Case Study

Transforming HPC Strategy with The AMD HPC Innovation Lab

Clarios accelerates HPC transformation with Oakwood's AMD Innovation Lab, gaining critical insight and confidence to move complex engineering workloads to Azure.

Microsoft Solutions partner

Introduction

Clarios, a global leader in advanced energy storage solutions, manufactures one in every three vehicle batteries in the world. As a company grounded in innovation, its engineering team relies heavily on high-performance computing (HPC) to simulate, test, and refine electrochemical systems for the next generation of smart and sustainable vehicles.

Despite their deep experience with HPC, Clarios had been cautious when it came to cloud adoption. Their existing on-premises infrastructure—custom-built and fine-tuned over time—continued to support the bulk of their simulation workloads, including those running complex solvers such as Ansys Fluent and LS-DYNA. While effective, the environment was becoming a constraint. Scaling compute, experimenting with new architectures, and accommodating shifting R&D priorities required a level of agility that their on-prem systems simply couldn't offer. Cloud HPC was already on the radar—but in their minds, it was a project for next year.

That all changed when Oakwood introduced the AMD HPC Innovation Lab. Developed in collaboration with AMD and Microsoft, the lab provided a tailored, no-risk environment where Clarios could explore Azure HPC in a hands-on, real-world context. It wasn't just a proof of concept—it was a strategic accelerator. By eliminating many of the traditional barriers to cloud adoption, the Innovation Lab gave Clarios a safe, structured way to experiment, test, and validate whether Azure could meet their most demanding compute needs.

Microsoft Solutions Partner
Data & AI
Azure

Microsoft Solutions Partner
Infrastructure
Azure

Microsoft Solutions Partner
Digital & App Innovation
Azure

Microsoft Solutions Partner
Modern Work

1001 Craig Rd. Suite 305 | St. Louis, MO 63146
p. 314.824.3000 | e. marketing@oakwoodsys.com

Business Challenge

For Clarios, the decision to move HPC workloads to the cloud wasn't just about performance metrics or financial modeling—it was about operational continuity. Their current systems were deeply embedded into the company's engineering workflows, with carefully orchestrated pre- and post-processing pipelines, licensed ISV software, and customized hardware configurations. Engineers were intimately familiar with the tools, and any deviation from the known environment risked introducing friction—or worse, downtime.

There was also skepticism around whether Azure's infrastructure, particularly when powered by AMD EPYC™ processors, could handle the nuanced demands of their simulations. Clarios needed to be sure that solver performance, I/O throughput, and scalability would meet or exceed what they had grown accustomed to. A traditional migration approach felt too disruptive. What they needed instead was a sandbox—a place to experiment without pressure, to measure without committing.

Additionally, Clarios wanted to understand the broader operational impact of working in the cloud. They weren't just testing compute speeds; they were evaluating the full user experience. How easy would it be to provision infrastructure? Would their engineers require retraining? Could jobs be managed independently of IT? Could cloud open the door to innovation—or would it simply become another layer of complexity?

These were the questions they needed answered before moving forward.

Solution

Oakwood's AMD HPC Innovation Lab was purpose-built to provide those answers.

Through a browser-based portal, Clarios engineers were given access to a curated set of "t-shirt-sized" environments—small, medium, and large templates designed around AMD EPYC-backed virtual machines running on Azure. Within these environments, Oakwood had pre-integrated the full Ansys solver suite, eliminating the need for Clarios to install, license, or configure any applications themselves. The lab was designed for rapid interaction: engineers could upload solver files, run simulations, monitor performance, and compare outcomes across different configurations, all within a matter of hours.

What set the lab apart was its turnkey nature. Everything—from ISV licensing to secure access, telemetry tracking, and backend provisioning—was managed by Oakwood. Clarios engineers didn't need to engage IT, request hardware, or navigate the typical roadblocks associated with building new environments. They could simply log in and begin testing.

The ISVs currently integrated include Altair (AcuSolve and Radioss), Ansys (Fluent, HFSS, LS-DYNA, Mechanical, Maxwell), and Siemens (Simcenter STAR-CCM+)

Outcome

The Innovation Lab delivered immediate value. For several key workloads, Clarios saw noticeable improvements in runtime performance. Simulations that once took hours to complete on-prem were now completing faster in the cloud, with AMD-backed instances proving especially effective for compute-

1001 Craig Rd. Suite 305 | St. Louis, MO 63146
p. 314.824.3000 | e. marketing@oakwoodsys.com

intensive models. This translated into tangible productivity gains—engineers were able to iterate more quickly, test more variables, and respond to design changes in less time.

