenging as scientific problems... HLRS has been one of the few centers – not just in Europe, but in the world – that really understands these things.”*

because political relations are uncertain. What does sovereignty mean, though? For Europe this has meant working to develop a homegrown supply chain, a process that is well underway. There are still some missing pieces, though. For example, if you're going to have a completely sovereign market that is walled off with trade barriers, you'd better have at least two competent vendors in each product category so that there is competitive bidding and innovation. Currently, Europe has just one major vendor of its own that is capable of building HPC systems. Processor initiatives are also very important, and Europe is still at the beginning of that trajectory.

Another important question is that if you have a sovereign market, what's the size of that market? How many vendors can that market sustain at a world-class level? And how many requirements within that geography can be incorporated into your product? There is a tension between protectionism and wanting your vendors to have as large a market as possible. Success means selling to a global market, which also means needing to address a wider range of requirements. One of the things that companies like IBM and Cray learned early is that the only way to produce a world-class product is get it into the hands of users around the world. This is how you identify requirements, which you can then embed in next-generation products.

In practical terms, complete sovereignty is unachievable. Nobody, for example, can build a processor without relying on non-indigenous capabilities, such as

manufacturing in Taiwan, supplies of materials like lithium, or advanced lithography from the Netherlands. In this sense, the goal can not be complete independence, but rather complete confidence that your local and nonlocal sources are as secure and as uninterrupted as possible. Pragmatism is also very important for sovereignty.

*You talked earlier about HPC for industry as an area where HLRS has been very active. How have interactions between the academic and industrial worlds changed over your career?*

I became interested in the topic of HPC for industry in about 2003, when I led National Science Foundation-funded studies for the Council on Competitiveness in Washington. Most clients at the NSF are small to medium-sized universities, and when we polled their HPC users and the businesses that were using their systems, we discovered that programs for industry access to HPC already existed and were wildly successful. Satisfaction scores, both for companies and the HPC centers serving them, were above 90 percent. This was not what the NSF wanted to hear, though. Their systems were all oversubscribed, with demand from the academic community often two to three times their existing capacity. The last thing they wanted to hear was that they should spend more energy marketing HPC usage to industry. I took this as an important lesson, though.

When we did another study for NSF in 2016-17, I pointed out that many universities were trying to attract indus-

try due to pressure from their local economic development councils and governments. And they were failing terribly because they didn't know how to do it. This led us to recommend conducting a study to gather best practices in providing HPC for industry. We told them that we were aware of HPC centers – including HLRS – that know how to do it and that it could be very helpful to disseminate that understanding. When we were working on that report, HLRS was very helpful in providing input.

*What benefits have you observed when academic HPC centers work with industry?*

When we began the study it was already clear that access to HPC gives companies the ability to develop superior products in shorter timeframes. But what about the benefits for HPC centers? We heard very consistent responses. The biggest advantage was that working with industry enabled the centers to identify new pathways for science. The second was that scientists love working on real-world problems, not just theoretical problems. Being able to incorporate industry applications into the mix helped the HPC centers attract and retain staff scientists and HPC center personnel.

We were shocked by these findings, because a lot of times the mindset within governments was that opening access to companies was a “mercy move.” They thought valuable resources were being wasted to address trivial problems. It turned out that industrial problems are often just as challenging as scientific problems, and that even the scientists endorsed this view. HLRS made important contributions to this study because it has been one of the few centers – not just in Europe, but in the world – that really understands these things.

*An important component of the European AI Factories initiative, which includes HammerHAI, is also to support industry, SMEs, and start-ups. What role do you see them playing?*

AI is still in a very exploratory stage and the number one question at the moment is what the borders are between frontier AI and enterprise AI. When you look

closely, frontier AI is running on HPC technology, using everything from HPC infrastructure to MPI, a classic parallel programming standard. Organizations involved in frontier AI have hired a lot of people with HPC backgrounds to run their programs, and so there's a tight and enduring connection. What is happening in typical business enterprises, though, is almost exclusively focused on increasing individual productivity, and rarely on accelerating new corporate initiatives. These companies will be looking to frontier AI organizations like HLRS and HammerHAI for new ideas to apply to corporate initiatives.

The other interesting thing about frontier AI is that different technologies are going to be well integrated. For companies like social networks that's not so important, as they are doing pure AI, but for the scientific and industrial research communities that have used HPC, sites like HLRS are going to play a very important role in combining technologies like AI and quantum computing in interesting ways.

*Are there any other big issues that you think could affect how HPC develops in the coming years?*

The last piece of the puzzle to me is what's happening in the United States. At the moment we've got a government that is actively combating facts and scientific research. That has expressed itself by making significant cuts to staffing and funding for scientific agencies that traditionally have done a lot to advance HPC and AI, including the agency that uses HPC to monitor America's nuclear weapons. If this continues unabated, then progress in the United States compared with Europe, China, and Japan could slow. That could change a lot of things, including investment incentives in Europe.

I happen to believe that the country or region that wins will be the one that is best at attracting the best and brightest people from throughout the world. All of a sudden, droves of international students who would have traveled to the US for graduate level education are not coming. It would be very interesting for Europe to develop an initiative to attract them.

In a collaboration with MACK One, the HLRS Visualization Department helped to develop a digital twin of Europa Park. Collaborative research in the AuralReality project is also investigating new ways to develop and implement immersive spatial audio for entertainment environments such as amusement parks, immersive exhibitions, escape rooms, and event locations.



# News Briefs



Michael Resch (Director, HLRS), Petra Olschowski (former Baden-Württemberg Minister of Science, Research and Art), Winfried Kretschmann (former Minister-President of the State of Baden-Württemberg), Cem Özdemir (former German Federal Minister for Education and Research), Peter Middendorf (Rector, University of Stuttgart (l-r)).

## Important Pillars of the German and European AI Ecosystem

During a visit to Stuttgart in March 2025, Baden-Württemberg Minister-President Winfried Kretschmann, interim German Federal Research Minister Cem Özdemir, and former Baden-Württemberg Science Minister Petra Olschowski highlighted the significance of HLRS and the AI Factory HammerHAI. The visit was part of Minister Özdemir's tour of leading research centers in Baden-Württemberg. Commenting on the new AI Factory, Minister-President Winfried Kretschmann said, "Baden-Württemberg is at the forefront in the fields of supercomputing and AI, and is internationally competitive. HLRS plays a key role in this. The new Hunter supercomputer and the European AI Factory HammerHAI are groundbreaking milestones for our high-tech region, our economy, and our technological sovereignty ... As a central driver of knowledge transfer to the commercial sector, HLRS is making a decisive contribution." Minister Cem Özdemir added "HLRS and the AI Factory being established here with other partners are important pillars of the German and European AI ecosystem. They show that the necessary conditions already exist here that will enable AI 'made in Europe' to remain globally competitive." Presentations made during the visit focused on efforts to accelerate uptake of AI in industry, to improve environmental sustainability, and to support engineering and city planning using digital twins.

## Event Series Explores AI on HPC Systems for SMEs

In collaboration with Baden-Württemberg networking and knowledge-transfer platform bwcon and artificial intelligence company Seedbox.ai, HLRS launched a five-part web series focused on artificial intelligence on high-performance computing systems. With a focus on increasing understanding of the opportunities HPC and AI offer for small and medium-sized enterprises (SMEs), the series brought together AI experts at HLRS and companies in Baden-Württemberg across multiple sectors. By highlighting successful applications, the ongoing series demonstrated how AI is changing engineering, manufacturing, and other industries. Talks focused on HLRS and its capabilities for HPC, AI, and hybrid workflows; applications of large language models in business settings; AI for more efficient energy system management; AI-driven manufacturing and automation; and AI in precision engineering. The final event will be held in conjunction with the inauguration of a new, AI-optimized, EuroHPC Joint Undertaking supercomputer that will be installed at HLRS in conjunction with the AI factory HammerHAI.

## HLRS Joins German AI Association

The German AI Association is the country's largest network for companies working in the field of artificial intelligence, offering a platform for information sharing, network building, and raising the visibility of German innovation using AI. It helps AI-focused businesses to address key challenges they face, and serves as a voice for the AI community on the national political level. Organized into regional chapters, it operates working groups focusing on topics such as data economy, data privacy, ethics, industry 4.0, mobility, and health, among many other topics. To extend its integration into the German AI community, HLRS became a member of the German AI Association in 2025. "This is a perfect forum for us to contribute our many years of experience in providing computing resources and support services for industry and science," said HLRS Managing Director Dr. Bastian Koller. "At the same time, we look forward to interactions within the German AI Association that will enable us to expand our knowledge and our portfolio continually." As one of Germany's three national high-performance computing centers, HLRS provides state-of-the-art infrastructure and expertise in artificial intelligence. HLRS is also coordinator of HammerHAI, established by the EuroHPC Joint Undertaking as Germany's first AI Factory.



HLRS Managing Director Bastian Koller and EPCC Director Mark Parsons signed the memorandum of understanding at the SC25 convention in St. Louis.

## HLRS and EPCC Sign Agreement to Strengthen Collaboration

At the SC25 Convention in St. Louis, Missouri, HLRS Managing Director Dr. Bastian Koller and Professor Mark Parsons, Director of EPCC at the University of Edinburgh, signed a formal memorandum of understanding that will enable closer collaboration on topics shaping the future of high-performance computing. These include numerical algorithms and emerging computer technologies for engineering applications, energy efficiency in supercomputing, quantum computing, technology transfer, and the provisioning of AI services within the framework of the European AI Factories program. Building on their complementary strengths and expertise, the partnership will support staff exchange, joint research and education projects, and the development of professional services. HLRS and EPCC share the distinction of being the first national supercomputing centers to be established in Germany and the United Kingdom, respectively. They have also enjoyed a long history of successful cooperation focused on improving access to and usage of HPC technologies in industry, in particular by SMEs and startups in the context of the Fortissimo project series. Under the EuroHPC Joint Undertaking's AI Antennas program, EPCC is also collaborating with HammerHAI – a EuroHPC JU AI Factory coordinated by HLRS – to bring high-performance AI capabilities to research and industry.



Spencer Starr, Dr. Andreas Stegmeir, and Dr. Karen Palacio-Rodriguez were named winners of the 2025 HLRS Golden Spike Awards.

## 2025 Golden Spike Awards

The 28<sup>th</sup> Annual Results and Review Workshop once again offered an opportunity to learn about the scientific research being done using HLRS's computing resources as well as approaches for maximizing the performance of research codes. Twenty-one talks and a poster session covered topics from a broad range of scientific disciplines, while many also reported on progress in porting research codes to HLRS's GPU-accelerated flagship supercomputer, Hunter. The event offered a chance to exchange insights on lessons learned and challenges faced in making this transition. At the conclusion of the meeting, Prof. Thomas Ludwig, Director of the German Climate Computing Center and Chair of the HLRS Steering Committee, announced the winners of the 2025 HLRS Golden Spike Awards, which recognize excellence in computational research and the use of HPC. Representing their respective projects, this year's Golden Spike Award winners were Spencer Starr (University of Stuttgart) for "Combining the discontinuous Galerkin solver FLEXI with deep reinforcement learning," Dr. Karen Palacio-Rodriguez (Max Planck Institute for Biophysics) for "Atomistic molecular dynamics simulations of the human nuclear pore complex," and Dr. Andreas Stegmeir (Max Planck Institute for Plasma Physics) for "Plasma turbulence simulations in the periphery of magnetic confinement fusion reactors."

## MPI Forum meets at HLRS on Path to MPI 5.0

In March 2025, members of the global community responsible for the continuing development of the Message Passing Interface (MPI) met at HLRS. Their goal was to finalize preparations for MPI 5.0, the next release of this widely used model for programming high-performance computing (HPC) applications. MPI provides ways to exploit the most essential feature of HPC using supercomputers: the ability to break a complex problem into many smaller parts and to run these processes in parallel across large numbers of processors. The new standard for the first time introduces an Application Binary Interface (ABI), which enables interoperability of MPI libraries from different MPI vendors and on different HPC platforms. In addition, MPI 5.0 improves portability and performance for MPI-based applications in containerized environments, an important factor when combining HPC and artificial intelligence. Recommendations finalized in the meeting were ratified in June, enabling the formal adoption of MPI 5.0. HLRS is a longtime member of the MPI Forum, having first contributed to the creation of the MPI 2.0 standard, released in 1996.



MPI Forum members met in the Rühle-Saal at HLRS to prepare the new standard for parallel programming.

### HiDALGO2 Clustering Event

HiDALGO2 is a EuroHPC Joint Undertaking Center of Excellence that is investigating how HPC and AI could help to address environmental and societal challenges, including the effects of climate change and the need to meet Europe's energy requirements. More specifically, research within the HiDALGO2 consortium aims to improve HPC and AI infrastructures, computational fluid dynamics methods, and uncertainty analysis in ways that take advantage of more powerful computing systems for modelling, simulation, data analysis, and visualization. HLRS is a partner in the project and on November 4, 2025 hosted a full-day conference in which researchers working within the consortium presented their work to industry experts, policymakers, and SME representatives. Under the motto "HPC and Big Data Technologies Addressing Global Challenges," the talks covered research focusing on urban air pollution, energy efficiency in buildings and cities, wildfire prediction and management, renewable energy forecasting, and water pollution and material transport models.



