

GPU Development with

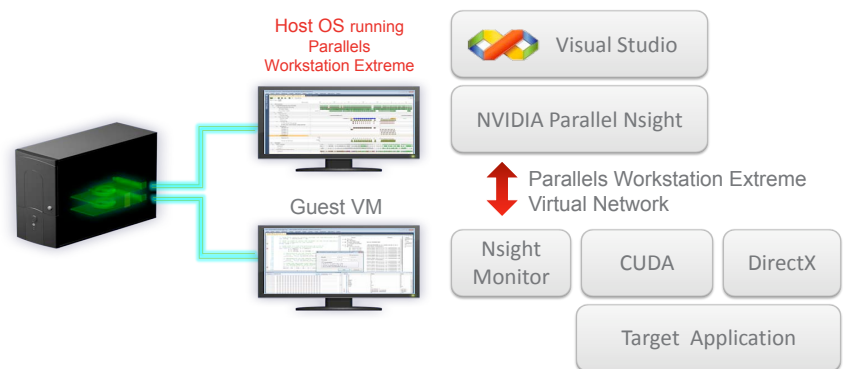
Parallels[®] Workstation *Extreme*

Parallels Workstation Extreme enables users to assign a GPU directly to the virtual machine, so they can run both host and target debug functions on a single workstation, without sacrificing performance or functionality.

The massive computing capabilities of modern graphics processing units (GPUs) are increasingly being tapped by developers of high performance application models in fields such as oil and gas, digital content creation, biosciences, manufacturing, and finance. This enables them to achieve significant performance gains over traditional programming models, with the applications they develop running as much as fifty times faster in some cases. As such, GPU development is expected to grow rapidly as more developers learn how to harness this power.

As with any kind of development, there are challenges; with GPU development one area of greatest expense and inefficiency is the need for both a “host” and a “target” machine. This is particularly relevant to visualization debugging, where it is necessary to set breakpoints, though doing so effectively halts the debugging process and a secondary system is required to reset the breakpoint and advance the process.

As a GPU drives the host operating system display, setting a breakpoint on that GPU can block the user from accessing the host user interface. Thus, a second workstation is needed to interact with the host operating system. Traditionally this has meant developers needed to run two separate workstations, each with its own GPU, connected via a network. Investing in double the amount of workstation hardware is costly, and switching between two machines is clearly time-consuming and inefficient.



Now, Parallels Workstation Extreme addresses this problem, enabling developers to run both the host and target machines together on one Intel Xeon-based workstation, each with its own dedicated GPU. Parallels Workstation Extreme is a powerful, next-generation virtualization platform that gives end-users dedicated graphics and networking resources for both host and guest workstation environments.

Using Parallels Workstation Extreme, developers can directly assign a secondary GPU directly to the virtual machine – designated as the target – without sacrificing native performance, even for graphically-intensive work. Through NVIDIA[®] SLI Multi-OS technology, the same Quadro graphics driver can be used on both the host operating system and the guest virtual machine, meaning any GPU-based application that uses OpenGL, NVIDIA CUDA, or OpenCL will run on either the native workstation or in the virtual environment without any changes.

Since Parallels Workstation Extreme allows users to assign a virtual machine its own dedicated GPU, developers can run the target virtual machine side-by-side with the native host machine on one workstation. As the virtual machine’s GPU will not be in “debug mode”, developers can use it to externally control and reset the host during breakpoints.

A network adapter can also be directly assigned to the virtual machine, ensuring high volumes of data can be transferred over the network without significant performance degradation. Yet the host and guest machines are completely isolated, so the development process is not compromised by any action performed in the other environment.

The cost and efficiency benefits of deploying Parallels Workstation Extreme during GPU development can be realized regardless of whether the user is developing applications for just one or for various platforms. Those running multiple

operating systems, such as Windows and Linux, will not need to multi-boot between the different environments, as they can be deployed and tested in a virtual machine with its own dedicated NVIDIA Quadro graphics card.

GPU Development for Linux on a Single Workstation


Parallels Workstation Extreme combines with NVIDIA's Linux GPU development tools, whether those tools are run on a Linux Host or in a Linux Guest VM. This enables GPU developers to function in multiple operating system environments at the same time to achieve their application development goals, thus enabling faster time-to-results.

GPU Development for Windows® on a Single Workstation

Parallels Workstation Extreme also combines with NVIDIA Parallel Nsight and Microsoft® Visual Studio® to create a comprehensive environment for Windows-based GPU development. NVIDIA Parallel Nsight is a tool offered by the GPU manufacturer to extend the debugging and performance analysis capabilities of Microsoft Visual Studio. NVIDIA's solution assists with difficult tasks such as debugging, performance tuning and designing parallel algorithms.

System Requirements for Parallels Workstation Extreme:

- Intel Xeon Processor 5500 or 5600 Series workstation platform with the Intel X58 or 5520 chipsets
- Two NVIDIA Quadro FX 1800, FX 3800, FX 4800, FX 5800 graphics cards; or two Quadro 2000, 4000, 5000, 6000 graphics cards
- Host Operating System
 - Windows 7 64-bit, Windows Vista 64-bit, Windows XP 64-bit
 - Red Hat Enterprise Linux 5.3, 5.4, 5.5 64-bit



Contact Us

For more information about using Parallels Workstation Extreme to develop or deploy GPU-based applications, please visit www.parallels.com/extreme

Parallels offers volume licensing to businesses and educational institutions. Please contact your local Parallels sales representative for more information or contact us directly at:

Phone: +1 425-282-6400

Fax: +1 425-282-6444

www.parallels.com