

# **GPU Consolidation for Image Generation**

**Tim Woodard**

## **ABSTRACT**

(500 words max)

One of the predominant trends in high-performance computing (HPC) is hardware consolidation, particularly with respect to graphics processing units (GPUs). Not only is greater performance-per-watt possible with modern GPUs, but also greater density. Data centers are increasingly built on systems that combine multiple GPUs in a single system to reduce operational costs and maximize computational throughput by eliminating bottlenecks.

As with HPC environments, image generators (IGs) for training and simulation can also benefit from hardware consolidation. By combining multiple GPUs in a single system, it is possible to efficiently drive large-scale display systems with far less hardware. Numerous factors must be taken into consideration to utilize consolidated GPU resources most effectively, and the ideal configuration is dependent on the specific use-case. In this paper, we discuss the advantages of increasing GPU density for image generation. We also cover the various technologies that are relevant, including motherboard topology, GPU affinity, OS considerations, window system management, and synchronization. Several high-level architectures are presented, along with the advantages and disadvantages of each.

## **BIO**

### ***PRIMARY AUTHOR***

Tim Woodard is a Senior Solutions Architect with NVIDIA's Pro-Viz group and has over 20 years of experience designing and developing software architectures for real-time training and simulation visual systems using modern OpenGL techniques, advanced C++, and Agile development processes. He has received patents for run-time simulator database generation technology and has published and presented papers at GTC, I/ITSEC, IMAGE, ASQ, and ITEC.