Over two days, Ukrainian and German scientists met at HLRS to discuss recent research.

### Conference Celebrates 52 Years of Scientific Partnership with Ukraine

Since 1973, the University of Stuttgart and Donetsk National Technical University (DonNTU) have maintained a close relationship that has supported Ukrainian science and the international exchange of expertise. The partnership entered a new phase in September 2025 with the launch of StuttWay4Ukraine, a DAAD-funded project that aims to support reconstruction in Ukrainian higher education once the ongoing war with Russia ends. In November 2025, HLRS hosted a two-day conference celebrating 52 years of German-Ukrainian scientific partnership. Scientists participated both at HLRS and streaming online from Ukraine, presenting their recent research. One area of emphasis was on applications of artificial intelligence for simulation and data analysis. Several talks also demonstrated the relevance of scientific research in the context of the war, considering topics such as mathematical modeling in energy resource management, artificial intelligence for threat assessment and risk analysis, scenario planning for a post-conflict Ukraine, and strategies for improving computer network security. Other presentations focused on topics in fields such as environmental protection, education, and linguistics. At the conference, DonNTU Vice President Prof. Iryna Shvets highlighted the importance of this partnership, saying, "When universities support one another, when colleagues reach out across borders, it gives us strength, it gives us hope, and it gives the energy to keep moving forward."

### PRACE HPC Excellence Award for Black Hole Simulations

Prof. Dr. Luciano Rezzolla, the Chair of Theoretical Astrophysics at Goethe University and Senior Fellow at the Frankfurt Institute for Advanced Studies, received the PRACE HPC Excellence Award in February 2025 at a ceremony in Brussels. The award recognizes his use of high-performance computing (HPC) to simulate physics of neutron stars and black holes. As a member of the international Event Horizon Telescope (EHT) collaboration, his team's work contributed to the generation of the world's first image of a black hole in 2019 (the M87 galaxy) and the first image of the black hole at the center of our galaxy (Sgr A) in 2022. Because black holes are impossible to see with the eye, the Event Horizon Telescope collaboration's image of M87 was widely acknowledged as a major accomplishment, and was named by *Science* magazine as its 2019 Breakthrough of the Year. Rezzolla and his team have been longtime users of high-performance computing resources at HLRS. "Generating these simulations of black hole plasma physics is impossible without high-performance computing," Rezzolla said. "I am proud to have received the PRACE HPC Excellence Award and am grateful for having had the opportunity to use the machines at HLRS to achieve this result."



Using HLRS computing resources, Prof. Luciano Rezzolla contributed to producing the first images of black holes.



FORTISSIMO  
PLUS

### FFplus Receives Overwhelming Response from SMEs and startups

Coordinated by HLRS, FFplus is a EuroHPC Joint Undertaking project that gives SMEs and startups the opportunity to explore how HPC and AI can improve their products and services. Built on a cascade funding model, FFplus supports two types of sub-projects: business experiments, which use HPC or AI to solve business challenges, and innovation studies, which develop new technologies using artificial intelligence, including optimized foundation models and applications of generative AI. In 2025, FFplus supported the launch 16 business experiments and 18 innovation studies, which together highlight the wide spectrum of industries that can benefit from HPC and AI capabilities. Following the first round of FFplus open calls, selected business experiments are testing applications of HPC and AI for agriculture, medicine, renewable energy, vehicle engineering, manufacturing, and business operations, while innovation studies are focusing on uses of AI in geography, journalism, software engineering, banking, and public administration, among other fields. A second open call for business experiments held in 2025 received more than 310 eligible proposal submissions, reflecting an enormous demand for HPC and AI capabilities. Following review, 21 of these sub-projects were chosen for funding in early 2026. The second open call for innovation studies took place in February 2026, receiving 250 submissions. Learn more about FFplus, including summaries of its funded sub-projects at [www.ffplus-project.eu](http://www.ffplus-project.eu).



The SAS25 conference brought together computational scientists, philosophers, and social scientists. Among the keynote speakers was Ben Recht, an expert in machine learning at the University of California, Berkeley.

### SAS25 Conference Explores Uncertainty in Simulation

HPC simulations and machine learning methods can deliver insights and predictions that are not accessible in any other way. Simultaneously, however, there is almost always a measurable gap between computational results and reality. In this sense, uncertainty can be understood as an integral element of the process through which simulation methods produce new knowledge. Uncertainty is not just a concept in HPC, of course, and other disciplines including economics, politics, philosophy, and science all have long traditions of reckoning with it in different ways. A three-day conference held at HLRS in July 2025 sought to draw on such multidisciplinary expertise, exploring how dialogue among scholars in the humanities, social sciences, and computational sciences could shed new light on the concept of uncertainty. This conference emerged from the HLRS research project Modeling for Policy, and was the latest in the series “The Science and Art of Simulation” (SAS25), an annual meeting that has taken place at HLRS since 2017. Speakers considered the concept from multiple perspectives, exploring the challenges of quantifying uncertainty in probabilistic models, how uncertainty can be productive and beneficial, and how to address uncertainty in the context of policy making. The conference proceedings will be collected in a forthcoming book.



### Remembering Prof. Egon Krause

In October, HLRS marked the loss of Prof. Egon Krause, former Chair of Fluid Mechanics and the Institute for Aerodynamics at RWTH Aachen. Prof. Krause was not just an outstanding scientist in the field of aerospace engineering but also one of HLRS’s founding fathers, recognizing the importance of high-performance computing in enabling German science to compete internationally. He supported HLRS’s creation, serving as its first steering committee chair and providing support, advice, and encouragement to the young HLRS team. Prof. Krause was also the first to understand the need for collaboration within the German HPC community, and so invited the chairmen of the steering committees of the Jülich Supercomputing Centre, the Leibniz Supercomputing Centre, and HLRS to meet regularly. These efforts laid the foundation for what later became the Gauss Centre for Supercomputing. In 1973 Prof. Krause was among the few scientists to establish connections with Soviet scientists, initiating the long-term exchange of people and ideas, and contributing to the melting period of the Cold War. HLRS will keep the memory of Prof. Egon Krause alive, aim to follow in his scientific footsteps, and continue his tireless efforts to support the next generation of young scientists.



Prof. Kiyoko Myojo and HLRS Director Prof. Michael Resch signed the new memorandum of understanding in Tokyo.

### HLRS and Seijo University to Collaborate on Textual Research

In an initiative aimed at exploring new applications of high-performance computing in the humanities, HLRS signed a memorandum of understanding with the Research Center for Textual Scholarship at Seijo University, Japan. Focusing on uses of HPC in the digital humanities and literary archives research, the agreement will facilitate staff exchange, joint research projects, and collaborations for education and industry. Prof. Michael Resch signed the agreement in Tokyo on behalf of HLRS, together with Prof. Kiyoko Myojo, a Professor in the Seijo University Faculty of Arts and Literature. Utilizing modern information technologies and digital libraries of manuscripts and other primary texts, the Research Center for Textual Scholarship is improving methods for scholarly editing, developing new approaches for reconstructing the processes through which literary works are written. Digital tools also make international textual scholarship possible, and the center aims to build bridges between Japanese and Western literary scholarship and resources. Prof. Myojo’s recent research, for example, is exploring at how digital methods could help to develop a more accurate scholarly edition of Franz Kafka’s novel, *The Trial*.



# User Research

## Particle Scattering Model Could Improve Low-Orbit Spaceflight

Using molecular dynamics simulations to train a machine learning model, scientists developed software that predicts and could help to exploit the effects of gas particle impacts on satellites in the upper atmosphere.

Over the past decade the space business has been booming, showing a dramatic increase in the number of new satellites launched into Earth's orbit for telecommunications, navigation, and observation. As this trend continues, many of these satellites will be deployed in very low Earth orbits (VLEO), a loosely defined region between 200 and 450 km above the Earth's surface. Flying satellites at these altitudes has many advantages, including reduced launch costs and the ability to observe Earth at higher resolution. This approach also reduces the accumulation of space junk that jeopardizes spacecraft at higher altitudes, because traces of the Earth's atmosphere present in VLEO naturally slow satellites down, causing them to re enter and burn up once their missions end.

This atmospheric drag is caused mainly by atomic oxygen. These particles continually strike satellites in VLEO, influencing their flight paths, causing surface erosion, and ultimately limiting their operational lifetimes. For this reason, a collaborative research center at the University of Stuttgart called ATLAS (Advancing Technologies for Low-Altitude Satellites) has been conducting fundamental research on interactions between rarefied high energy flows and spacecraft surfaces, developing concepts for utilizing the residual atmosphere, and exploring new design and operational strategies that will support improved VLEO satellite design, lifespan, and economic viability.

In one ATLAS subproject, Miklas Schütte of the University of Stuttgart's Institute of Space Systems and Stephen Hocker of its Institute for Functional Matter and

Quantum Technologies have been developing a method for better predicting how gas particles and surfaces of satellites interact at very low Earth orbits. "For us at ATLAS, the question is not just how we could optimize aerodynamics to minimize resistance, but how we could use the forces acting on satellites to control their orientation and orbit," Schütte explained. In a recent paper in the journal *Physics of Fluids* the team describes a new computational approach that models how gas particles reflect off surfaces by considering these interactions at the smallest of scales.

### HPC-generated data used to train a machine learning model

Close to Earth, the aerodynamics of automobiles or airplanes are typically simulated using computational fluid dynamics (CFD), which treats gases as continuous flows. Because the atmosphere at VLEO altitudes verges on the vacuum of space, however, it is very thin. This means that individual gas particles are located at greater distances from one another and CFD principles do not apply.

Researchers commonly use a method called Direct Simulation Monte Carlo (DSMC) to predict satellite drag in the upper atmosphere. Current DSMC implementations still rely on highly simplified models for gas-surface interactions, though. In most cases, models assume that reflections are either purely mirror-like or purely diffuse, with particles scattering in many directions. Experimental studies, on the other hand, consistently show that the actual distributions of reflected particles deviate significantly from these idealized assumptions.

A much more precise method for simulating particle-surface interactions is molecular dynamics (MD). Based on mathematics that accurately reproduce basic physical principles, MD simulates how molecules interact at the scale of individual atoms over very short periods of time. Achieving this resolution makes molecular dynamics simulations very computationally demanding, and they can only be done using high-performance computing (HPC) systems like those at the High-Performance Computing Center Stuttgart (HLRS).

In an ideal world one might use MD simulations to catalog every possible interaction of a gas particle and a surface, but this would be impractical even using today's fastest supercomputers. Moreover, modeling an entire satellite with many surfaces and angles, traveling through space at 8,000 meters per second, would be impossible.

Instead, the ATLAS team used molecular dynamics to support a data-driven, generative machine learning approach. To create their dataset, they used HLRS's Hawk supercomputer to simulate 225,000 particle-surface impacts in VLEO, investigating five different velocity magnitudes each at nine different incident angles. The dataset does not nearly cover all possible interactions, but provides sufficient coverage of a spectrum of potential angles and velocities that would be typical for a satellite traveling in very low Earth orbits. Using 128 cores on Hawk, it took approximately one month to generate the dataset.

Based on the results, the team then trained a machine learning algorithm on the MD data. The resulting model is able to interpolate and extrapolate from the dataset to automatically predict particle reflections for any other situation within the spectrum of the VLEO regime, including particle-surface interactions not specifically simulated using MD. When the investigators checked the resulting model's accuracy they found that its results closely replicate those seen in training and validation data, suggesting that it is much more effective at making reliable predictions than current state-of-the-art models.

Schütte developed this model into a particle scattering kernel that he then integrated into the DSMC simulation method in the plasma simulation software PICLas. DSMC, in turn, can be used to simulate at a larger scale how rarefied flows (composed of isolated particles and not continuous flows) interact with surfaces in space. "Integrating a scattering kernel into DSMC methods brings these extremely precise simulations of particle reflections at the microscopic level up into the macroscopic or mesoscopic scale that is needed to actually simulate a satellite," Schütte explained.

### Simulation opens new opportunities for improving satellite design

Schütte says that his team's successes so far make it possible to ask new questions within the ATLAS project. For one, they will look more closely at how impacts can lead to the adsorption of atomic oxygen on the satellite surface. A second question concerns the effects of roughness on the scattering of gas particles. The work so far has assumed that surfaces are flat, but it is expected that a more realistic representation of surface topography could produce different results. And finally within the broader ATLAS project, researchers will be able to use the improved models of particle reflection to optimize materials selection for satellite construction, offering better reflection capabilities for controlling spaceflight and orbit stability.

Another tantalizing idea would be to design propulsion systems that collect oxygen molecules in VLEO environments, and use them as fuel in propulsion systems to counteract the effects of drag. "Having an accurate surface model is a critical step in being able to design an intake that could capture particles in this way," Schütte said. "Right now this is still theoretical, but if it becomes possible it could be a real game changer that could dramatically reduce the cost and extend the lifecycles of VLEO satellites."

**Related publication:**  
Schütte M, Hocker S, Lipp H, et al. 2025. A machine learning framework for scattering kernel derivation using molecular dynamics data in very low Earth orbit. *Phys Fluids*. 37: 093609.

