GPU IN THE CLOUD: BRING YOUR COMPANY TO THE NEXT LEVEL

OUTSCALE

cisco Powered

Excellence as a Service

Introduction

Graphics virtualization will bring virtual desktop infrastructure into industries that rely on graphics performance to get work done. Virtual desktops have instigated a revolution in the way IT departments think about enterprise desktop deployments. They've resulted in more efficient software rollouts, more secure end-user systems, and have increased the productivity of enterprise IT departments by orders of magnitude. Nvidia, the leading provider of GPU hardware and software frameworks for graphics virtualization, expects that the market for GPU virtualization in the Cloud will reach $1 Billion in the next two years.

How can you use GPU in the Cloud to set virtual desktop infrastructure and much more?

Graphics Virtualization Takes Virtual Desktops to the Next Level

Over the last few years, virtual desktops have become increasingly important to business. As businesses move away from in-house IT infrastructure management and towards a Cloud model that outsources infrastructure to Cloud vendors, it makes sense to move away from the old enterprise desktop strategy and the consistent software and security headaches it causes.

Virtual desktop infrastructure (VDI) allows businesses to centrally manage their enterprise desktop portfolio, reduce costs, rationalize the use of IT equipment, and empower employees to work flexibly, whether within an office setting, out in the field, or remotely from their homes.

But until relatively recently, virtual desktop platforms had one big limitation: they were solely CPU-based. Virtualization technology hadn't advanced to the level at which it could reliably virtualize the GPU-processing that many modern desktop applications depend on. Virtual desktops were great for running groupware, collaboration, and productivity software, but not so good at running any sort of graphically intensive application, such as video and photo editing, 3D modeling, or design.

Graphics chips are designed differently to CPUs. They are optimized for parallel processing and graphics-intensive workloads. However powerful a virtual CPU is, it isn't as good as a dedicated GPU.

That left businesses with a requirement for graphical applications out in the cold where virtual desktops are concerned. Fortunately, over the last couple of years, that has changed. With the advent of Cloud virtualization technologies, forward-looking Cloud platforms have evolved to the point at which they are capable of supporting the most intensive graphical applications within virtual desktops.

According to Ovum analyst Roy Illsley, "The use of rich 3D graphics on virtual desktops was a barrier to some organizations as the user experience was less than satisfactory. NVIDIA's GRID vGPU removes this barrier and effectively enables IT managers to virtualize a whole new class of users."

Virtual desktops aren't the only advantage that Cloud platforms with graphics virtualization have. Because they allow applications to offload graphics processing to the Cloud, application developers don't have to worry about the limitations of the devices on which users access software. In a world where mobile devices dominate the workplace, that's important.

However impressive mobile devices and their embedded Systems-on-a-Chip become, they are likely to lag behind the graphics processing power of server and desktop GPUs, not because power conservation is the priority on mobile.

GPU virtualization allows companies to move the bulk of the graphics processing, or indeed any processing that can take advantage of GPU's parallel processing power, into the Cloud, which massively increases the potential of smartphones and tablets. We can imagine a situation in which an architect makes changes to plans on a tablet while out in the field and is able to leverage Cloud GPU processing to have 3D renders of those changes delivered to her device in real time.

Virtual desktop infrastructure has been popular in enterprise for a few years, but with the increasing utility of Cloud platforms with graphical virtualization capabilities, we can expect to see their penetration expand into industries that have previously been unable to rely on traditional virtual desktops because of graphical limitations.

GPU Virtualization Is Changing The Face Of Virtual Desktops

The migration to virtualized, centrally managed, available-everywhere virtual desktops has changed the traditional office into a more flexible, mobile and modern way to work.

However, until fairly recently, virtual desktops have been somewhat limited. They are great for simple office tasks, for collaboration, and for BYOD and on-the-road access to files and software. But they have also been seriously constrained in their graphical capability. Running intensive video, CAD, or image editing software on virtual desktops has been largely impossible because CPU virtualization wasn't up to the task.

The introduction of powerful GPU virtualization technology, like what's provided by Nvidia's GRID system, has changed all that. GPU virtualization allows virtual desktop users to access the full capabilities of graphics processing units on their virtual desktops, opening up a huge range of new possibilities that has the potential to transform Cloud desktop solutions from simple providers of centralized productivity and groupware application deployment platforms to full-powered workstation solutions.

