Nvidia Smashes Visual Record with New GPU Server

Combines Four Quadro Gpus in a Standard 1-U Server Configuration for Remote Graphics and Offline Rendering Applications

August 08, 2007 - 01:49

NVIDIA Corporation launched NVIDIA Quadro® Plex Visual Computing System (VCS) Model S4, a graphics server featuring a record number of graphics processing units (GPUs) in a standard 1U server form factor. The new server contains more Quadro graphics processing power in less space than has ever been achieved before, enabling interactive visualization of extremely large-scale data sets by multiple users with a minimal footprint.

Designed for data-centric applications such as off-line rendering, remote graphics applications and embedded visual computing, the Quadro Plex Model S4 server takes visual compute density to new extremes by incorporating four NVIDIA Quadro® FX 5600 professional GPUs in an industry-standard 1U rack configuration; dynamically allocating compute, geometry, shading and pixel processing power for optimized GPU performance. GPU computing for visualization solves programming functionality available in the NVIDIA CUDA™ tool suite.

The Quadro Plex Model S4 is designed to serve high-quality interactive visualization of large models to multiple users. Its 6 GB frame buffer (1.5 GB per GPU) can process large textures with full-screen antialiasing while Shader Model 4.0 vertex and pixel programmability delivers ultra-realistic effects for OpenGL and DirectX 10 applications. The server's accelerated 3D texture performance provides greater interactivity for the visualization of large volumetric datasets often required in energy and scientific research.

The Quadro Plex Series

Until now, NVIDIA's Quadro Plex product line has consisted of desk side units that can also be mounted side-by-side in a 3U rack space. The Model S4 server complements the rest of the product line by offering extreme compute density in 1U housing for solving the largest visualization and high-performance computing problems.

The Quadro Plex VCS Model II is intended primarily for powering large arrays of monitors. Featuring four Quadro FX 4500 GPUs with a 512 MB frame buffer per GPU, the Model II can power up to eight frame-synched, dual-link DVI displays.

The Quadro Plex Model III adds high-definition SDI output by replacing the GPUs in the Model II with two Quadro FX 5500s with a 1 GB frame buffer per GPU and dual HD SDI graphics-to-video output boards, powering two dual-link DVI displays plus two HD SDI outputs.

The Quadro Plex Model IV is intended primarily for processing and visualization of large data sets desk side, in a workbench environment. It features two Quadro FX 5600 GPUs with a 1.5 GB frame buffer per GPU for processing of larger textures and can power up to four frame-synched dual-link DVI displays.

The Quadro Plex Model S4 is an industry standard form factor optimized for large-scale server deployments. Featuring four Quadro GPUs in a high-density 1U chassis for high-performance remote

graphics and offline rendering applications, the Quadro Plex Model S4 offers system monitoring, thermal control and fault notification for efficient integration into data centers.

NVIDIA Quadro Plex VCS Models I, II, III, and IV are currently available through NVIDIA and distribution partners worldwide, including PNY Technologies, Leadtek, and Elsa. The Quadro Plex VCS Model S4 will be available in the fourth quarter of 2007.

More Information at www.nvidia.com