## Large-Scale Simulation Reveals Novel Insights about Turbulence

Using HLRS's Hawk supercomputer, University of Stuttgart scientists for the first time produced high-resolution simulation data characterizing the transition from low to high Reynolds numbers in turbulent boundary layers.

Scientists at the University of Stuttgart's Institute of Aerodynamics and Gas Dynamics (IAG) have produced a novel dataset that will improve the development of turbulence models. With the help of the Hawk supercomputer at the High-Performance Computing Center Stuttgart (HLRS), investigators in the laboratory of Dr. Christoph Wenzel conducted a large-scale direct numerical simulation of a spatially evolving turbulent boundary layer. Using more than 100 million CPU hours on Hawk, the simulation is unique in that it captures the onset of a canonical, fully-developed turbulent state in a single computational domain. The study also identified a previously undefined inflection point at which the outer region of the turbulent boundary layer begins to maintain a self-similar structure as it moves toward high Reynolds numbers. The results appear in a paper published in the *Journal of Fluid Mechanics*.

"Our team's goal is to understand unexplored parameter regimes in turbulent boundary layers," said Jason Appelbaum, a PhD candidate in the Wenzel Lab and leader of this research. "By running a large-scale simulation that fully resolves the entire development of turbulence from an early to an evolved state, we have generated the first reliable, full-resolution dataset for investigating how high-Reynolds-number effects emerge."

### Connecting the dots

During flight, turbulence causes excess shear stress on the surface of an aircraft. The resulting drag can reduce flight performance and fuel efficiency. To predict this effect, aerospace engineers rely on computational

models of the turbulent boundary layer, the millimeters-thin region where the surface of the aircraft interacts with free-flowing air.

For industrial applications, turbulence models do not need to replicate physics down to the finest details; they must only be accurate enough for practical use and be capable of running smoothly on modest computing resources. Before engineers can use such simplified models, however, scientific research using high-performance computing systems is necessary to provide the data on which they are based. This is why the Wenzel Lab has long used HLRS's supercomputers to run NS3D, direct numerical simulation software it created to investigate fundamental physical properties of turbulent boundary layers at extremely high resolution.

Scientists in the field of computational fluid dynamics (CFD) use a figure called the Reynolds number to characterize the developmental state of a turbulent boundary layer. The Reynolds number is the ratio of inertial forces to viscous forces in a fluid flow, which governs the local range of turbulent eddy sizes. At low Reynolds numbers, which occur early in the motion of a surface through air, nonlinear convective instabilities responsible for turbulence are quickly damped by viscous action at small scales. With increasing Reynolds number, the turbulent boundary layer becomes thicker. Large, coherent structures emerge, creating a more complex turbulent system that is not simply an extrapolation of trends at low Reynolds numbers, but has its own distinct properties.

In the past, CFD simulations have generated rich datasets for understanding turbulence at low Reynolds numbers. This is because the computational domain size and the number of time steps involved are at this stage still relatively small, meaning that simulations are not prohibitively expensive. Laboratory experiments also provide invaluable data for turbulence research. For quantities relevant to the present study, however, they have only focused on the high Reynolds number regime due to physical limitations. Sensors can only be machined so small, and some fundamental physical quantities, such as shear stress, are notoriously difficult to measure in the lab with high accuracy.

As a result, scientists have accumulated a wealth of simulation data for low Reynolds numbers and reliable experimental data for high Reynolds numbers. What has been missing, however, is a clear picture of what happens in between, as both simulation and experimental methods are of limited use. Appelbaum and his collaborators in the Wenzel Lab set out to attack this problem directly.

#### A sharp swerve

Using HLRS's Hawk supercomputer, Appelbaum ran a series of simulations that, when the results were stitched together, replicated the entire evolution of a turbulent boundary layer from low to high Reynolds numbers. Although the "real-life" situation the simulation represented might seem vanishingly small – traveling at Mach 2 for approximately 20 milliseconds – the campaign required large-scale computing power. The team used 1,024 computing nodes (more than 130,000 cores) on Hawk – one-fourth of the entire machine – for hundreds of short runs, each of which lasted 4 to 5 hours. In total, the simulation required more than 30 days of computer runtime.

"Most research groups would not take the risk of spending so much computational time on a problem like this, and might instead look at other interesting research problems that aren't as expensive," Appelbaum said. "We're the weird ones who put all of our eggs in this one basket to investigate a long-standing gap in the research."


The investment paid off. In their large-scale simulation the investigators focused on the skin friction coefficient, a value that represents the proportion between shear stress at a solid surface in a moving fluid in comparison to free momentum of the flow. It is a key parameter describing the shape of the mean velocity profile and is fundamental in determining the viscous drag.

Appelbaum used the results of the simulation to show how the previously separate datasets for low and high Reynolds numbers blend together. Whereas past research could only estimate through interpolation how the datasets might intersect, the IAG team's results reveal a sharp turn. Notably, they identified a change in skin friction scaling that is linked to the establishment of a fully-developed state in the outer 90% of the boundary layer. This self-similar state is a milestone in the turbulent boundary layer's development, signaling that scaling behavior continues in a predictable way as it evolves to industrially relevant Reynolds numbers.

"To understand self-similarity, it helps to think about the aspect ratio of a photograph," Appelbaum explained. "If I have a rectangular picture where the lengths of the sides have a ratio of 1:2, it doesn't matter whether the picture is the size of my hand or if I scale it to the size of a bus. The relationships among the elements in the photo remain self-similar, regardless how large it is. Our work confirms that the outer region of the turbulent boundary layer takes on the same kind of self-similarity once the system reaches a specific Reynolds number. Importantly, this state is coupled with the change in scaling behavior of the skin friction coefficient, which experiments have shown to remain in effect until very high Reynolds numbers seen in aerospace applications. This allows us to get an early, but realistic glimpse of the turbulent behavior in this ultimate regime of turbulence."

#### Increased performance offers new opportunities for research and engineering

This new dataset offers a unique resource that will better enable researchers in the computational fluid dynamics community to investigate turbulent boundary layers at moderate Reynolds numbers. For the Wenzel



After a certain point in the development of a turbulent flow – for example, as air moves over a wing in flight – the outer region of the turbulent boundary layer (where blue dominates) maintains a persistent, self-similar physical structure.

Lab, the next step will be to dive deeper into the physics behind the inflection point they identified. Appelbaum says that the team already has some ideas about this and plans to publish a follow-up paper soon.

In other ongoing work in the Wenzel Lab, the scientists have been busy porting the NS3D code to GPUs on HLRS's newest supercomputer, Hunter. With the help of user support staff at HLRS and computing processor manufacturer AMD, they have already verified that the code remains physically accurate and performant using the new, GPU-accelerated system. In the coming months they will be optimizing NS3D to ensure that it takes full advantage of the increased performance that Hunter offers.

"We anticipate being able to simulate larger domains at even higher turbulent states," Appelbaum said. "More computing performance will also make it more feasible to do studies in which we might run several simulations to investigate the scaling behavior of two or more parameters simultaneously."

In work that points toward this future, Tobias Gibis, a member of the Wenzel Lab and co-author of the pres-

ent work, recently defended a thesis in which he unified the scaling behavior of heat transfer and pressure gradients in turbulent boundary layers. Appelbaum added, "Building on Christoph and Tobias's transformative work on the influence of heat transfer and pressure gradients to include Reynolds number effects would undoubtedly have very high scientific value. The support and resources from HLRS are the bedrock for this type of heavy computational work."

In the meantime, the team's dataset at moderate Reynolds numbers will contribute to a pool of wall-bounded turbulent flow data and could aid in the development of more comprehensive, more accurate turbulence models. This will give engineers new capabilities for optimizing aircraft designs for a wider range of operating conditions, and for improving other kinds of machines like fans or automobiles whose efficiency relies on managing effects in turbulent boundary layers.

**Related publication:**  
Appelbaum J, Gibis T, Pirozzoli S, Wenzel C. 2025. The onset of outer-layer self-similarity in turbulent boundary layers. *J Fluid Mech.* 1015: A37.

# Evolutionary Molecular Dynamics Delivers Insights into Cholesterol Attraction

A new computational method combines machine learning and physics-based simulations to reveal a surprising explanation of cholesterol-protein attraction in cell membranes.

In the nineteenth century Charles Darwin developed his theory of evolution to explain how species change over time. Since then the concept of evolution has found many other applications – including as inspiration for so-called genetic algorithms in the computational sciences. In recent research, for example, Prof. Herre Jelger Risselada and colleagues at the Technical University Dortmund and Leiden University have been testing and improving a software package they call evolutionary molecular dynamics (Evo-MD). Developed using supercomputing resources at the High-Performance Computing Center Stuttgart (HLRS) and other high-performance computing (HPC) centers, this method has started to deliver unique insights into molecular interactions that could one day reveal strategies for new drugs.

In a recent paper published in the journal *Nature Communications*, the researchers introduce Evo-MD and demonstrate its ability to help explain important biophysical processes at the plasma membrane in mammalian cells. The paper also challenges a long-held hypothesis about how cholesterol accumulates near plasma membrane proteins, proposing an unusual mechanism that could better explain this phenomenon.

## The problem of cholesterol attraction

Cholesterol is a major component of plasma membranes in mammals, and when it interacts with proteins is instrumental in processes that transmit biochemical signals into and out of cells. Evidence suggests that

when cholesterol accumulates abnormally, it can also play a role in neurological diseases and cancer. Gaining a better understanding of how cholesterol interacts with proteins in the plasma membrane could thus help to identify new strategies for drug design.

Since 2009, data produced using a variety of scientific approaches have suggested that membrane proteins attract cholesterol at the amino acid motif CRAC (Cholesterol Recognition / Interaction Amino Acid Consensus) and its inverse, CARC. The observation posits that a genetic sequence and its resulting 3D protein structure form a good fit for cholesterol molecules, attracting them and holding them in place. As interesting as this hypothesis has been, however, these motifs have only been loosely defined and research has not been successful in demonstrating conclusively that cholesterol can bind to proteins at such sites.

“The concept that cholesterol molecules bind to CRAC and CARC motifs is based on the idea that there is a specific attraction between ligands and proteins,” Risselada explained. “Using x-ray crystallography, one can observe how cholesterol binds with other things, but resolving structure in CRAC / CARC binding has not been possible. When we started our research we wanted to prove that Evo-MD works by recovering a known motif. Surprisingly, however, it revealed that that the motif was not responsible for cholesterol attraction. Instead, it led to a better hypothesis of why it occurs.”

## Simulating evolution with machine learning and molecular dynamics

Risselada and postdoctoral scientist Dr. Jeroen Methorst set out to better understand the nature of cholesterol accumulation using Evo-MD. Specifically, they explored the thermodynamic forces that govern how a cholesterol molecule becomes attracted to a signaling protein that spans a plasma membrane.

In a computer algorithm, Evo-MD began by generating a large set of random amino acid sequences – the genetic building blocks of proteins. Evo-MD then evaluated the ability of each sequence to attract cholesterol, assigning it a fitness score. Those found to be most likely to interact with cholesterol receive a higher score, while those that are less likely to do so receive a lower one.

After a round is complete, the entire process repeats, but this time the best performers from the previous round become seeds for the next. The algorithm repeatedly selects high-fitness sequences and recombines them to create a new population. This process, called a genetic algorithm, mimics the passing of genetic material from parents to children, while also introducing random alterations in the amino acid sequences that prevent a mere copying of genetic information between generations.

To generate the fitness score for each amino acid sequence, the researchers use molecular dynamics (MD), a computationally demanding method that uses physical principles to simulate interactions among molecules down to the atomic level. Specifically, Risselada’s team investigated the ability of cholesterol and transmembrane motifs to interact using force fields, quantifying the energy within the system that results from the arrangement of molecules and atoms in space.

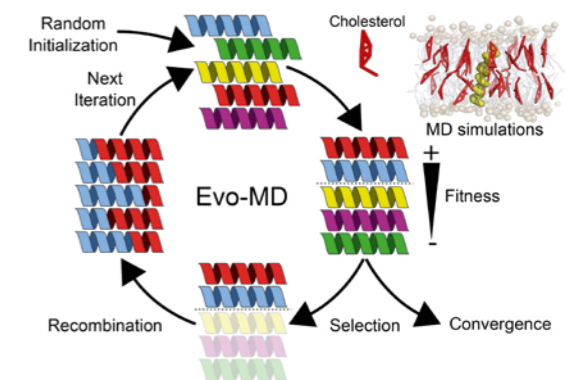
Like a species adapting successfully to its environment, repeating this cycle numerous times guides evolution toward the goal of an optimal fitness solution. By using random sampling and this data-driven reinforcement learning approach, Evo-MD is much faster and more

efficient at maximizing protein-cholesterol attraction than would be possible by individually screening every possible combination of amino acid sequences (In the example studied, this would require evaluating  $20^{10}$  possibilities!). At the same time, the incorporation of molecular dynamics simulations within the algorithm ensures that well established biophysical principles provide a scientifically reliable foundation for predicting how cholesterol and proteins interact.


