UNIVA

CASE STUDY

# Information Sciences Institute maximizes distributed data resources with Univa Grid Engine

## ISI's VISTA lab conducts machine learning research faster, more efficiently, and with lower overall costs.

USC Viterbi School of Engineering

The Information Sciences Institute (ISI) is a world leader in research and development of advanced information processing, computer 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 by more than 350 engineers, research scientists, graduate students and staff currently includes machine learning and artificial intelligence, cyber security, novel electronics, high performance computing architectures and quantum computing.

### THE CHALLENGE

**Scale resources to meet growing demand for machine learning computations**

Stephen Rawls, programmer and research analyst at ISI, makes extensive use of machine learning in a variety of application areas such as facial identification, natural language processing, and handwriting recognition using optical character recognition (OCR). As a Ph.D. student, Stephen's focus is OCR for handwritten documents—this technology enables the conversion of different types of documents such as scanned paper documents, PDF files, or images into editable and searchable data.

Rawls is a member of the Video, Image, Speech and Text Analytics (VISTA) lab at ISI. They share a compute cluster collectively with a team comprised of seven to ten disparate research groups. When Rawls arrived at VISTA four years ago, the organization distributed project resources on an ad hoc and manual basis. The segregated groups were assigned to dedicated machines and when jobs were not in queue, the resources sat idle.

As demand increased, it became evident that a more sophisticated way to distribute and manage resources among its multiple users was required. His challenge was to build an infrastructure that could scale and optimize comprehensive data simulations.

Rawls had used an open source version of Grid Engine at his previous employer. When VISTA installed an open source scheduler they immediately experienced difficulties with their expensive GPU resources. The team tried a few open source modifications to work with GPUs, but all had various downsides. A significant amount of time was being taken away from research to trouble shoot an aging open source scheduler.

"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.

## THE SOLUTION

**VISTA selected Univa Grid Engine to manage growing infrastructure and accelerate its machine learning research**

Rawls cites key contributing factors for VISTA's decision to transition to Univa Grid Engine over other vendors: built-in advanced GPU support, detailed documentation and ongoing product upgrades. Rawls also appreciated Univa'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 require 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 Nvidia 100 GPUs in total, with many machines having eight 1080TI GPUs each and various other configurations. Univa Grid Engine allows the organization to configure a project-based share tree and scale GPU-based frameworks.

ISI's VISTA Group has spent the past three years advancing the state of research and publishing results 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 facial landmarks. In ISI's system there are 68 landmarks such as 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.

Specific experiments required weeks of compute time to run terabytes of data. For instance, one experiment involved the image processing of over 3 million images. VISTA utilized Univa Grid Engine to set up the parallel processing and manage dependencies for the entire process without fail.

"With Univa 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, ISI


## About Univa

Univa is the leading independent provider of software-defined computing infrastructure and workload orchestration solutions. Univa's intelligent cluster management software increases efficiency while accelerating enterprise migration to hybrid clouds. Millions of compute cores are currently managed by Univa products in industries such as life sciences, manufacturing, oil and gas, transportation and financial services. We help hundreds of companies to manage thousands of applications and run billions of tasks every day. Univa is headquartered in Chicago, with offices in Toronto and Munich. For more information, please visit **www.univa.com**.

UNIVA

**Univa Corporation** 2300 North Barrington Road, Suite 400, Hoffman Estates, IL, 60195 USA
Tel: +1.647.478.5901 **www.univa.com**