AMBER | WARR rocketry

CASE STUDY

From 7 Hours to 20 Minutes: How WARR Rocketry Simulations Took Off with NVIDIA Hardware

WARR team at the Technical University of Munich accelerates rocket simulations with NVIDIA GPU-powered HPC, achieving full-scale analyses in minutes

NVIDIA Elite Partner

EXECUTIVE SUMMARY

Challenges Facing the Industry

WARR, a leading aerospace student team at the Technical University of Munich (TUM), faced a critical bottleneck in their rocket development: complex simulations were consuming over 7 hours each (often crashing due to memory limits) or were not feasible on their previous hardware. This drastically limited their ability to analyze designs and iterate quickly.

In early 2025, WARR partnered with AMBER, a high-performance computing solutions provider, to deploy a state-of-the-art on-premises GPU cluster powered by NVIDIA technology. The impact was immediate and profound. Simulations that once took an entire workday now run in ~20 minutes, a ~19× speed-up that has revolutionized WARR's workflow. The team can now routinely run full-geometry, high-fidelity rocket simulations in-house, leveraging ~45 million cell models that capture physics with unprecedented detail.

This case study examines how AMBER's solution addressed WARR's challenges and delivered tangible benefits: more iterations in less time, higher accuracy in design data, and faster innovation cycles, ultimately helping WARR build lighter, more reliable rockets.

"Innovation needs the right infrastructure and the right people. By equipping WARR with a cutting-edge NVIDIA A40 GPU server, we empower the next generation of aerospace engineers to push the boundaries of what's possible. High-performance computing is the catalyst that allows their engineering dreams to take flight."

Michael Rechenmacher, CEO of AMBER

INTRODUCTION TO AMBER AND WARR

HPC Meets Rocketry

AMBER and WARR: Transforming aerospace simulations with high performance computing

WARR

WARR (Scientific Workgroup for Rocketry and Spaceflight) is the student aerospace team at the Technical University of Munich. The team develops cutting-edge rocketry, satellite, and space robotics projects, and has earned multiple international competition wins, making WARR one of the most successful student aerospace groups in Europe. In recent years, WARR's rocketry division achieved record-breaking milestones, including winning the European Rocketry Challenge in 2023 and 2024, in several categories. These accomplishments underscore WARR's reputation for innovation in student-led space engineering.

AMBER

AMBER AI & Data Science Solutions GmbH is a leading provider of cutting-edge full-stack AI solutions and high-performance computing. As an NVIDIA Elite Partner, AMBER collaborates with industry leaders to deliver state-of-the-art, tailored solutions that accelerate innovation and reduce costs. With ISO-certified quality standards, AMBER ensures reliability and excellence in every project. Founded in 2006 as FluiDyna, AMBER has been at the forefront of artificial intelligence innovation since 2008.

With a shared passion for technological excellence, AMBER and WARR partnered to overcome the computational barriers in WARR's ambitious aerospace simulations.

CHALLENGES

Why WARR Needed On-Prem HPC

Despite WARR's engineering talent and ambition, their computational resources fell short
Key challenges included:

Insufficient Compute Power

WARR's complex CFD and FEA simulations were often impossible on their existing hardware. They either had to simplify models or rely on external university clusters, which introduced long queue times and inflexibility. In-house, they could only run partial simulations (e.g. 1/8th of a rocket cross-section) due to limited performance. This lack of on-demand horsepower was stifling their progress.

Prohibitive Simulation Times

When simulations did run on a personal workstation, they frequently took over 7 hours for a single cooling-channel flow analysis, and even then ran the risk of crashing due to insufficient RAM. Such lengthy runtimes made iterative testing impractical and delayed critical design feedback.

Limited Fidelity & Accuracy

Because of hardware limits, WARR resorted to one-eighth symmetry models and low-resolution meshes. This meant important non-symmetric phenomena were missed, and results were less realistic. Key aerodynamic properties (e.g. drag coefficients) had to be estimated from handbook curves rather than computed directly, undermining confidence in the data.

Hindered Design Iteration

The combination of slow runs and coarse models meant design iterations were severely constrained. Simulations could only be performed after major design decisions, instead of guiding them. WARR couldn't readily test multiple design variants or refine concepts in real-time, slowing their R&D cycles.