## A new explanation of cholesterol accumulation

The investigators found that after simulating approximately 128 individual amino acid sequences over 40 iterations, Evo-MD begins to converge to an optimal fitness score in which the top sequences facilitate a stable interaction between protein and cholesterol molecules. “The interesting thing about using a data driven approach is that it finds a solution that is the best fit for the available data,” Risselada explained. “But then the question is, why is that the solution?”

Surprisingly, when looking at the optimal protein-cholesterol interactions, the investigators realized that CRAC / CARC motifs can not by themselves attract cholesterol, because they contain amino acid sequences



In evolutionary dynamics, molecular dynamics simulations are used to assess the ability of random amino acid sequences to attract cholesterol. As in evolution, sequences with the highest scores are recombined in an iterative manner, until an optimal attraction is reached.



Cholesterol molecules (red) attract to transmembrane proteins (yellow) in plasma membranes. Abnormal attraction can interfere with normal biochemical signaling pathways and lead to disease.

that actively repel it. Instead, the results suggest that there is a fine balance in protein structures surrounding CRAC / CARC between those that attract cholesterol and those that repel it. This exerts a more complex set of thermodynamic forces that holds the lipid in place. From this perspective, it is not that cholesterol binds to CRAC / CARC motifs, but that neighboring structures in the protein facilitate cholesterol accumulation. Experiments using nuclear magnetic resonance scans and cellular assays confirmed that this explanation is more likely, and provided more detail about the exact amino acid residues involved.

“The typical idea is that molecules bind because their shapes fit together, but in this case it’s clear that the mechanism is not normal,” Risselada said. “It makes a lot more sense to think about this as a membrane-mediated effect that fixes some compositional features but also enables a lot of flexibility. The protein induces a high energy state in the membrane, which causes an effective cholesterol attraction.”

### HPC offers new research opportunities

Evo-MD needs to run large pools of amino acid sequences simultaneously in order to capture the diversity of all potential solutions. This means running many simulations in parallel, using up to thousands of compute cores. For this reason, doing this research efficiently was only possible using HLRS’s Hawk supercomputer and other HPC systems. By conducting the research on multiple computers and receiving consistent results, the investigators gained confidence that their method is sound.

As this research project developed, the investigators also improved the Evo-MD algorithm to use supercomputing capabilities more efficiently. Completing all simulations for each generation of amino acid sequences before moving to the next generation meant that initially many cores would spend time idling while waiting for others to finish their calculations. Today, the system uses a more complex workflow where generations do not evolve in a strictly chronological fashion but jump among pools of potential amino acid sequences in ways that improve load balancing and make better use of HPC’s parallel computing capabilities.

In ongoing work, Risselada’s team has been applying Evo-MD to new problems, including using evolutionary schemes to look at interactions among small molecules, which form much more complex systems. “Many research groups that use machine learning are very pragmatic, focusing on structure-based optimization for pharmaceutical applications,” Risselada explained. “Evo-MD allows us to measure very dynamic properties, and is enabling us to pursue research that is more unique and fundamental in nature.”

#### Related publication

Methorst J, Verwei N, Hoffmann, et al. 2025. **Physics-based evolution of transmembrane helices reveals mechanisms of cholesterol attraction.** Nat Commun. 16: 9275.

## Selected Publications by Our Users in 2025

Ajmani A, Kirchhof L, Rouhi A, Mehring C. 2025. **Assessing the effect of air ventilation on the dispersion of exhaled aerosol particles in a lecture hall: simulation strategy and streamlined workflow.** Fluids. 10(5): 132.

Albanesi S, Gamba R, Bernuzzi S, et al. 2025. **Effective-one-body modeling for generic compact binaries with arbitrary orbits.** Phys Rev D. 112: L121503.

Albanesi S, Rashti A, Zappa F, et al. 2025. **Scattering and dynamical capture of two black holes: synergies between numerical and analytical methods.** Phys Rev D. 111(2): 024069.

Andreolli A, Hutchins N, Frohnapfel B, Gatti D. 2025. **Temporal decay of secondary motions in turbulent channel flows.** J Fluid Mech. 1007: A19.

Ansorge C, Wurps H. 2025. **Wind veer and speed in turbulent Ekman flow part I: scaling analysis and velocity profile model.** Boundary-Layer Meteorology. 191: 44.

Appelbaum J, Gibis T, Pirozzoli S, Wenzel C. 2025. **The onset of outer-layer self-similarity in turbulent boundary layers.** J Fluid Mech. 1015: A37.

Barone A, Djukanovic D, von Hippel G, et al. 2025. **Isoscalar octet axial form factor of the nucleon from lattice QCD.** Phys Rev D. 112: 014503.

Barone A, Djukanovic D, von Hippel G, et al. 2025. **The isoscalar non-singlet axial form factor of the nucleon from lattice QCD.** In: Proceedings of the 41<sup>st</sup> International Symposium on Lattice Field Theory (LATTICE2024).

Beckers MF, Schollenberger M, Lutz T, Krämer E. 2025. **Numerical investigation of energy harvesting wingtip-mounted propellers.** J Aircraft. ePub Aug 31.

Bergström H, Liu SJ, Bandaru V, et al. 2025. **Introduction of a 3D global non-linear full-f particle-in-cell model for runaway electrons in JOREK.** Plasma Phys Control Fusion. 67(3): 035004.

Berizzi N, Gatti D, Soldati G, et al. 2025. **Aerodynamic performance of a transonic arfoil with spanwise forcing.** J Fluid Mech. 1010: A18.

Bernuzzi S, Fontbuté J, Albanesi S, Zenginoglu A. 2025. **Perturbative hyperboloidal extraction of gravitational waves in 3 + 1 numerical relativity.** Phys Rev D. 112: 084036.

Bocchini A, Gerstmann U, Schmidt WG. 2025. **Microscopic origin of gray tracks in KTiOPO<sub>4</sub>.** Phys Rev B. 111: 104103.

Bocchini A, Kollmann S, Gerstmann U, et al. 2025. **Phosphonic acid adsorption on  $\alpha$ -Bi<sub>2</sub>O<sub>3</sub> surfaces.** Surface Sci. 760: 122776.

Bocchini A, Rüsing M, Bollmers L, et al. 2025. **Mg dopants in lithium niobite (LiNbO<sub>3</sub>, LN): defect models and impact on domain inversion.** Phys Rev Mater. 9: 074402.

Borsányi S, Capitani S, Fodor Z, et al. 2025. **Taste breaking in the minimally doubled Karsten-Wilczek action and its tree-level improvement.** Phys Rev D. 111: 074510.

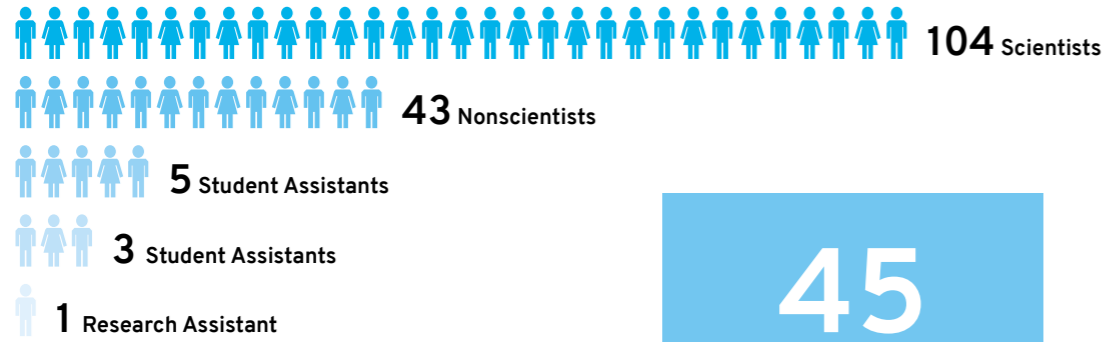
- Borsányi S, Fodor Z, Guenther JN, et al. 2025. **Lattice QCD constraints on the critical point from an improved precision equation of state.** Phys Rev D. 112: L111505.
- Conigli A, von Hippel G, Kuberski S, et al. 2025. **The hadronic contribution to the running of  $\alpha$  and the electroweak mixing angle.** In: Proceedings of the 41<sup>st</sup> International Symposium on Lattice Field Theory (LATTICE2024).
- Cook W, Daszuta B, Fields J, et al. 2025. **GR-Athena++: general-relativistic magnetohydrodynamics simulations of neutron star spacetimes.** Astrophys J. 277(1): 3.
- Dahale R, Cho I, Moriyama K, et al. 2025. **Origin of the ring ellipticity in the black hole images of M87\*.** Astron Astrophys. 699: A279.
- Deng X, Müller-Plathe F. 2025. **Underwater adhesion of oil droplets to surfaces grafted with hydrophilic polymer brushes: many-body dissipative-particle simulations.** Langmuir. 41(22): 14287-14299.
- Diederich J, Paszuk A, Ruiz Alvarado IA, et al. 2025. **Ultrafast electron dynamics at the water-modified InP(100) surface.** Adv Mater Interfaces. 12(16): e00463.
- Djukanovic D, von Hippel G, Kuberski S, et al. 2025. **The hadronic vacuum polarization contribution of the muon  $g - 2$  at long distances.** J High Energy Phys. 2025: 98.
- Ecker C, Gorda T, Kurkela A, Rezzolla L. 2025. **Constraining the equation of state in neutron-star cores via the long ringdown signal.** Nat Communications. 16: 1320.
- Folson D, Blanco C, Lisanti M, et al. 2025. **Dark matter velocity distributions for direct detection: astrophysical uncertainties are smaller than they appear.** Phys Rev Lett. 135: 211004.
- Fontbuté J, Bernuzzi S, Rettegno P, et al. 2025. **Gravitational scattering of two neutron stars.** Phys Rev D. 112: L121501.
- Fraas S, Tismer A, Riedelbauch S. 2025. **Comparison of different strategies to include structural mechanics in the optimization process of an axial turbine's runner blade.** Int J Turbomach Propuls Power. 10(4): 38.
- Gajo F, Heintz A, Keßler M. 2025. **A comparison of numerical and experimental data from wake decay in hovering rotors.** CEAS Aeronaut J. ePub Dec 11.
- Garmirian F, Pfeiffer M. 2025. **Implementation of asymptotic preserving discrete velocity methods into the simulation code PICLas.** Comput Phys Commun. 314: 109648.
- Gieg H, Schienchi F, Ujevic M, Dietrich T. 2025. **Role of muons in binary neutron star mergers: first simulations.** Phys Rev D. 112: 023036.
- Glaser T, Nolte GF, Bohamud T, et al. 2025. **Through stronger hindrance to higher reactivity: influence of the alkyl chains on the activation energy of ether cleavage on silicon.** Angew Chemie Int Ed. 65(1): e19990.
- Goddi C, Carlos DF, Crew GB, et al. 2025. **First polarization study of the M87 jet and active galactic nuclei at submillimeter wavelengths with ALMA.** Astron Astrophys. 699: A265.
- Göbel S, Peters JA, Glaser T, et al. 2025. **Adsorption of methanol on Si(001) revisited: alternative reaction pathways at increased surface temperature.** J Phys Chem C. 129(4): 10911-10916.
- Gutiérrez EM, Bhattacharya M, Radice D, et al. 2025. **Cocoon shock breakout emission from binary neutron star mergers.** Phys Rev D. 111(6): 063031.
- Heinzelmann P, Baum S, Riedmüller KR, et al. 2025. **A comprehensive benchmark dataset for sheet metal forming: advancing machine learning and surrogate modelling in process simulations.** MATEC Web Conf. 408: 01090.
- Heinzelmann P, Baum S, Riedmüller KR, et al. 2025. **A data-driven surrogate model for predicting springback in deep drawing processes.** J Phys Conf Ser. 3104: 012060.
- Höfler K, Görler T, Happel T, et al. 2025. **Milestone in predicting core plasma turbulence: successful multi-channel validation of the gyrokinetic code GENE.** Nat Commun. 16: 2558.
- Huez G, Bernuzzi S, Breschi M, Gamba R. 2025. **Gravitational waves from eccentric binary neutron star mergers: systematic biases and inadequacy of quasicircular templates.** Phys Rev D. 112: 084054.
- Imbrogno M, Meringolo C, Cruz-Osorio A, et al. 2025. **Turbulence and magnetic reconnection in multispecies plasmas.** Astrophys J Lett. 990: L33.
- Imuetinyan H, Saric D, Bazile JP, et al. 2025. **Equilibrium Thermophysical Properties of CO<sub>2</sub> + Cyclohexanol + Toluene Mixtures at  $T = 293.15\text{--}353.15$  K and  $p = 0.1\text{--}30$  MPa.** J Chem Eng Data. 70(4): 1638-1651.
- Jiang JL, Ng HHY, Chabanov M, Rezzolla L. 2025. **Long-term impact of the magnetic-field strength on the evolution and electromagnetic emission by neutron-star merger remnants.** Phys Rev D. 111: 103043.
- Klinner J, Munoz Lopez EJ, Hergt A, Willert CE. 2025. **The unsteady shock-boundary layer interaction in a compressor cascade—part I: measurements with time-resolved particle image velocimetry.** J Turbomach. 147(9): 091001.
- Klose BF, Morsbach C, Bergmann M, et al. 2025. **The unsteady shock-boundary layer interaction in a compressor cascade—part II: high-fidelity simulation.** J Turbomach. 147(9): 091002.
- Knodel MM, Nägel A, Wittum G. 2025. **Solving nonlinear virus replication PDE models with hierarchical grid distribution based GMG.** In Sequeira A, Silvestre A, Valtchev SS, Janela J, eds. Numerical Mathematics and Advanced Applications ENUMATH 2023, Vol 2. Springer Nature.
- Koehn H, Wouters T, Pang PTH, et al. 2025. **Efficient Bayesian analysis of kilonovae and gamma ray burst afterglows with FIESTA.** Astron Astrophys. 704: A55.
- Koppitz B, Saxena T, Bernhardt F, et al. 2025. **Second harmonic generation contrasts of ferroelectric domain structures and composition in lithium niobate-tantalate mixed crystals.** J Appl Phys. 138(3): 034101.
- Kostelecky J, Ansoorge C. 2025. **Surface roughness in stratified turbulent Ekman flow.** Boundary-Layer Meteorol. 191: 5.
- Kurz M, Kaushik R, Blind M, et al. 2025. **Invariant control strategies for active flow control using graph neural networks.** Computers Fluids. 303: 106854.
- Legeza Ö, Menczer A, Ganyecz A, et al. 2025. **Orbital optimization of large active spaces via AI-Accelerators.** J Chem Theory Comput. 21(13): 6545-6558.
- Lin YC, Matsumoto R, Ghorbani-Asl M, et al. 2025. **Na intercalation in bilayer graphene: formation of a close-packed trilayer.** ACS Nano. 19(47): 40612-40619.
- Liu Q, Sun H, Lin YC, et al. 2025. **Defect migration and phase transformations in two-dimensional iron chloride inside bilayer graphene.** ACS Nano. 19(4): 4845-4854.

