



## Argonne Deploys Custom AMAX ClusterMax SuperG GPGPU Cluster

### THE CHALLENGE

[Argonne National Laboratory](#) is one of the U.S. Department of Energy's oldest and largest national laboratories for science and engineering research. Since 1990, Argonne's mission has been to apply a unique mix of world-class science & engineering to deliver innovative research and technologies with the goal of addressing the most important scientific and societal needs of the nation. Argonne develops transformational energy storage systems that enable and enhance electric-drive vehicles, is establishing the next generation of alternative energy sources to promote energy independence, and engineers advanced reactor and fuel cycle systems for nuclear energy including advanced modeling and simulation methods. This type of research depends heavily on computational performance and left Argonne with the need to deploy a large multi-GPGPU cluster solution.

“The AMAX SuperG GPGPU cluster allows national labs like Argonne to take full advantage of the game-changing benefits of GPU computing, enabling them to achieve results much faster than ever before, and allowing them to pursue new, more complex research projects that otherwise would have been impossible.

Sumit Gupta, Senior Director  
of the Tesla Business, NVIDIA

### THE PROBLEM

Until recently, Argonne engineers have been utilizing a CPU-based cluster, spanning multiple racks which required a substantial power & cooling footprint to crunch their large-scale computations – mainly dynamics differential equations which help to describe the time evolution of large systems as well as fluid dynamics and simulations for super conductivity. These equations are extremely compute-intensive and Argonne engineers would typically have to wait months to see the results. Having tested their computations with smaller-scale GPU workstations running on consumer-based GPUs, Argonne's engineers were able to confirm their hopes – that GPUs would exponentially speed up their equations. Understandably, the workstations could only manage to run a few positives at a time. The solution goal then for Argonne was to engineer a high-performance, extremely-dense GPGPU cluster solution that could speed up their computations by 100x while being simple to manage and easy to use with ample room to scale.

### THE SOLUTION

AMAX provided all the necessary engineering expertise to deliver a complete turnkey working solution to Argonne based on its unique specifications of very high-density GPUs with minimal CPUs. In order to ensure a successful deployment, AMAX's team of engineers and technical consultants worked with Argonne engineers to fully understand the requirements of the project, conducting a thorough analysis of Argonne's application and usage to come up with a suitable configuration. Gathered from the discussions, AMAX designed a GPGPU cluster with 42 NVIDIA® Tesla® C2070 GPUs based



AMAX was the only vendor that could provide a computational GPGPU cluster with our exact specifications. They engineered a solution that not only met, but exceeded our wildest expectations.

Andreas Glatz, Argonne National Laboratory



off of the proven [AMAX ClusterMax SuperG](#) GPU platform, designed to not only match but optimize Argonne's specific requirements for density, performance, budget and space with sensible scalability in mind. The SuperG GPGPU cluster was built and tested in AMAX's ISO:9001 facility by a specialized team of certified HPC engineers with years of experience in network & infrastructure design and integration expertise, who helped to fine-tune the multi-GPU cluster solution specifically for Argonne's applications before deployment. Once the crated GPU cluster arrived, Argonne's engineers found deployment both simple and painless.

What mattered most to Argonne, however, was that this system would perform as promised. In fact, the AMAX GPGPU cluster exceeded the desired performance Argonne had hoped to achieve. "When we started this project, we were running all codes on CPUs and now we've converted those codes to run on the GPUs," said Andreas Glatz, Systems Engineer, Argonne. "With the AMAX SuperG GPGPU cluster, the performance factor increase has been roughly 120-150x!" Andreas went on to remark, "Our computational equations typically took months to return a result, and now with the AMAX SuperG GPGPU cluster we can get those results within hours." Since deploying the AMAX cluster, Argonne has published two papers filled with computational work that would have taken their CPU-based cluster over half a year to complete, while it took the AMAX SuperG GPGPU cluster less than 1 week to complete all computations. "This is a great advancement, and this performance has been without us tweaking much of anything. The AMAX SuperG cluster has significantly sped up our research and ultimately our publications come out much faster."

"The AMAX SuperG GPGPU cluster allows national labs like Argonne to take full advantage of the game-changing benefits of GPU computing, enabling them to achieve results much faster than ever before, and allowing them to pursue new, more complex research projects that otherwise would have been impossible," remarked Sumit Gupta, Senior Director of the Tesla Business, NVIDIA.

AMAX was chosen as the supplier for the Argonne National Laboratory due to excellent sales support, a complete understanding of Argonne's requirements to create an optimized design, an excellent reputation for manufacturing quality and very competitive pricing. AMAX's well-trained IT consultants conducted a thorough analysis of Argonne's application requirements to propose an optimized solution to find the solution sweet spot that would satisfy all requirements – from density & performance to price & reliability. "At Argonne, we are interested in all kinds of properties of energy from energy transport and energy conversion, to all types of problems related to sorting and creation of energy for large systems," added Glatz. "These are exactly the kinds of problems that we can now easily solve with the AMAX SuperG GPGPU cluster."