NEEDS

Building the Infrastructure Rocketry Demanded

Dedicated on-prem GPUs. Speed in minutes. Full-geometry fidelity. Seamless access and support.

Dedicated On-prem GPU Power

WARR needed a server-class GPU cluster they fully controlled. They aimed to eliminate queues, quotas, and surprises. It had to run big CFD and FEA jobs reliably on demand. The goal was to own the schedule and the results. NVIDIA-class accelerators with ample memory were non-negotiable. This became the backbone that made everything else possible.

Speed that Turns Simulation into a Design Tool

Runtimes had to drop from hours to minutes. Designers needed to iterate multiple variants in a single work session. DOE sweeps and quick A/Bs had to become routine. Feedback needed to land while decisions were still being made. That was how simulation shifted from validation to creation. Speed removed hesitation and changed the culture.

Full-Geometry Fidelity at Scale

They needed to stop slicing models to fit weak hardware. The team had to run entire rockets at high resolution. Shocks, loads, turbulence, and interactions needed to be resolved, not assumed. That delivered real drag, pressure drop, and thermal numbers. Better inputs drove better designs. Precision replaced guesswork.

Reliable Partners for Implementation and Support

Everyone who built needed to launch jobs remotely, securely, and fast. The environment had to integrate with the current toolchain and data flow. Templates, job scripts, and storage policies had to reduce friction. Monitoring and maintenance needed to be predictable. Experts had to be on call when runs failed or needed tuning. Knowledge transfer kept the capability alive after handover.

AMBER'S SOLUTION

Powering WARR Rocketry with a Turnkey HPC Stack

Built for full-geometry CFD and FEA. Faster runtimes. Remote-first. Scalable and supported.

AMBER delivered a turnkey HPC stack tailored for simulation-heavy work. In January 2025, the team installed a dedicated multi-GPU server at WARR. It used NVIDIA A40 accelerators (48 GB per GPU) with high-core CPUs, ample RAM, and fast SSDs for pre and post processing. Compared to the old PCs, the new cluster provided orders-of-magnitude more parallel compute, so CFD and FEA finally ran at full speed.

The server joined the university network with proper security and was dedicated to WARR, which removed queues and shared slowdowns. It was built to scale, allowing future GPU or node additions. AMBER also provided ongoing support. Monitoring, updates, and troubleshooting kept the system reliable. In short, AMBER delivered hardware plus integration and expertise, enabling WARR to push limits immediately.

Reasons for a A40-Based HPC Stack

- GPU memory headroom
- Multi-GPU throughput
- CFD/FEA optimized
- Local data control
- Queue-free capacity
- Remote, multi-user
- Built to scale
- Enterprise-grade reliability

BENEFITS

From Estimates To Evidence

To summarize the benefits of the new NVIDIA GPU-powered HPC solution for WARR, the key improvements are outlined below:

Accelerated Iteration

A seven-hour CFD run became 22 minutes, or 19× faster. This meant that ideas were tested the same day, not the next morning. Simulation moved into the design meeting and started steering choices in real time.

Full-Geometry Fidelity

WARR stopped slicing rockets into eighths. 45-million-cell models captured shocks, fin interactions, and turbulence across the entire vehicle. Decisions were based on what the physics actually did, not what symmetry assumptions hoped for.

Queue-Free Control

The A40 stack lived on-prem, so the team owned the schedule. No LRZ queues, no surprise throttling. Students submitted from laptops, pulled results quickly, and kept momentum.

Lighter Parts, Higher Efficiency

Precise aero loads translated into leaner structures with confidence. Verified pressure-loss data guided pump and engine tweaks that saved energy. Weight and performance trades became evidence-driven instead of safe guesses.

Reliable Flight Predictions

Transonic through supersonic behavior was modeled in-house with the needed resolution. Drag and stability estimates tightened, and apogee forecasts stopped drifting. Mission planning stepped up from probable to predictable.

Team-Wide Productivity

HPC turned into a daily habit. Templates, tuned job scripts, and a clean remote workflow removed friction. The result was a faster learning loop across WARR, and a team that ships better engineering sooner.

RESULTS

Results That Move Rockets

19× faster runtimes, ~45M-cell models, full-geometry CFD in-house, and an integrated daily workflow

The NVIDIA A40 GPU server delivered a step-change in compute. A first 45-million-cell cooling-channel CFD finished in ~22 minutes instead of ~7 hours on a PC. That ~19× runtime reduction let the team move from single overnight runs to same-day iteration, and it unlocked full-geometry rocket cases that previously sat in a queue or failed on memory.