- Lütge S, Krebs M, Risselada HJ. 2025. **Toward the evolutionary optimisation of small molecules within coarse-grained simulations: training molecules to hide behind lipid head groups.** *J Phys Chem B.* 129(9): 2482-2492.
- Mace EJ, Vrabec J. 2025. **Thermodynamic properties of ammonia–water mixtures in the vapor–liquid, critical, and Widom regions.** *J Chem Phys.* 163: 184305.
- Mastrikov YA, Gryaznov D, Chesnokov A, et al. 2025. **Oxygen defect formation and migration in Sr<sub>2</sub>FeO<sub>4-δ</sub>: Insights from first principles DFT calculations with the PBE+*U* functional.** *Phys Rev Materials.* 9: 105801.
- Mehlich K, Chagas T, Davies FH, et al. 2025. **Kirkendall voids in monolayer Mo<sub>x</sub>Ta<sub>y</sub>S<sub>2</sub> alloys on Au(111).** *J Phys Chem Lett.* ePub Dec 26.
- Methorst J, Verwei N, Hoffmann C, et al. 2025. **Physics-based evolution of transmembrane helices reveals mechanisms of cholesterol attraction.** *Nat Commun.* 16: 9275.
- Meringolo C, Camilloni F, Rezzolla L. 2025. **Electromagnetic energy extraction from Kerr black holes: ab initio calculations.** *Astrophys J Lett.* 922(1): L8.
- Mohamed AGA, Lörch D, Euchner H, et al. 2025. **CO<sub>2</sub> reduction on liquid GaInSn-metal: dynamics of the electrode-electrolyte interface.** *ChemCatChem.* 17(7): e202401740.
- Mohanty B, Pradhan L, Satpati B, et al. 2025. **Structural and compositional optimization of bimetallic NiCo alloy nanoparticles for promotion of alkaline hydrogen evolution reaction.** *J Power Sources.* 625: 235641.
- Molignini P. 2025. **Beyond-mean-field phases of rotating dipolar condensates in the strongly-correlated regime.** *J Phys Condens Matter.* 37(44): 445401.
- Molignini P. 2025. **Stability of quasicrystalline ultracold fermions to dipolar interactions.** *Phys Rev Research.* 7: L032026.
- Molignini P, Chakrabarti B. 2025. **Stability of dipolar bosons in a quasiperiodic potential.** *Phys Rev Research* 7: 023237.
- Mosnier J-P, Kennedy ET, Cubaynes D, et al. 2025. **L-shell photoionisation cross sections in the S<sup>+</sup>, S<sup>2+</sup>, S<sup>3+</sup> isonuclear sequence.** *J Phys B.* 58(7): 075002.
- Mossier P, Jöns S, Chiocchetti S, et al. 2025. **Numerical simulation of phase transition with the hyperbolic Godunov-Peshkov-Romenski model.** *J Comput Phys.* 520: 113514.
- Mossier P, Oestringer P, Jöns S, et al. 2026. **Tackling compressible turbulent multi-component flows with dynamic hp-adaptation.** *Computers Fluids.* 306: 106928. (ePub Nov 26, 2025)
- Mudimadugula R, Schianchi F, Neuweiler A, et al. 2025. **Employing deep-learning techniques for the conservative-to-primitive recovery in binary neutron star simulations.** *Euro Phys J A.* 61: 193.
- Munoz Lopez EJ, Hergt A. 2025. **Experimental and numerical investigations of SBLI and flow control on a transonic compressor cascade.** In: Flaszynski P, Babinsky H, Doerffer P, eds. *Towards Effective Flow Control and Mitigation of Shock Effects in Aeronautical Applications.* 297-324. Springer.
- Munoz Lopez EJ, Hergt A, Klinner J, et al. 2025. **The unsteady shock-boundary layer interaction in a compressor cascade—part III: mechanisms of shock oscillation.** *J Turbomach.* 147(9): 091003.
- Musolino C, Rezzolla L, Most ER. 2025. **On the impact of neutrinos on the launching of relativistic jets from “magnetars” produced in neutron star mergers.** *Astrophys J Lett.* 984(2): L61.
- Nathanail A, Mizuno Y, Contopoulos I, et al. 2025. **The impact of resistivity on the variability of black hole accretion flows.** *Astron Astrophys.* 693: A56.
- Neher NS, Faulhaber E, Berger S, et al. 2026. **Robust and efficient pre-processing techniques for particle-based methods including dynamic boundary generation.** *Comput Phys Commun.* 318: 109898.
- Neuhauser J, Metsch JH, Gatti D, Frohnapfel B. 2025. **Solution of the extended Graetz problem for nonuniform heat flux.** *Int J Heat Mass Transfer.* 249: 127198.
- Neuhauser J, Schmidt C, Gatti D, Frohnapfel B. 2025. **Predicting the global drag of turbulent channel flow over roughness strips.** *Int J Heat Fluid Flow.* 115: 109848.
- Neuweiler A, Dietrich T, Brüggemann B. 2025. **Magneto-hydrodynamic simulations of eccentric binary neutron star mergers.** *Phys Rev D.* 112: 023033.
- Ng HHY, Musolino C, Tootle SD, Rezzolla L. 2025. **Accurate muonic interactions in neutron star mergers and impact on heavy-element nucleosynthesis.** *Astrophys J Lett.* 985(2): L36.
- Nitzke I, Guevara-Carrion G, Saric D, et al. 2025. **ms2: a molecular simulation tool for thermodynamic properties, release 5.0.** *Comput Phys Commun.* 310: 109541.
- Onyedikachi A, Kruklinkii D, Ghorbani-Asl M, et al. 2025. **When defects become the crystal: periodic mirror-twin boundary phases in substoichiometric 2D-MoTe<sub>2</sub>.** *Nano Lett.* ePub Dec 23.
- Palmetshofer P, Saha R, Bryce J, et al. 2025. **Droplet impact onto sub-millimeter microstructures fully and barely submerged in a liquid film.** *Phys Fluids.* 38: 082130.
- Pfeiffer F, Krenz M, Sanna S. 2025. **Theoretical investigation and simulated imaging of halogenated tetracenes on coinage metal (111) surfaces.** *J Phys Chem C.* ePub Dec 1.
- Picker J, Ghorbani-Asl M, Schaal M, et al. 2025. **Atomic structure and electronic properties of Janus SeMoS monolayers on Au(111).** *Nano Lett.* 25(8): 3330-3336.
- Pionteck MN, Roeper M, Koppitz B, et al. 2025. **Second-order nonlinear piezo-optic properties of single crystal lithium niobate thin films.** *Phys Rev B.* 111: 064109.
- Polgar B, Homes S, Vrabec J. 2025. **Injection of a cyclopentane jet into a dense nitrogen environment by molecular dynamics.** *Phys Fluids.* 37: 114103.
- Pollinger T, Hurler M, Van Craen A, et al. 2025. **DisCoTec: Distributed higher-dimensional HPC simulations with the sparse grid combination technique.** *J Open Source Software.* 10(106): 7018.
- Ramasamy R, Zhang H, Geiger J, et al. 2025. **A comparison of low-n Mercier unstable Wendelstein stellarators and quasi-interchange modes in tokamaks.** *J Plasma Phys.* 91(4): E119.
- Rashti A, Gamba R, Chandra K, et al. 2025. **Binary black hole waveforms from high-resolution gr-athena++ simulations.** *Phys Rev D.* 111(10): 104078.
- Roy R, Alon OE. 2025. **Assessing small accelerations using a bosonic Josephson junction.** *Phys Rev A.* 111: 043307.
- Roy R, Alon OE. 2025. **Dynamics and transport of Bose-Einstein condensates in bent potentials.** *J Chem Phys.* 163: 224306.
- Ryu R, Sills A, Pakmor R, et al. 2025. **Magnetic field amplification during stellar collisions between low-mass stars.** *Astrophys J Lett.* 980: 2.

- Ryu T, Dessart L. 2025. **Light curves and spectra for stellar collisions between main-sequence stars in galactic nuclei.** *Astron Astrophys.* 698: A255.
- Ryu T, Sari R, de Mink SE, et al. 2025. **Binary mass transfer in 3D: mass transfer rate and morphology.** *Astron Astrophys.* 702: A61.
- Schramm J, Tonner-Zech R. 2025. **Unveiling correlations in metal-organic interface properties: a computational exploration of alternant and non-alternant  $\pi$ -electron systems.** *ChemPlusChem.* 90(6): e202400771.
- Schuchart J, Bouteiller A, Herault T, et al. 2025. **Scalable block-sparse matrix multiplication using template task graphs.** In: *Asynchronous Many-Task Systems and Applications: Proceedings of the Third International Workshop WAMTA.* Diehl P, Cao Q, Herault T, Bosilca G, eds. Springer Nature. Pages 120 – 132.
- Schütte M, Hocker S, Lipp H, et al. 2025. **A machine learning framework for scattering kernel derivation using molecular dynamics data in very low Earth orbit.** *Phys Fluids.* 37: 093609.
- Schwarz A, Kempf D, Keim J, et al. 2025. **Comparison of entropy stable collocation high-order DG methods for compressible turbulent flows.** *Computers Fluids.* 303: 106874.
- Scialabba G, Davidovic M, Attili A, Pitsch H. 2025. **Direct numerical simulation of soot break-through in turbulent non-premixed flames.** *Combustion Flame.* 275: 114093.
- Secchi F, Gatti D, Piomelli U, Frohnafel B. 2025. **A framework for assessing the Reynolds analogy in turbulent forced convection over rough walls.** *J Fluid Mech.* 1006: R3.
- Shamooni A, Cheng R, Zirwes T, et al. 2025. **Super-resolution of turbulent velocity fields in two-way coupled particle-laden flows.** *Phys Fluids.* 37: 093383.
- Shamraienko V, Friedrich R, Subakti S, et al. 2025. **Weak spots in semiconductor nanoplatelets: from isolated defects toward directed nanoscale assemblies.** *Small.* 21(5): 2411112.
- Shao X, Albers M, Meinke M, Schröder W. 2025. **Compressibility effects on drag reduction by spanwise traveling transversal surface waves in turbulent boundary layers.** *Phys Rev Fluids.* 10: 094603.
- Stober JL, Schulte K. 2025. **What governs droplet spreading on a wetted wall?** *Phys Fluids.* 37: 032124.
- Su Y, Saric D, Guevara-Carrion G, et al. 2025. **Diffusion coefficients of aqueous alkali-halide solutions.** *Indust Eng Chem Res.* 64(41): 20045-20061.
- Su Y, Saric D, Guevara-Carrion G, et al. 2025. **Fick and Maxwell-Stefan diffusion of the liquid mixture cyclohexane + toluene + acetone + methanol and its subsystems.** *Chem Eng Sci.* 301: 120662.
- Tang W, Moelzl M, Lehnen M, et al. 2025. **Non-linear MHD modelling of shattered pellet injection in ASDEX Upgrade.** *Nucl Fusion.* 65: 116003.
- Topolski K, Tootle SD, Rezzolla L. 2025. **Black hole-neutron star binaries with high spins and large mass asymmetries: I. Properties of quasiequilibrium sequences.** *Phys Rev D.* 111: 064022.
- Topolski K, Tootle SD, Rezzolla L. 2025. **Black hole-neutron star binaries with high spins and large mass asymmetries: II. Properties of dynamical simulations.** *Phys Rev D.* 111: 064023.
- Tugend L, Homes S, Vrabec J. 2025. **Topologies of nanoscale droplets upon head-on collision from large molecular dynamics simulations.** *Langmuir.* 41(2): 1480-1490.
- Vorspohl J, André L, Rüttgers M, Schröder W. 2025. **Drag correlations for multiphase flows using artificial neural networks.** In: *Proceedings of the 35<sup>th</sup> International Conference on Parallel Computational Fluid Dynamics.*
- Wang Y, Rinn N, Eberheim K, et al. 2025. **The  $\pi$ -trap approach for obtaining crystal structure data of inherently amorphous cluster compounds.** *Nat Commun.* 16: 7903.
- Wegmann T, Niemöller A, Meinke M, Schröder W. 2025. **Parallel Eulerian-Lagrangian coupling method on hierarchical meshes.** *J Comput Phys.* 521: 113509.
- Wengryzn O, Keßler M, Krämer E, et al. 2025. **Tight coupling of helicopter airframe including the influence of the gearbox.** *CEAS Aeronaut J.* ePub Dec 15.
- Wichai K, Deng X, Gorges, et al. 2025. **Structural and adhesive properties of cellulosic mucilage: a molecular-level investigation.** *Biomacromolecules.* 26(8): 5258-5268.
- Wittig H, Blum T, Conigli A, et al. 2025. **Machine-learning techniques as noise reduction strategies in lattice calculations of the muon  $g - 2$ .** In: *Proceedings of the 41<sup>st</sup> International Symposium on Lattice Field Theory (LATTICE2024).*
- Wulfmeier H, Yakhnevych U, Boekhoff C, et al. 2025. **Demonstration of domain wall current in MgO-doped lithium niobate single crystals up to 400 °C.** *Solid State Ionics.* 429: 116949.
- Wurps H, Steinfeld G, Ansorge C. 2025. **Wind veer and speed in turbulent Ekman flow part II: comparison to LES for atmospheric Reynolds numbers.** *Boundary-Layer Meteorology.* 191: 41.
- Yang Z, Meinke M, Krug D, Schröder W. 2025. **Numerical investigation on noise reduction of distributed propellers by phase synchronization.** *Proceedings of the 11<sup>th</sup> Convention of the European Acoustics Association.*
- Ye Q, Tiedje O, Domnick J. 2025. **Influence of droplet charge on spray painting with a high-speed rotary bell.** *J Electrostatics.* 134: 104042.
- Ye Q, Tiedje O, Shen B, Domnick J. 2025. **Impact of viscous droplets on dry and wet substrates for spray painting processes.** *Fluids.* 10(5): 131.
- Zare Pour MA, Shekarabi S, Ruiz Alvarado IA, et al. 2025. **Exploring electronic states and ultrafast electron dynamics in AlInP window layers: the role of surface reconstruction.** *Adv Func Mater.* 35(34): 2423702.
- Ziese F, Dehnen S, Sanna S. 2026. **C atoms versus Si atoms at the bridgehead positions of phenyl-decorated adamantane-type clusters: influence on the nonlinear optical response.** *Phys Chem Chem Phys.* ePub Dec 2, 2025.

