



MAXIMIZING DATA RESOURCES

ISI'S VISTA LAB SPEEDS UP MACHINE LEARNING RESEARCH AND LOWERS EXPENSES

About the Customer

The Information Sciences Institute (ISI) is a world leader in research and development of advanced information processing, computing, and communications technologies. A unit of the University of Southern California's highly ranked Viterbi School of Engineering, ISI is one of the nation's largest, most successful university-affiliated computer research institutes. A pacesetter for 45 years, ISI invented and developed the Domain Name System (.com, .net, etc.) and the text-to-numerical-IP-address system that are essential to internet functionality. Work conducted at ISI includes machine learning and AI, cyber security, novel electronics, HPC architectures, and quantum computing.



With Altair Grid Engine, we have an infrastructure that schedules workloads to GPUs. We operate our infrastructure at 95% capacity with lower overall costs.

Stephen Rawls, Research Analyst, Information Sciences Institute



Their Challenge

Stephen Rawls, programmer and research analyst at ISI's Video, Image, Speech and Text Analytics (VISTA) lab, makes **extensive use of machine learning** in a variety of application areas such as facial identification, natural language processing, and handwriting recognition. VISTA users share a compute cluster with several other research groups. When Rawls joined VISTA, they had been distributing project resources on an ad hoc and manual basis, with resources often sitting idle. Increasing demand required a more sophisticated way to distribute and manage resources among multiple users. When VISTA installed an open-source scheduler they experienced difficulties with expensive GPU resources, and a significant amount of time was being taken away from research to troubleshoot. "We needed a **reliable, powerful workload management platform** that would enhance performance and have the ability to run complex, diverse workloads across multiple users within the entire ISI organization," said Rawls.

Our Solution

VISTA selected Altair® Grid Engine® to manage growing infrastructure and accelerate its machine learning research. Rawls cites key contributing factors for VISTA's decision to transition to Altair Grid Engine over other vendors: **built-in advanced GPU support, detailed documentation, and ongoing product upgrades**. Rawls also appreciated Altair's **customer support**. "The basis for artificial intelligence and machine learning research is to create neural networks to help solve problems," said Rawls. "But training artificial neural networks requires a lot of data and a significant amount of GPU time for tuning up parameters and running multiple sets of experiments simultaneously." The ISI infrastructure currently has over 100 total NVIDIA GPUs. Altair Grid Engine allows the organization to configure a project-based share tree and scale GPU-based frameworks.

Results

ISI's VISTA team is **advancing the state of research for facial recognition**, a technology with huge implications for security and commerce. One of the ways ISI scientists are teaching computers how to recognize faces is by extracting 68 facial landmarks including eyebrows, nose, and mouth. Using code and algorithms, the project uses deep learning to teach computers to mimic how neurons in the brain talk to each other, and experiments required weeks of compute time to run terabytes of data. One experiment involved processing over 3 million images. VISTA utilized Altair Grid Engine to set up the parallel processing and **manage dependencies for the entire process without fail**.

ISI needed to scale and optimize its resources to meet growing demand for machine learning computation