Software that depends on the GPU is an essential part of the modern workplace. It's not limited to power users and graphical artists. Most of the knowledgable workers interact with graphing software, image editors, video conferencing and rich media internet services as part of their daily routine. Virtual desktops have traditionally delivered an unsatisfactory experience to this segment of the workforce. Even something as simple as streaming video to a virtual desktop was frustrating.

Those limitations have excluded a significant number of workers and companies from the virtual desktop revolution. With virtualized graphics processing, users can now have access to the same rich media capabilities on their virtual desktop as they experience on any laptop or desktop PC, making virtual desktops a viable solution for companies that had previously considered virtual desktops, but dismissed them because of a lack of graphics processing power.

Of course, the power of GPU virtualization is not limited to knowledgable workers. Any software, even the most graphically intensive can be deployed on GPU-equipped virtual desktops. Outscale client Dassault Systemes recently showcased how they are using Outscale's GPU-accelerated

Cloud platform to revolutionize CAD modeling with real time realistic rendering using CATIA 3D.

GPU virtualization brings the power, flexibility, efficiency, and cost savings of virtual desktops to media companies, product designers, publishers, photographers, video editors and more.

Isn't Just About Graphics Performance

As seen previously, Nvidia is the leading of GPU hardward. That's no surprise to anyone who keeps a close eye on the Cloud market. Cloud gaming has grown massively over the last few years as businesses seek to leverage the cost and performance benefits of Cloud infrastructure to provide gamers with a Netflix like experience. Predictably, much of the press around GPU virtualization has focused on gaming, but gaming is far from the only area in which we can expect to see huge growth in Cloud use powered by graphics virtualization.

As discussed in this eBook previously, desktop virtualization solutions have been limited in their capability by a lack of GPU virtualization platforms. Running relatively simple groupware and collaboration applications is reasonable without GPU virtualization, but for the video and image processing applications that creatives require, desktop virtualization was not an option. That's changing, bringing the benefits of desktop virtualization to a market segment that was previously limited to in-house cluster and workstation deployments.

However, not all uses of GPU virtualization are concerned with graphics. There are many areas of computation that can leverage the specific hardware capabilities of GPUs to provide processing speeds that are orders of magnitude faster than equivalent CPU-based processing.

Because of the particular requirements of gaming, GPU development followed a different path than that of CPUs. GPU development has focused on incorporating massive numbers of cores onto chips, and that's not just good for gaming, it's good for any workload that requires parallel processing.

I'd like to look at three examples of industries leveraging the power of GPU-based parallel processing to increase the speed and efficiency of computational workloads.

Bioinformatics

Bioinformatics is a rapidly growing field that applies computational modeling and analytics to biological problems that include protein sequencing, drug development, and evolutionary biology. The biotech startup space is booming, and it depends on the ability to process massive amounts of data.

As we've mentioned before, protein analysis, which includes DNA analysis, requires the sequencing of extremely large data sets. Fortunately, sequencing is an excellent candidate for parallel processing, which means sequencing workloads are orders of magnitude more efficient when run on graphics processing units.

"Sequence comparison is the lifeblood of Bioinformatics. Weiguo Liu et al. ran the key Smith-Waterman algorithm on a high end GPU. They demonstrated a reduction by a factor of up to sixteen in the lookup times for most proteins. Smith Waterman has also been ported to the Sony PlayStation 3 and the GeForce 8800 (CUDA). Trapnell and Schatz also used CUDA to port another sequence searching tool (MUMmer) to another G80 GPU and obtained speed ups of up to 13 when matching short DNA strands against much longer sequences."

Big Data Analytics

Data analytics offers an essential competitive advantage to businesses, and it is intimately related to Cloud computing. Processing and storing huge amounts of ephemeral data is a perfect application for highly elastic Cloud platforms.

As with bioinformatics, big data analytics is an area that can benefit enormously from parallel processing on GPUs. Although big data has been hailed as transformative to business, its potential has always been limited by the time and expense involved in its processing. Leveraging Cloud with graphics virtualization can make Big Data cheaper and faster, and therefore a practical option for business.