More importantly, the lab helped Clarios make sense of what to move first. Through analysis of their telemetry data, the Oakwood team helped Clarios identify which workloads were ideal for cloud bursting—such as infrequent, high-demand simulations that didn't justify permanent infrastructure investment. Conversely, more tightly coupled MPI workloads were flagged as candidates for future refinement, giving Clarios a roadmap that balanced impact and feasibility.

The engagement also surfaced insights that extended beyond raw performance. Engineers commented on the simplicity and intuitiveness of the lab interface, noting that it allowed them to focus entirely on simulation outcomes rather than infrastructure management. The experience gave Clarios a glimpse into what an operational cloud HPC model could look like governed by self-service, supported by automation, and scalable by design.

As shown in the chart to the right, the Innovation Lab delivered a dramatic improvement in computation time. Workloads that previously took 13–14 hours to run on Clarios' on-premises infrastructure were completed in just a few hours — a reduction of over 85%!

Engineers were especially impressed by the lab's accuracy during Newton-Raphson iterations on large matrices, and the availability of journal files proved invaluable for debugging and error tracing.

**These results to the right are based off of data from previous runs of smaller workloads extrapolated for predictions in larger model run times and has proven accurate in actual sample testing for workload projection times. The actual test model run times are reflected in the tables below.*

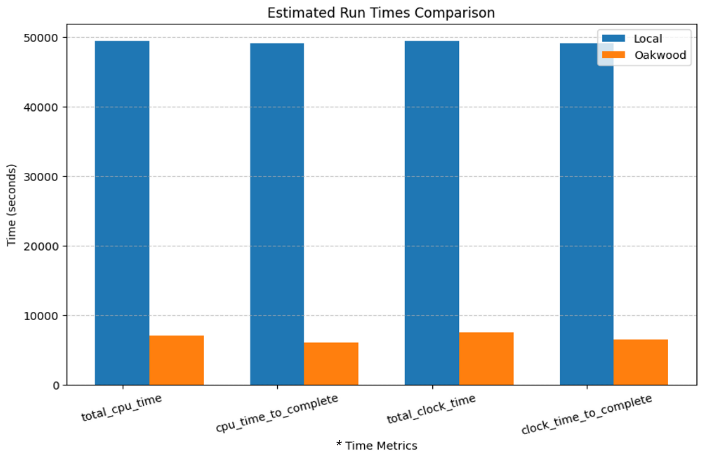

The table below presents a representative subset of solver runs executed during the Clarios engagement. Each entry highlights the specific ISV used, the selected workload size, solver duration within the Azure environment, and the corresponding "all-in" Azure consumption costs. This data proved essential for Clarios in evaluating the performance-to-cost ratio of cloud-based HPC compared to their existing on-premises systems. By analyzing run times and associated costs across different configurations, Clarios was able to identify which workloads were best suited for cloud bursting versus those better retained on-prem. This sampling provided a data-driven foundation for shaping their future hybrid HPC strategy.

ISV / Solver	Workload Size	Output File Size (MB)	Duration (min:sec)	Cost (USD)	SKU
Ansys Fluent	Large	20.8 MB	1:15.19	$3.20	Standard_HB176-96rs_v4
Ansys Fluent	Small	34.96 MB	1:06.17	$3.77	Standard_HB176-24rs_v4
Ansys LS-DYNA	Large	47.68 MB	2:04.45	$23.69	Standard_HB176rs_v4
Ansys LS-DYNA	Large	8.34 MB	1:27.38	$29.88	Standard_HB176rs_v4

Table: Sample Solver Run Performance & Cost

The configurations used consisted of three workload sizes or "T-Shirt sizes" for ease of use for the current system and are easily customizable to any number of configurations or SKU's available. For the current system each solver shared 3 choices varying by core count utilizing AMD HBv4 VMs which feature up to 176 AMD EPYC™ 9V33X ("Genoa-X") CPU cores with AMD's 3D V-Cache, clock frequencies up to 3.7 GHz, and no simultaneous multithreading.

Configuration	SKU	Cores	Memory	Local Storage
Small	Standard_HB176-24rs_v4	24	768 GB	2280 GB – temp+NVMe
Medium	Standard_HB176-48rs_v4	48	768 GB	2280 GB – temp+NVMe
Large	Standard_HB176-96rs_v4	96	768 GB	2280 GB – temp+NVMe

Table: Sample VM configurations