With in-house capacity, WARR ran entire vehicles at higher resolution. Shocks, fin–body interactions, turbulence, and pressure gradients were captured across transonic and supersonic regimes, not approximated. Simulation shifted into the design phase, so variants and DOE sweeps informed choices in real time, and validation no longer lagged design by days.

On-prem deployment gave predictable, queue-free throughput and local data control. Remote multi-user access integrated cleanly into student workflows. With AMBER's tuning and support, the stack stayed optimized, which cut friction and kept the learning loop tight. The outcome was practical engineering wins: lighter fin-can structures from precise aero loads, lower pressure losses in pumps and plumbing, and tighter apogee and stability predictions that improved mission planning.

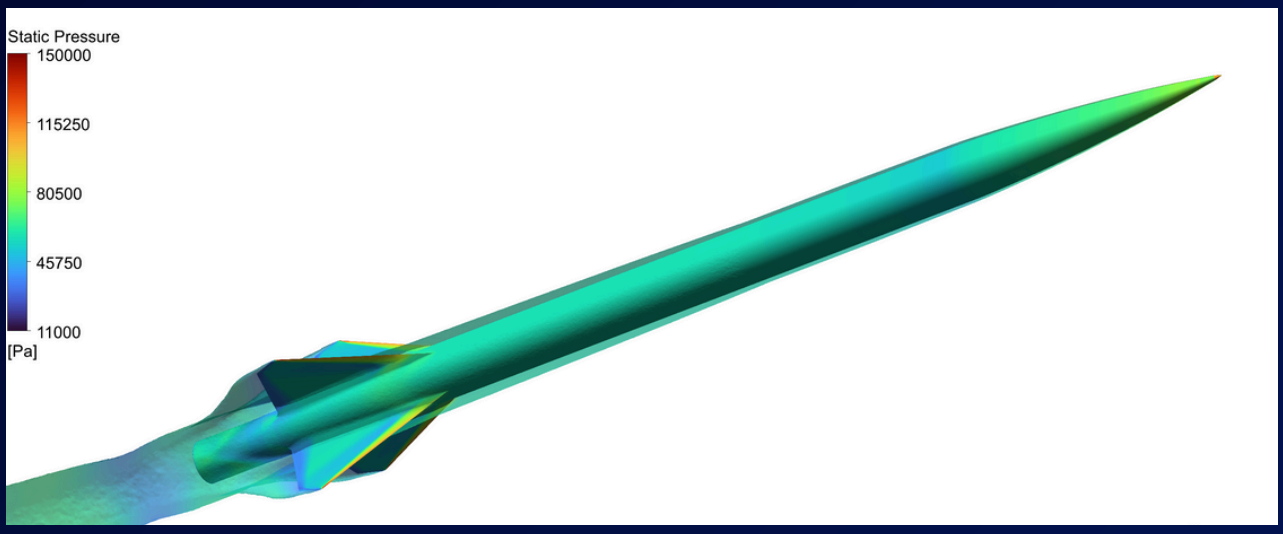

A high-fidelity rocket aerodynamics simulation that WARR can now perform in-house on the NVIDIA GPU-powered cluster. Such detailed full-rocket simulations were previously only achievable on an external supercomputing cluster, due to the limitations of WARR's old hardware.

TESTIMONIALS

Voices

Insights from WARR & AMBER on the transformative impact of the NVIDIA A40 GPU:

WARR rocketry

"Access to AMBER's HPC environment made it possible for us to simulate rocket configurations we couldn't handle locally. The results were faster, more accurate, and helped us validate our design with confidence."

Richard Emeder, Former Head of Rocketry, WARR

AMBER

"AMBER's collaboration with WARR highlights how industry-grade technology can empower young innovators. Working with WARR showcases how cutting-edge computing can drive student innovation. By deploying NVIDIA A40 GPUs and providing tailored support, we're proud to assist such a talented team with our expertise."

Michael Rechenmacher, Founder, CEO

CONTACT

Your Contact for More Information

AMBER AI & Data Science Solutions GmbH

Grünwalder Weg 32
D-82041 Oberhaching

T +49 89 413 2733-00
F +49 89 413 2733-09
E info@amber.eu