TAKING HPC TO NEW HEIGHTS

For over 20 years, the High-Performance Computing Center Stuttgart (HLRS), located at the University of Stuttgart, has helped leading companies like Daimler and Porsche address their complex engineering problems via the strategic use of high-performance computing (HPC) resources. Today, HLRS is preparing these and other companies to optimize their use of HPC across the business.* Dimensions *recently spoke with Michael Resch, director of HLRS, about current and future trends in HPC.

DIMENSIONS: What services does HLRS provide — and who are your customers?

MICHAEL RESCH: HLRS was established in 1996 as the first national German high-performance computing (HPC) center. It is a research and service organization affiliated with the University of Stuttgart, offering HPC services to both academic users and corporate product development teams. Not only do we offer computing resources that allow scientists and engineers to run massive simulations, but we consult with them about how to solve their problems strategically. We want to ensure that they are achieving fast solve times and cost benefits, but also that they are attaining the high degree of accuracy and fidelity needed to deliver product reliability and safety.

We have a wide range of customers, with two of our largest being Daimler and Porsche. Obviously, automakers face incredibly complex, numerically large engineering problems, especially as the number of electronic components in cars increases. At the same time, car manufacturers work under very stringent regulatory and safety standards. They rely heavily on engineering simulation to solve these problems without the high cost of repeated physical tests. This makes automakers ideal customers for HLRS — but we can help any business that wants to leverage simulation and HPC to combine a high degree of speed with a high degree of product confidence and reliability.

DIMENSIONS: For companies just beginning to leverage the power of high-performance computing, what is the most common pitfall?

MR: Without a doubt, the biggest mistake companies make when they invest in HPC is focusing on the number of technology components, or the size of the cluster. If the product development team is not solving problems fast enough, the typical reaction is, "We need to add more computing nodes or processors." While a certain level of resources is obviously required to run a large simulation, what's more important to achieving speed is the manner in which those resources are deployed.

At HLRS, we have pioneered advanced numerical methods, parallel processing schemes, big data management practices and other innovations that allow any business to maximize the effectiveness of its HPC efforts. We help customers define their computing problems and create a strategic plan for solving them in the most time-, cost- and resource-effective manner.

DIMENSIONS: In addition to lacking a strategic approach, what are some other obstacles engineering teams face in implementing HPC?

MR: Many businesses, especially small and mid-sized companies, are challenged to manage user licenses for all the specialized tools needed for product development. Solutions like ANSYS software are growing increasingly easy to use — and the best practice of pervasive simulation means they are being applied by more people across the business. For technology managers, this means providing flexible access to the software and ensuring it is available when people need it. Fortunately, ANSYS has pioneered a concept called elastic licensing to address this.

The growing popularity of cloud computing is also helping these same small and mid-sized companies access technology on a flexible basis. But that brings its own challenges, primarily the need for uninterrupted connectivity and fast data-streaming speeds.

To maximize their use of high-performance computing, companies of all sizes need a plan for sustained performance over time. In today's competitive market, product development can't be held back by employee downtime caused by a lack of licenses, or by failed internet connections. These are administrative issues our experts can help manage, whether we are licensing and managing the technology for them, or simply devising an optimal plan for their internal ownership of resources.

DIMENSIONS: What do you see as the single greatest trend in HPC today?

MR: Probably the most important trend we're witnessing is the democratization of high-performance computing, as more people across the business need access to HPC resources to support pervasive simulation — but also to carry out other work. High-performance computing used to be the domain of engineers and researchers, but today the whole enterprise is becoming more digital in nature. Employees in many functions now need to manage large amounts of data and solve complex problems — and that means they need processing power and speed.

With only about 20 to 30 supercomputing centers worldwide and fast-growing demand for HPC resources, I expect cloud services to expand to answer this market need. As access to high-performance computing increases, costs will be driven down. In three to five years, I expect that most professionals will have daily access to HPC resources, including those at small and mid-sized companies. Instead of being seen as a specialized activity, HPC will be integrated into many daily business activities and will be a part of many corporate functions — not just product development.

DIMENSIONS: Looking ahead, do you see the speed of high-performance computing increasing significantly?

MR: I believe that processing speeds may continue to increase for the next several years, but in the 2020s we are actually going to see a stabilization of high-performance computing speeds. Just as auto engineers gradually shifted their focus from horsepower gains to overall performance gains, computing experts are going to focus more on managing data and solving problems more strategically — not just on processing speeds.

There is so much data available today, and that volume is only going to grow as the world becomes ever more digital in nature. As one example, factories are going to be much more automated, which will generate a huge volume of numerical data about production costs and efficiency. These insights will have enormous value across the

EMPLOYEES IN MANY FUNCTIONS NOW NEED TO MANAGE LARGE AMOUNTS OF DATA AND SOLVE COMPLEX PROBLEMS — AND THAT MEANS THEY NEED PROCESSING POWER AND SPEED.

business, including the product development function. In keeping with the concept of pervasive simulation, this manufacturing data will be considered during product engineering to keep production costs low and efficiency high. But which data is important? And which is trivial?

The same applies to the emerging concept of the digital twin — a simulated product system that mirrors a real product that is working in the field. As engineers apply real-world physical forces and environmental conditions to the model, a very large amount of performance data is generated. Someone needs to look at that information and figure out what is important for predictive maintenance, repair, future product development and other applications.

Looking ahead, the high-performance computing industry is going to focus on managing and applying data more strategically, resulting in smarter simulations, not necessarily faster ones. Today, the emphasis for many companies is on speed — "How soon can we get the results?" — but, in the near future, more emphasis will be placed on the value of the results. Not just for the engineering team, but the business as a whole.

DIMENSIONS: What role do you see HLRS playing in the evolution of HPC?

MR: I see HLRS playing an essential role because we partner with leading software companies like ANSYS to understand simulation technology, which relies on HPC, as well as with leading companies that use ANSYS solutions in an HPC environment on a daily basis. By helping to optimize the performance of ANSYS software and other HPC-enabled tools, we can meet industry's real-world needs.

As businesses grow increasingly digital and rely on HPC more and more, HPC centers must support sound decision-making for our corporate partners. Companies need to address data security, cloud access, software licensing and other technical issues — but, even more important, they need to ensure they're getting the maximum return from their technology investments.

Experts at HPC centers have worked with some of the world's leading companies to solve these problems already. There are a number of best practices that can be applied to help companies make significant improvements in their product development processes. For example, HLRS recently helped a German manufacturer use HPC-enabled simulation to reduce physical testing costs by two-thirds. Over the course of 20 years, HLRS has accumulated many "lessons learned" that we can leverage to help customers make similar improvements.

HLRS serves as a bridge between how companies are using HPC today versus what is possible in the future. That's a role that will become even more important as the number of processors worldwide grows and HPC becomes ever more commonplace in the business world.

HLRS at a Glance

Founded in 1996

Number of Employees: 125

About Michael Resch

Since 2003, Professor Michael Resch has been director of the High-Performance Computing Center Stuttgart (HLRS), one of the fastest civil computing systems in Europe. He also heads the University of Stuttgart's Institute for High-Performance Computing. Resch has received numerous awards, including the National Science Foundation's Award for High-Performance Distributed Computing, the HPC Challenge Award and the Microsoft Early Contributor Award.