# HLRS by Numbers

## 156 Staff



**45**  
Talks by  
HLRS Staff

## System Usage\*

**402.8** million  
Total Compute Unit Hours Produced

## Research

**357.2** million  
Compute Unit Hours

140 User Projects    125 User Publications

## Staff Publications



43 Papers in Journals, Books,  
and Conference Proceedings

2 Books

## Industry

**45.6** million  
Compute Unit Hours

76 Industrial Customers


\* Please note that both Hawk and Hunter were in operation in 2025. In order to compare CPU-based and GPU-accelerated systems, HLRS no longer uses core hours as a unit for tracking system usage, but rather compute unit hours.

## Education and Training



**2,181**  
Visitors at HLRS

## Third-Party Funds



**7,300,000 €**

- EU 35%
- State Gov./Land 29%
- Federal Gov./Bund 23%
- DFG 8%
- Foundations 3%
- Industry 2%

# About Us



## Inside Our Computing Room

### HPE Cray EX4000 (Hunter)

HLRS's flagship supercomputer, Hunter, went into service in February 2025. Based on the same technology as the world's three verified exascale systems, it enables large-scale, state-of-the-art applications of simulation, artificial intelligence, and data analytics, including hybrid computing approaches that combine diverse methods into powerful workflows. At its core is the AMD Instinct™ MI300A accelerated processing unit (APU), which combines CPUs, GPUs, and high-bandwidth memory in a single package. With a theoretical peak performance of 48.1 petaflops, Hunter's speed is nearly double that of its predecessor, Hawk, while slashing energy requirements at peak performance by approximately 80%. In Hunter's first appearance in the Green500 List in June 2025, it was ranked 12<sup>th</sup> in energy performance among the world's fastest supercomputers, placing it among the top 3% globally. It was also ranked 53<sup>rd</sup> for overall peak performance in the TOP500 List. Hunter is conceived as a stepping-stone system that is enabling system users to prepare for HLRS's next flagship supercomputer, Herder, which is scheduled to arrive in 2027.

### Processors

APUs: 752 AMD Instinct MI300A

CPUs: 512 AMD EPYC 9374F

### Memory technology:

APU: 512 GB HBM3 (~5.3 TB/s)

CPU: 768 GB DDR5-4800

### Storage:

Cray ClusterStor E2000 (25 PB capacity on 2,120 disks)


### Networking:

HPE Slingshot 11 Dragonfly (APU: 4 x 200 Gbps per node)

### Theoretical peak performance:

48.1 petaflops

Funding for Hunter was provided by the Baden-Württemberg Ministry of Science, Research and Arts, and by the German Federal Ministry of Research, Technology and Space through the Gauss Centre for Supercomputing (GCS). Hunter is part of the GCS national supercomputing infrastructure.



### HPE Apollo 6500 Gen10 Plus (Hawk AI Expansion)

This partition was initially installed in 2021 as part of HLRS's previous flagship system, Hawk. Although Hawk was taken out of service following the arrival of Hunter, the AI expansion continues to offer NVIDIA GPUs suitable for machine learning and artificial intelligence applications.

CPU Type: AMD EPYC

GPU Type: NVIDIA A100

Number of GPUs: 192

Performance: 120 petaflops AI performance

### NEC Cluster (Vulcan)

This standard PC cluster was installed in 2009. Its configuration has been continually adapted to meet increasing demands and provide requirement-optimized solutions, including CPU, GPU, and vector computing components. The current configuration is as follows.

#### Intel Xeon Gold 6248 @ 2.5 GHz (CascadeLake)

Number of nodes: 96

Memory per node: 128 GB

#### Intel Xeon Gold 6138 @ 2.0 GHz (SkyLake)

Number of nodes: 100

Memory per node: 192 GB

#### AMD Radeon

CPU: Intel Xeon Silver 4112 @ 2.6 GHz (Skylake)

Number of nodes: 6

Memory per node: 96 GB

CPU: 1 × AMD Radeon Pro WX8200

CPU memory: 8 GB

#### AMD Genoa

CPU: AMD Genoa EPYC 9334 @ 2.7 GHz base,

32 cores

Number of nodes: 156

Memory per node: 768 GB

#### NVIDIA A30

Number of nodes: 24

CPU: 2 × AMD EPYC 9124 Genoa, 3.0 GHz base, 16 core

Memory per node: 768 GB DDR5

CPU: 1 × NVIDIA A30

CPU memory: 24 GB HBM2e

#### Large Memory Nodes

Number of nodes: 2

CPU: 2 × AMD EPYC 9124 Genoa, 3.0 GHz base, 16 core

Memory per node: 3,072 GB DDR5

#### Interconnects

Infiniband EDR / HDR

## User Profile 2025

The following is a summary of usage of HLRS's supercomputer in calendar year 2025.

### Gauss Centre for Supercomputing (GCS)

#### Awards for large-scale projects \*

Number of new large-scale projects	9
AI node hours (millions)	205.9

#### GCS projects – system usage in 2025

Number of active GCS projects	140
AI node hours (millions)	146.96
CPU node hours (millions)	1.72

#### Other scientific usage

AI node hours (millions)	3.18
CPU node hours (thousands)	412.35

#### Total scientific usage

Compute unit hours (millions)	357.2
-------------------------------	-------

### Usage by Industry

#### HLRS Industrial Clients

Number of industry clients	76
Active users in 2024	41
Compute unit hours (millions)	17.24

#### HWW Usage

Compute unit hours (millions)	28.37
-------------------------------	-------

#### Total Industrial Usage

Compute unit hours (millions)	45.61
-------------------------------	-------

#### Total System Usage in 2025

Compute unit hours (millions)	402.77
Percentage science	88.67%
Percentage industry	11.33%

### System Usage by Scientific Discipline (Hunter)

CFD	57,24%
Physics	40,71%
Computer Science	1,61%
Transportation and Climate	0,23%
Reactive Flows	0,15%
Materials Simulation & Materials Research	0,04%
Chemistry	0,02%
Solid State Physics	0,01%
Other	< 0,01%
Structural Mechanics	< 0,01%

### System Usage by State (Hunter)

North Rhine-Westphalia	78,72%
Baden-Württemberg	16,40%
Bavaria	1,93%
Federal Research Center	1,28%
Hesse	0,83%
Thuringia	0,66%
Brandenburg	0,15%
Berlin	0,02%
Lower Saxony	< 0,01%
Rheinland-Palatinate	< 0,01%
Saxony	< 0,01%

\* Please note that both Hawk and Hunter were in operation in 2025. In order to compare CPU-based and GPU-accelerated systems, HLRS no longer uses core hours as a unit for tracking system usage, but rather compute unit hours.

# Third-Party Funded Research Projects

In addition to providing supercomputing resources for scientists and engineers in academia and industry, HLRS conducts its own funded research on important topics relevant for high-performance computing (HPC), artificial intelligence, visualization, and high-performance data analytics. These activities, many of which are conducted in collaboration with investigators at other institutes and in industry, address key problems facing supercomputing and are opening up new opportunities for addressing key German, European, and global challenges. The following is a list of funded projects in 2025.

For more information about our current projects, visit [www.hlrs.de/projects](http://www.hlrs.de/projects).

## 3xa

November 2022 – December 2025 BMFTR

Will develop scalable methods for the simulation of three-body interactions in particle systems, applying vectorized kernels, dynamic load balancing approaches, and adaptive resolution schemata.

## AI Alliance BW

August 2023 – March 2026 MWK

Aims to develop a data platform for AI-relevant data exchange among academic and industrial stakeholders in Baden-Württemberg.

## AuralReality ■

September 2025 – August 2027 BMWE

AuralReality is creating an integrated European solution for development and implementation of spatial audio in location-based entertainment environments such as amusement parks or immersive exhibitions.

## BEGIN HPC+

April 2024 – April 2027 MWK

BEGIN HPC+ complements funding from the Digital Europe Programme (EuroHPC JU) and the BMFTR for the CASTIEL 2 and EuroCC 2 projects, providing for additional resources (staff, travel, and equipment) and supporting activities aligned with the projects' objectives.

## bwHPC-S5

July 2018 – January 2026 MWK

Coordinates support for HPC users in Baden-Württemberg and the implementation of related measures and activities, including data intensive computing and large-scale scientific data management.

## CASTIEL 2

January 2023 – March 2026 JU

CASTIEL 2 facilitates collaboration among the EuroCC 2 and EuroCC4SEE National Competence Centers, and the EuroHPC Joint Undertaking Centers of Excellence, promoting the development of HPC expertise and the adoption of leading codes across Europe.

## CEEC

January 2023 – December 2026 BMFTR, JU

The Center of Excellence in Exascale CFD will improve European state-of-the-art computational fluid dynamics algorithms to prepare them for efficient performance on exascale supercomputers.

## ChEESE-2P

January 2023 – December 2026 JU, BMFTR

Focusing on critical applications for the prediction of geohazards, the Centre of Excellence for Exascale in Solid Earth aims to become a hub for HPC software within the solid earth community.

## DECICE

December 2022 – November 2025 EU

DECICE developed an open and portable cloud management framework that will enable the automatic and adaptive optimization of software applications for heterogeneous computing architectures.

## EE-HPC

September 2022 – November 2025 BMFTR

EE-HPC tested an approach for improving energy efficiency in HPC systems by automatically regulating system parameters and settings based on current job requirements.

## EuroCC 2

January 2023 – March 2026 JU, BMFTR

Supported by the EuroHPC Joint Undertaking, EuroCC 2 manages a European network of National Competence Centers (NCC) for high-performance computing and related technologies, promoting a common level of expertise across the participating countries.

## EuroCC4SEE

April 2024 – April 2026 JU, BMFTR

As sister project of EuroCC 2, EuroCC4SEE manages a subset of NCC for HPC and related technologies in five countries of Southeastern Europe, promoting a common level of expertise across the participating countries.

## EVITA ■

April 2025 – March 2029 JU

A collaboration among leading institutions to build a unified framework for standardizing training modules in high-performance computing and related technologies.

