SOLUTION PAPER

HIGH-PERFORMANCE COMPUTING WITH LEAN CLIENTS

One bottleneck less: with the help of virtualized, shared graphics resources (vGPU), even workstations in architecture, engineering, planning and design requiring particularly high levels of computing power can be provisioned efficiently and securely via thin clients.

Contents:

The trend towards IT standardization continues uninterrupted. More and more applications and desktops are being virtualized and provisioned centrally or hosted in the cloud. Instead of PCs, users are increasingly relying on energy and cost-efficient thin or zero clients. After office and standard workstations, another traditional PC domain is now crumbling – workstations requiring high levels of graphics and video performance. This is possible thanks to shared graphics resources in the data center. The technology is still relatively expensive, so why are companies and government agencies so interested in it?

Great need for data protection and IT efficiency

Workstations that require a high level of graphics performance usually handle sensitive or security-relevant data. This is particularly true in product design or the engineering of machines, components and tools. However, typical uses also include architecture, building planning (including cabling and pipework) and film editing. To ensure efficient teamwork, unified communication solutions should also be integrated.

a potential target for industrial espionage attacks. However, the need for IT standardization is growing at the same time. As a logical consequence of both of these facts, more and more companies are opting to host their IT systems in a suitably secured data center – either within their own organization or somewhere in their own country (or in a country with strict legal data protection regulations).

Consolidation and cost transparency

As the level of consolidation in the data center increases, so do the cost benefits. The more applications are provided centrally or as a cloud service, the simpler IT management becomes. Phasing out remaining island solutions with fat clients by replacing them with thin or zero client workstations therefore makes sense because the client-side management, service and support outlay can be reduced on a permanent basis. As part of Client as a Service (ClaaS) contracts, many service providers offer even graphics intensive applications such as CAD or GIS including thin client leasing and management as a cost-transparent rental package per workstation.