In a paper on the use of GPUs to accelerate text mining, a common Big Data task, researchers from North Carolina State University discovered that:

“We achieve up to a 30-fold speedup over CPU-based algorithms for selected phases of the problem solution on GPUs with overall wall-clock speedups ranging from six-fold to eight-fold depending on algorithmic parameters. This experiment demonstrates the potential of GPUs to accelerate even integer oriented, branch-dominated massive data text mining algorithms.”

Computational Finance

Modern finance has seen an explosion in the amount of data available to it, along with an increase in the cost of infrastructure required for its processing. Financial quants have specific requirements of data processing: it has to be able to handle massive amounts of data and it has to do it very quickly.

According to a series of case studies delivered by Hanweck Associates, a provider of data analytics to exchanges, banks, and brokers, GPU processing of financial computation workloads has the potential to both reduce the cost and increase the speed of financial modeling, risk assessment and analytics.

“GPUs are the way forward. Major financial institutions are using them for quant finance. Performance gains of more than 10x times, dollar for dollar, are achievable in practice in many common use cases, which is generally sufficient to offset the costs of new development.”

GPU and The Cloud

Until recently, the bulk of parallel processing with GPUs in industry has leveraged private infrastructure. With the advent of powerful graphics virtualization Cloud platforms, businesses can combine the benefits of GPU processing with the established benefits of Cloud infrastructure.

TO GO A STEP FURTHER

CUDA In The Cloud Makes GPU Computing A Cinch

Traditionally, the industry hasn't cared much about graphics processing in enterprise data centers with the exception of a limited set of applications like render farms. Servers are usually headless machines and they haven't often required graphics accelerators to do their work.

The need for graphics acceleration in enterprise servers changed with the emergence of general purpose GPU computing. While the majority of enterprise servers still don't make much use of graphics processing, there are a large number of application where GPU processing has proven incredibly powerful and useful. GPUs evolved to be ideally suited for parallel processing workloads , much more so than CPUs, which with their limited number of cores when compared to GPUs, can be much slower for parallel workloads.

But using GPUs for general purpose computing was ferociously complex in the early days. It required extensive knowledge of graphics programming, which was thin on the ground in enterprise computing. Most enterprise developers did not have the experience or knowledge to leverage graphics technologies like Direct3D.

NVIDIA's introduction of CUDA changed all that. CUDA is a parallel computing platform and API that lets developers use languages they are familiar with — C, C++, and Fortran — to build general purpose applications for graphics processing units.

It's hard to overestimate the impact that general purpose graphics computing is having on enterprise IT. As David Cardinal recently wrote on ExtremeTech:

"From spending a few days at last year's Nvidia Global Tech Conference (GTC) it is very clear that the uses for GPUs have exploded -- they have become an essential element in dozens of computing domains. As one attendee suggested to me, GPUs could now be better described as application coprocessors."

Until fairly recently, the Cloud had the same approach to general purpose computing with GPUs as traditional data centers. Which is to say they had very little to do with them. Virtual machines were all about the CPU and graphics virtualization was not well developed. In that last couple of years, largely due to work by NVIDIA, graphics processing on the Cloud has become a reality — and that includes CUDA.

Just as general purpose GPU computing instigated a sea-change across dozens of areas of computing, so Cloud GPU virtualization ushers in a new way of deploying and managing powerful GPU-computing platforms.

We are all familiar with the benefits of the Cloud, but they're worth re-iterating because GPU virtualization is game changing. No longer do companies have to buy and manage expensive infrastructure to reap the benefits of GPU computing. They can deploy GPU-equipped servers to run their CUDA applications as easily as they have been able to spin up other servers in the Cloud.

That makes it cheaper and faster to built GPU-computing applications, which is likely to drive a greater expansion of GPU use in the enterprise. Big data analytics, finance, scientific computing, and many other areas could benefit from GPU virtualization. The space will become even more interesting as disruptive startups leverage GPU virtualization for their workloads.

Conclusion

GPU virtualization brings the power, flexibility, efficiency and cost savings of virtual desktops to media companies, product designers, publishers, photographers, video editors, and more.

As with Bioinformatics and Computational Finance, Big Data analytics is an area that can benefit enormously from parallel processing on GPUs.

The new era of GPU in the Cloud is here!