## EXCELLERAT P2

January 2023 – December 2026 BMFTR, JU

EXCELLERAT P2 is developing advanced applications for engineering in the manufacturing, energy, aeronautics, and automotive sectors, focusing on use cases that demonstrate the importance of HPC, HPDA, and AI for European competitiveness.

## FFplus

May 2024 – April 2028 JU

FFplus supports European small and medium-sized enterprises (SMEs) and startups in testing new applications of high-performance computing and artificial intelligence.

## GreenShift ■

October 2025 – September 2029 EU

GreenShift will establish a joint European Master's program in artificial intelligence and high-performance computing focused on green digital innovation in the transportation sector.

■ New project ■ Grant awarded, starts in 2026

### Funder Abbreviations:

BMFTR – Federal Ministry of Research, Technology and Space | BMWE – Federal Ministry for Economic Affairs and Energy | CZS – Carl Zeiss Foundation | DAAD – German Academic Exchange Service | DFG – German Research Foundation | EU – European Union | JU – EuroHPC Joint Undertaking | MSA – Bavarian State Ministry of Science and the Arts | MSC – Lower Saxony Ministry of Science and Culture | MWK – Baden-Württemberg Ministry for Science, Research, and Art

### HammerHAI ■

April 2025 – March 2028

JU, BMFTR, MWK, MSA, MSC

HammerHAI is building a state-of-the-art “AI Factory” that will provide European businesses and researchers with secure, scalable, and easily accessible AI resources and support, lowering barriers needed to deploy AI-driven solutions.

### HANAMI

March 2024 – February 2027 JU

HANAMI fosters collaboration between Europe and Japan to develop applications for future generations of supercomputers across diverse scientific fields, including environmental sciences, biomedicine, and materials science.

### HiDALGO2

January 2023 – December 2026 BMFTR, JU

HiDALGO2 is addressing challenges caused by climate change, focusing on technical issues related to scalability on HPC and AI infrastructures, the use of computational fluid dynamics methods, and uncertainty analysis.

### HPC SPECTRA

January 2024 – March 2026 JU

HPC SPECTRA will promote the development of HPC expertise across Europe by building a comprehensive online platform of training opportunities, making it easy for trainees to find courses that fit their interests and needs.

### HPCTRAIN ■

January 2026 – December 2030 JU

Focusing on companies using HPC, HPCTRAIN provides young professionals across Europe professional traineeships in HPC within the HPCTRAIN consortium, EuroHPC Hosting Entities, and other European projects and HPC stakeholders.

### InHPC-DE

November 2017 – December 2026 BMFTR

InHPC-DE furthers the federation of Germany's three national HPC centers, addresses new requirements such as security, and evaluates the Gaia-X ecosystem in the context of high-performance computing.

### Inno4scale

July 2023 – June 2025 JU

Inno4scale identified and provided funding to support the development of advanced algorithms and applications for upcoming European exascale systems.

### MERIDIONAL

October 2022 – September 2026 EU

Developing a tool for assessing the performance and loads experienced by onshore, offshore, and airborne wind energy systems.

### NFDI4Cat

October 2020 – September 2025 DFG

As a participant in the German National Research Data Infrastructure initiative, this consortium created a national platform for data integration in catalysis and chemical engineering research.

### POP3

November 2024 – December 2026 EU, BMFTR

The Performance Optimization and Productivity Centre of Excellence in HPC provides services for academic and industrial codes to assess and improve their efficiency on large-scale computing systems.

### SimTech

July 2019 – June 2025 DFG

An interdisciplinary Excellence Cluster at the University of Stuttgart developed simulation technologies to enable integrative systems science. HLRS supported the development of efficient methods for uncertainty quantification and management.

### Simulated Worlds

January 2011 – August 2028 MWK

Offers students opportunities to develop and execute simulation projects in collaboration with HLRS scientists.

### SiVeGCS

January 2017 – December 2032 BMFTR, MWK

Coordinates and ensures the availability of HPC resources of the Gauss Centre for Supercomputing, addressing issues related to funding, operation, training, and user support across Germany's national HPC infrastructure.

### SRI DiTEnS

April 2023 – March 2029 CZS

SRI DiTEnS is developing methods for discursive transformation in local energy systems, using urban digital twins involving virtual reality to support decision making among stakeholders.

### S+T+ARTS E(C)HO

January 2024 – December 2026 EU

This artist residency program supports collaboration between artists and experts in science and technology, focusing on promoting ecologically conscious and human-compatible digital technologies.

### StuttWay4Ukraine ■

July 2025 – June 2029 DAAD

Led by the University of Stuttgart, this project is focused on deepening academic collaboration between Germany and Ukraine, and strengthening the Ukrainian higher education system in the context of postwar reconstruction.

### TargetDART

October 2022 – December 2025 BMFTR

Developed a task-based approach for highly scalable simulation software that mitigates load imbalance on heterogenous systems through dynamic, adaptive, and reactive distribution of computational load across compute resources.

### TOPIO

November 2022 – December 2025 BMFTR

Focusing on a large-scale, high-resolution Earth system model, TOPIO investigated read and write rates for large amounts of data on high-performance file systems, as well as approaches that use compression to reduce the amount of data without causing a significant loss of information.

### UKAIFA ■

February 2026 – January 2029 JU

In a partnership with EPCC at the University of Edinburgh, the AI Factory HammerHAI is serving as a bridge between the United Kingdom and the European AI Factory network, accelerating the uptake of artificial intelligence in engineering, the automotive sector, and industry.

### URBANOME

March 2021 – August 2025 EU

URBANOME aims to promote health, well-being, and quality of life in cities by systematically integrating health considerations in city politics and the activities of city residents.

### WindHPC

October 2022 – March 2026 BMFTR

In the first ever project to connect computers in wind farms with an HPC center, WindHPC developed strategies to reduce energy consumption by improving efficiency in simulation codes, HPC workflows, and data management.

# HPC Training Courses in 2025

HLRS offered 41 courses in 2025, providing continuing professional education on a wide range of topics relevant for high-performance computing. The courses took place over 164 course-days online, in Stuttgart, and in cooperation with other institutes in Germany and internationally. A total of 1,623 trainees participated in these activities.

For a current listing of upcoming courses, please visit [www.hlrs.de/training](http://www.hlrs.de/training).

Date	Location	Topic	Host
Jan 13-31	HLRS/online	Hackathon: Porting and Optimization for Hunter <sup>NEW</sup>	HLRS/AMD/HPE
Jan 21-23	HLRS/online	Hybrid Programming in HPC - MPI+X	HLRS
Feb 3-Mar 7	online/self-paced	HPC Cluster Design, Cost & Sustainability	HLRS (SCA)
Feb 17-21	HLRS	Iterative Solvers and Parallelization	HLRS
Mar 3-7	online	Introduction to Computational Fluid Dynamics	HLRS/DLR
Mar 10-12	online	Data analytics for engineering data using machine learning	F.SCAI/HLRS
Mar 13-16	HLRS/online	Modern C++ Software Design (Intermediate)	HLRS
Mar 17-21	ZIH Dresden	Parallel Programming with MPI & OpenMP and Tools	ZIH/HLRS
Mar 24-28	HLRS/online	Fortran for Scientific Computing	HLRS
Apr 1-4	ZDV Mainz	Parallelization with MPI and OpenMP	ZDV/HLRS
Apr 8-9	online	Bootcamp: N-Ways to GPU Programming *	EuroCC/NVIDIA
May 5-8	online	AMD Instinct GPU Training <sup>NEW</sup>	HLRS/AMD
May 9-23	HLRS/online	Hackathon: Porting and Optimization for Hunter <sup>NEW</sup>	HLRS/AMD/HPE
May 13-16	HLRS/online	Modern C++ Software Design (Advanced)	HLRS
May 27-28	online	Bootcamp: AI for Science *	EuroCC/NVIDIA
Jun 3-6	online	Node-Level Performance Engineering	HLRS/NHR@FAU
Jun 3-24	online/self-paced	Natural Language Processing <sup>NEW</sup>	HLRS (SCA)
Jun 16-17	online	Efficient Parallel Programming with GASPI	HLRS/F.ITWM
Jun 17-18	online	Bootcamp: Multi-GPU Programming *	EuroCC/NVIDIA
Jun 23-28	HLRS	The Gray Scott HPC Summer School <sup>NEW</sup>	HLRS/LAPP
Jun 30-Aug 1	online/self-paced	HPC Cluster Design, Cost & Sustainability	HLRS (SCA)

Date	Location	Topic	Host
Jul 1-4	HLRS/online	Modern C++ Software Design (Intermediate)	HLRS
Jul 9-10	online	Bootcamp: AI Profiling * <sup>NEW</sup>	EuroCC/NVIDIA
Aug 18-Sep 26	online	Introduction to SYCL2020 <sup>NEW</sup>	HLRS/INTEL
Aug 20-29	online	Six-day Course in Parallel Programming with MPI/OpenMP	ETH/HLRS
Sep 1-Oct 10	online/self-paced	Data Analysis with HPC <sup>NEW</sup>	HLRS (SCA)
Sep 8-12	HLRS	Introduction to Computational Fluid Dynamics	HLRS/DLR
Sep 10-11	LRZ	Iterative Solvers for Linear Systems	LRZ/HLRS
Sep 23-26	HLRS	Julia for High-Performance Computing	HLRS
Oct 13-17	HLRS	Parallel Programming Workshop (with TtT)	HLRS
Oct 20-Nov 10	online/self-paced	Time Series Forecasting <sup>NEW</sup>	HLRS (SCA)
Oct 21-22	online	Introduction to OpenMP Offloading with AMD GPUs	HLRS/AMD
Oct 23-24	HLRS	Scientific Visualization	HLRS
Oct 28-31	HLRS	Multi-GPU Deep Learning	HLRS
Nov 3-5	online/self-paced	Parallel Programming with MPI	HLRS (SCA)
Nov 3-Dec 5	online/self-paced	Data Management	HLRS (SCA)
Nov 11-20	online	GPU Programming using CUDA <sup>NEW</sup>	HLRS
Nov 20-Dec 5	HLRS/online	Hackathon: Porting and Optimization for Hunter <sup>NEW</sup>	HLRS/AMD/HPE
Nov 25-28	HLRS/online	Modern C++ Software Design (Advanced)	HLRS
Dec 1-4	online	Advanced Parallel Programming with MPI and OpenMP <sup>NEW</sup>	JSC
Dec 8-12	HLRS/online	Fortran for Scientific Computing	HLRS

- Parallel Programming
- Simulation and Computational Fluid Dynamics (CFD)
- Performance Optimization and Debugging
- Data in HPC, Deep Learning, Machine Learning and AI
- Programming Languages for Scientific Computing
- Scientific Visualization
- Compute Cluster - Usage and Administration
- Community- and domain-specific content
- Hardware Accelerators

\* EuroCC 2 courses: HLRS is a member of the Gauss Centre for Supercomputing (GCS). SIDE is the German National Competence Centre (NCC) for High-Performance Computing. This course is provided within the framework of EuroCC 2.

TtT: Train the Trainer Courses | NEW: New or significantly modified course

#### Institute Abbreviations

AMD - Advanced Micro Devices, Inc. | DLR - German Aerospace Center | ETH - Scientific IT Services, ETH Zurich | EuroCC - EuroCC | F.ITWM - Fraunhofer Institute for Industrial Mathematics | F.SCAI - Fraunhofer Institute for Algorithms and Scientific Computing | HLRS - High-Performance Computing Center Stuttgart | HLRS (SCA) - Supercomputing Academy of the HLRS | HPE - Hewlett Packard Enterprise | INTEL - Intel Corporation | JSC - Jülich Supercomputing Centre | LAPP - Annecy Particle Physics Laboratory | LRZ - Leibniz Supercomputing Centre | NHR@FAU - Erlangen National High Performance Computing Center | NVIDIA - Nvidia Corporation | ZDV - Data Center, University of Mainz | ZIH - Center for Information Services and High Performance Computing (TU Dresden)

# Workshops and Conferences in 2025

## March 3–6 MPI Forum

Members of the global community responsible for the continuing development of the Message Passing Interface (MPI) met at HLRS to make final preparations for the next release of this critical parallel programming standard.

## May 27–28 39<sup>th</sup> Workshop on Sustained Simulation Performance

Organized in cooperation with NEC and Tohoku University, this annual meeting brings scientists, application developers, and hardware designers from different continents together to discuss hardware architectures, programming styles, and strategies for achieving the highest possible sustained application performance.

## July 28–30 SAS24 – Uncertainty

Organized by the HLRS Philosophy of Computational Sciences department, this Science & Art of Simulation conference encouraged interdisciplinary reflection on how to recognize and react to the gaps between computational models and reality.

## October 7–8 HPC Quantum AI Conference

Organized in collaboration with the Quantum Effects 2025 convention, this international academic meeting explored opportunities and challenges at the intersection of quantum computing, high-performance computing (HPC), and artificial intelligence (AI).

## October 9–10 28<sup>th</sup> Results and Review Workshop

Scientists and engineers, including users of HLRS's computing infrastructure, presented research using high-performance computing and discussed challenges and best practices in maximizing performance of research codes on HPC systems.

## October 21 2<sup>nd</sup> SIDE Forum

Organized by the German national competence center for HPC (SIDE), this meeting invited SMEs, startups, industrial companies, and public sector agencies to discover how HPC technologies can help solve challenges, optimize processes, and prepare for the future.