GRID VGPU TECHNOLOGY FROM NVIDIA

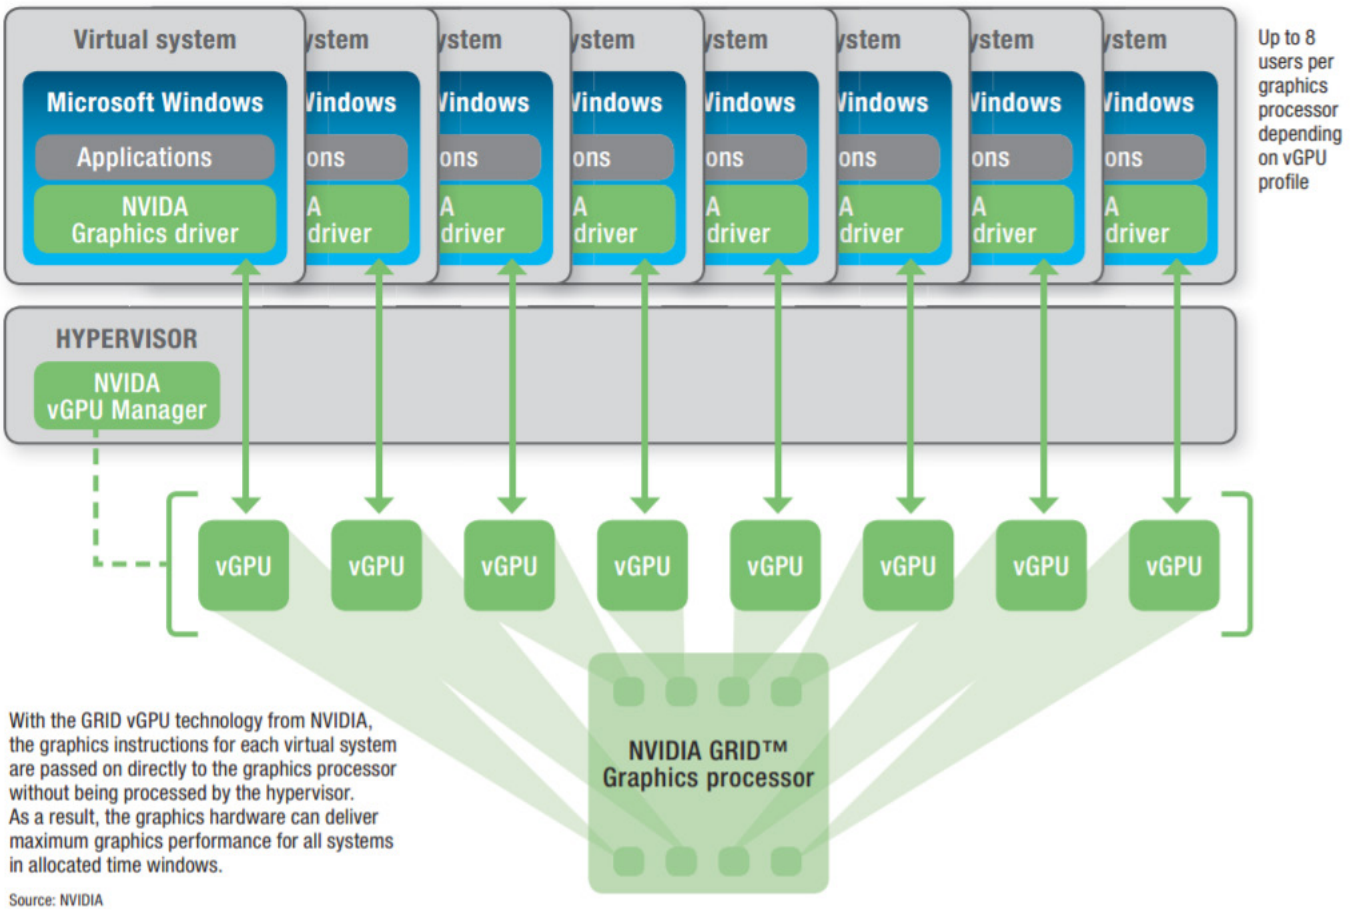

TABLE 1: MANUFACTURER-DEPENDED ROUTES FOR HARDWARE-BASED GRAPHICS ACCELERATION IN VARIOUS ECOSYSTEMS IN COMBINATION WITH NVIDIA GRID GRAPHICS CARDS

Ecosystem	Shared GPU (1:n)	Dedicated GPU for virtualized desktops (1:1)	Virtual GPU for virtualized desktops (1:n)	Supported versions and hypervisors (as at: 03/2015)
Citrix XenApp	X (GPU Pass Through)	–	–	From Citrix XenApp 6.5 with OpenGL 4.3 or XenApp 7.5 + NVIDIA GRID (GPU Pass Through), in conjunction with XenServer or VMware ESXi server as hypervisor
Citrix XenDesktop	X	X	X	Dedicated with Citrix XenDesktop 5.6 or 7, with XenServer or VMware ESXi server as hypervisor vGPU from XenDesktop 7, only in conjunction with XenServer
VMware Horizon	X (vSGA)	X (vDGA)	– (vGPU)	Dedicated from Horizon View 5.3 with vDGA, only with VMware hypervisor (incl. VMware vSphere) vGPU from Horizon 6.0 + NVIDIA GRID (vSGA)

Graphics sharing – a key technology

In the past, the biggest hurdles when it comes to the central or cloud-based provision of graphics-intensive applications were latencies in transferring content and the lack of a facility to virtualize and divide up physical graphics resources in the data center. High demands could only be met via 1:1 connections. Low-performance, inefficient graphics card emulation by the server CPUs (soft GPU) was not an option for 3D applications, for example. With hardware-based graphics acceleration by split graphics processors (vGPU), there is now another way of allocating graphics performance to users. Having been further developed accordingly, the Citrix, VMware and Microsoft IT environments allow virtual applications or desktops to take advantage of a discreet graphics memory and the high performance standards of DirectX and OpenGL (Fig. 1)

Up to 32 users per graphics card

With the graphics cards from the technological leader NVIDIA, up to eight users can be allocated to each graphics processor. Because an NVIDIA GRID K1 card, for example, has four physical GPUs, up to 32 users can access it together. When it comes to allocation, a distinction is made between three types of hardware-based graphics acceleration which are not available in every ecosystem (Table 1). With GPU sharing, the GPU is made available to a number of virtualized applications (Citrix XenApp) or desktops (VMware vSGA1). For virtual desktops with particularly demanding requirements (e.g., for 3D), a 1:1 relationship is an option (dedicated GPU for Citrix XenDesktop or VMware vDGA2). 1:n allocation to virtual desktops via virtualized GPUs (vGPU) is much more flexible. In combination with suitable hypervisors, NVIDIA with vGPU supports both Citrix XenDesktop and VMware Horizon.

Requirements as regards the thin client

To ensure efficient transfer to the thin client, the leading software providers Citrix, VMware and Microsoft have each developed new high-performance protocols (Table 2). These are HDX 3D Pro for Citrix XenApp and XenDesktop, the "Pixel Perfect" protocol PCoIP in combination with the Build-to-Lossless function for VMware Horizon and RemoteFX for the Remote Desktop Services for Microsoft Windows Server (from Version 2008 R2). In order to meet the relevant requirements, the thin client manufacturers, too, have further developed their solutions and are taking part in certification programs such as Citrix HDX Ready or HDX Ready Premium. Thanks to the high-performance thin client hardware, a 2D workstation can now be operated at low cost via a standard thin client with a dual core processor, for example an IGEL UD5. For professional 3D applications, there is particularly high-performance hardware required such as the IGEL UD6 with quad core CPU and 2GB RAM.