## November 4 HIDALGO2 EU Clustering Event

Organized by the HIDALGO2 center of excellence, this meeting highlighted innovative research that is using HPC to address social challenges, including urban air pollution, building construction, wildfires, renewable energy, and material transport in water.

## November 6–7 16<sup>th</sup> International Parallel Tools Workshop

Organized in partnership with the Center for Information Services and High Performance Computing in Dresden (ZIH), this annual workshop is a key forum to discuss trends and advances in parallel computing tools for debugging, performance analysis, programming environment development, and emerging architectures.

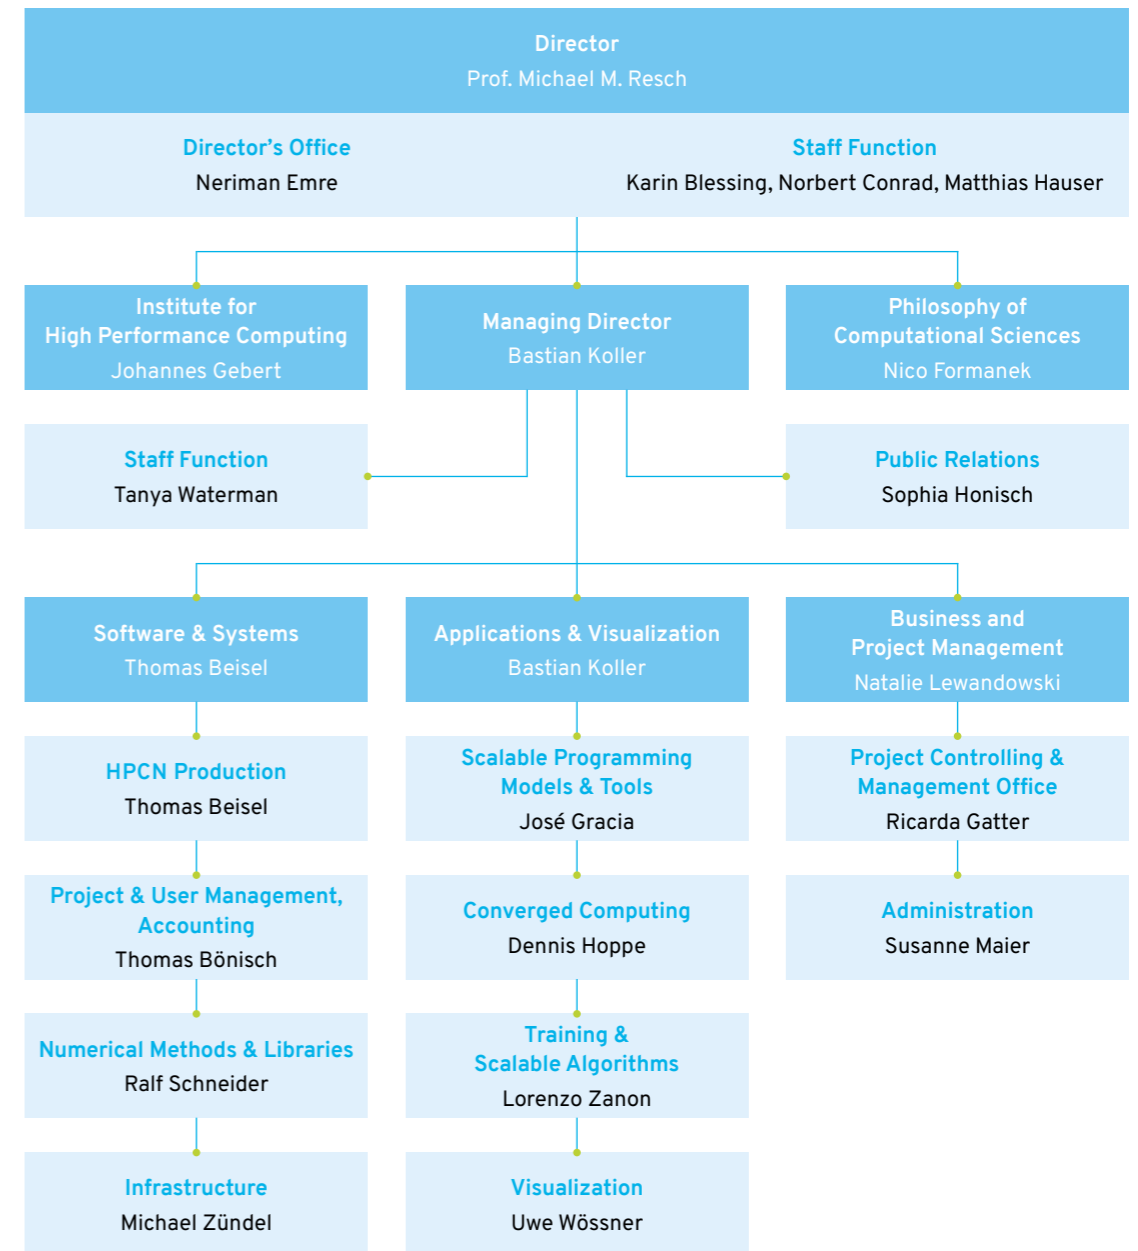
## November 18–19 52 Years of Scientific Cooperation

This two-day conference celebrated a longstanding partnership between the University of Stuttgart and Donetsk National Technical University, Ukraine, bringing together German and Ukrainian scientists from across many disciplines to share recent research.

## December 2 9<sup>th</sup> Industrial HPC User Round Table

The annual iHURT meeting facilitates dialogue between HLRS and its industrial user community, focusing on innovative applications of HPC for research and development as well as challenges that industry faces in using HPC.

# Organization Chart



# Departments

## Administration

*Leader: Susanne Maier*

Manages issues related to the day-to-day operation of HLRS. Areas of responsibility include financial planning, controlling and bookkeeping, financial project management and project controlling, legal issues, human resources development, personnel administration, procurement and inventory, and event support.

## Converged Computing

*Leader: Dennis Hoppe*

Evaluates novel technologies that will have a strong impact on the future use of high-performance computing. These technologies include artificial intelligence, cloud and edge computing, as well as quantum computing. For example, the group promotes the convergence of high-performance computing and artificial intelligence with the aim of incorporating AI methodologies into classical simulations to create hybrid HPC / AI workflows. This includes the development of AI solutions, especially in a business context, using cutting-edge technologies for Big Data, machine learning, and deep learning. The group also researches related virtualization technologies such as containers, orchestration, and job scheduling. By exploiting synergies between virtualization and HPC, the group has gained expertise in developing and operating dynamic and scalable federated cloud computing services.

## High-Performance Computing Network – Production (HPCN Production)

*Leader: Thomas Beisel*

Responsible for the operation of all platforms in the compute server infrastructure. This department also operates the network infrastructure necessary for HPC system function and is responsible for security on networks and provided platforms.

## Infrastructure

*Leader: Michael Zündel*

Plans, develops, and operates the technical facilities and infrastructure of HLRS. This department ensures the reliable and efficient operation of HLRS's high-performance computing systems, provides a comfortable working environment for researchers and HLRS administration, and supports all aspects of energy-efficient HPC operations. The team also supports HLRS's sustainability program. It collects and monitors energy consumption data and drives continuous improvement through energy efficiency measures.

## Numerical Methods and Libraries

*Leader: Dr.-Ing. Ralf Schneider*

Provides numerical libraries and compilers for HLRS computing platforms. The department has expertise in implementing algorithms on different processors and HPC environments, including vectorization based on the architecture of modern computers. Department members also conduct research related to the simulation of blood flow and bone fracture in the human body, and are responsible for training courses focused on programming languages and numerical methods that are important for HPC.

## Philosophy of Computational Sciences

*Leader: Nico Formanek*

Examines both how computer simulation and machine learning are changing science and technology development, and how society and politics react to these changes: Does simulation and machine learning change our understanding of knowledge and how we justify scientific results? How can computer-based methods help to overcome uncertainties about the future? And how do we deal with the uncertainties of simulation and machine learning itself?

## Project Controlling and Management Office

*Leader: Dr. Ricarda Gatter*

The Project Controlling and Management Office (PCMO) is responsible for the controlling and quality assurance of current research projects at HLRS or with HLRS as a beneficiary, and the management of large-scale third-party funded projects, including coordination and business development tasks. The PCMO also assists coordination at the proposal planning and writing stage and acts as a supporting and coordinating entity between the HLRS management, department heads, and HLRS administration in project-related matters.

## Project and User Management, Accounting

*Leader: Dr. Thomas Bönisch*

Responsible for user management and accounting, including creating and maintaining web interfaces necessary for (federal) project management and data availability for users. The department also conducts activities related to data management. This involves operating and continually developing the high-performance storage system as well as conceiving new strategies for data management for users and projects working in the field of research data management.

## Public Relations

*Leader: Sophia Honisch*

Responsible for all areas of HLRS's external communications, from media relations to the management of HLRS's website and social media accounts: It is the main contact point for press and the broader public. The PR department communicates about HLRS's wide range of scientific and engineering disciplines, its research (projects) as well as its services, and disseminates results, new findings, and insights gained.

## Scalable Programming Models and Tools

*Leader: Dr. José Gracia*

Conducts research into parallel programming models and tools to assist development of parallel applications in HPC. Currently the focus is on transparent global address spaces with background data transfers, task-parallelism based on distributed data-dependencies, collective off-loading of I/O operations, and parallel debugging. As a service to HLRS users, the group also maintains part of the software stack related to programming models, debugging, and performance analysis tools.

## Training and Scalable Algorithms

*Leader: Dr.-Ing. Lorenzo Zanon*

The Department of Training and Scalable Algorithms (TASC) organizes and implements HLRS's training activities focusing on a variety of topics in high-performance computing, artificial intelligence, and modeling and simulation. These include compact, high-intensity courses, flexible online learning modules, and public outreach activities. TASC provides an outstanding learning experience by offering training on relevant topics, with up-to-date and audience-focused content, and given by highly-qualified instructors. Besides teaching and outreach activities, TASC conducts research on the development of efficient algorithms for scientific computing applications, and also hosts the User Project Coordination unit.

## Visualization

*Leader: Dr.-Ing. Uwe Wössner*

Supports engineers and scientists in the visual analysis of data produced by simulations on high-performance computers. By providing technologies capable of immersing users in visual representations of their data, the department enables users to interact directly with it, reducing analysis time and enabling new kinds of insights. The department is developing tools for visualization in virtual reality, augmented reality, and has designed a software system for integrating processing steps spread across multiple hardware platforms into a seamless distributed simulation and visualization environment.

© 2026

#### High-Performance Computing Center Stuttgart (HLRS)

University of Stuttgart  
Nobelstraße 19  
70569 Stuttgart, Germany

Tel: +49 711 685-87269  
Fax: +49 711 685-87209  
Email: info@hlrs.de  
Web: www.hlrs.de

#### Director, HLRS

Prof. Dr.-Ing. Dr. h.c. Dr. h.c. Prof. E.h. Michael M. Resch

#### Head, Department of Public Relations

Sophia Honisch

#### Writer and Editor

Christopher M. Williams

#### Photography and images

Unless otherwise indicated, all images property of HLRS.  
p. 9, 10 Benthem Crouwel Architects, in collaboration with Birk Heilmeyer und Frenzel Architekten; p. 16 HPE; p. 18 Openchip;  
p. 21 Courtesy of Everllence, © Sebastian Vollmert;  
p. 25 BMUKN / Annette Riedl; p. 28 NCHC; p. 30 Slaven Vilus;  
p. 31 Tohoku University Cyberscience Center;  
p. 34 / 35 iStock.com / Vadym Sukhorukov; p. 42 Spyro Nita / EPCC (bottom); p. 45 Goethe University Frankfurt;  
p. 47 stock.adobe.com / EpicFade; p. 48 DLR; p. 53 IAG, University of Stuttgart; p. 55 Risselada Lab, TU Dortmund;  
p. 56 Risselada Lab, TU Dortmund

#### Printing

Buch- und Offsetdruckerei H. Heenemann GmbH & Co. KG


#### Design

GROOTHUIS. Gesellschaft der Ideen und Passionen mbH  
für Kommunikation und Medien, Marketing und Gestaltung;  
groothuis.de

#### Institutional affiliations



#### Funding for Hunter and Hawk provided by:



This magazine is printed on paper that has been certified by FSC, the EU Ecolabel, and the Blue Angel Ecolabel.



Follow us on:



LinkedIn



Bluesky

**High-Performance Computing  
Center Stuttgart**

[www.hlrs.de](http://www.hlrs.de)

Cover:

The year 2025 saw the start of construction of HLRS III, a new data center at the High-Performance Computing Center Stuttgart. Designed to meet the more demanding requirements of upcoming generations of super-computers, HLRS III is also conceived with sustainability in mind. A waste heat processing facility will capture the heat produced by HLRS's systems and distribute it to the University of Stuttgart's district heating network, supporting the university's decarbonization efforts.

HLRS is certified for environmental management under the Eco-Management Audit Scheme (EMAS). This magazine has been printed climate-neutral on paper that has been certified by FSC®.