TABLE 2: HIGH-PERFORMANCE PROTOCOLS FOR CITRIX, VMWARE, MICROSOFT ENVIRONMENTS & REQUIREMENTS AS REGARDS THE CLIENTS

Virtualization solution	Protocol	Properties / approach	Requirements as regards client
Citrix XenApp / XenDesktop	HDX 3D Pro	Includes various technologies for graphics acceleration in order to optimize the virtualization of rich graphics applications	Citrix Receiver 13 Recommended hardware: 2 GB RAM (min. 1 GB) and 2 GHz quad core CPU (at least 2 GHz dual core)
VMware Horizon	PCoIP with Build-to-Lossless function	So-called "Pixel Perfect" protocol with the option of quality adaptation for high, medium and low bandwidths	VMware Horizon Client from version 3.0
Microsoft RDS	RemoteFX	Extension to the Remote Desktop Protocol (RDP) for improving graphics performance	RDP client including RemoteFX 8 for Microsoft Windows Server (from 2008 R2)

Compensating for extra costs with thin clients

Naturally, the special server cards from NVIDIA, AMD or Intel allowing hardware-based graphics acceleration through GPU virtualization increase costs in the data center. The more users share a card, the more cost-effective the solution is. However, absolute high-end users may require a card and two servers in the backend just for themselves. The extra costs in the computer center compared to a local PC solution can be at least partly compensated for by the savings achieved through the expanded thin client infrastructure. If island solutions with a PC are replaced by a combination of a thin client and a private cloud, the overall workstation IT costs can on average be reduced even further[3]. Companies and government agencies can achieve a low-cost migration with the help of software thin clients[4] which, as a temporary solution, run very efficiently on the high-performance PC hardware. Greater work mobility results in further cost benefits. With a thin client concept, users are no longer tied to a specific place in the office (desk sharing) subject to the availability of suitable hardware. Provided that there is adequate bandwidth, thin clients are also an option for secure home workstations in the field of development and design. Last but not least, thin clients improve data protection, an issue that is of considerable importance to many companies.

7 REASONS FOR HIGH-END GRAPHICS WITH THIN CLIENTS

1 **High level of data protection**: central data, secure connections (VPN), central USB control, secure authentication (smartcard, badge, USB token etc.)

2 **Cost savings through workstation standardization**: robust, long-lasting hardware with unified managment[5]

3 **Rollout, service and support**: reduced outlay thanks to profile-based configuration (displays, dualview, etc.)

4 **Compliance**: highly-available management if required (UMS HA option)

5 **"Gentle" migration**: existing PCs continue to run as software thin clients (IGEL UDC2)

6 **Option of IT outsourcing**: reduction in balance sheet and cost transparency thanks to Client as a Service (ClaaS)

7 **Mobility**: desk sharing and secure home office for development / design (if adequate bandwidth is available)

Criteria when selecting a provider

How much extra efficiency and security is gained ultimately depends on the choice of thin client provider and the solutions offered. The software is often what makes the key difference. For example, all IGEL Technology models come supplied with the IGEL Universal Management Suite (UMS). Compatible with all operating systems[6] and setups, it allows comprehensive remote management for anything from the simple IGEL zero client to the UDC2 software thin client and even the high-performance multiprotocol UD6 thin client. With its profile-based structure, the UMS is particularly impressive in conjunction with dualview workstations. IGEL also includes numerous functions to optimize data protection, for example, a certificatebased, encrypted remote management system, support for various VPN and authentication solutions or advanced USB control.

Verdict

With GPU virtualization and high-performance thin clients, graphicsintensive workstations can be provided in a secure, efficient manner in order to encourage workstation standardization, mobility or IT outsourcing as necessary. Given that the level of investment required is still high, it is important to identify exactly what additional savings and security features the thin client solution offers.

1 Virtual shared graphics acceleration
2 Virtual dedicated graphics acceleration
3 Thin clients 2011: Ecological and economic aspects of virtual desktops, Fraunhofer Institute for Environmental, Safety and Energy Technology UMSICHT: www.igel.com/uploads/media/thinclients2011-de-2.pdf
4 Conversion using IGEL Universal Desktop Converter 2 (UDC2)
5 IGEL Universal Management Suite (UMS) supplied; universal OS compatibility for hardware with IGEL Linux and Windows Embedded)
6 Windows Embedded and IGEL Linux

Visit us online at **igel.com**

IGEL®
ENDPOINT MANAGEMENT

Revolutionary in its
Simplicity