

Teradici APEX™ 2800: Addressing the next VDI bottleneck

By: Mike Laverick

Table of Contents

Executive Summary	3
A Brief History of VDI Roadblocks	4
Legacy Methods Will Not Scale	8
Teradici APEX 2800 Business Benefits	9
Conclusion	13

Executive Summary

Virtual desktop infrastructure (VDI) is gaining momentum in the marketplace, and market data suggests sizable growth over the next few years. According to IDC research, VDI will grow from \$2.3 billion in 2011 to more than \$3 billion by 2015.¹ It's clear that VDI will continue to grow and mature as a viable technology. Customer research by Virsto.com indicates that more than 60% of those polled will deploy VDI within the next 12 months.²

This growth is being driven in part by the deployment of VDI in new segments of the organization, namely for office workers and power users. However, these new VDI users represent a more demanding set of users compared with the task worker, creating a new set of challenges for IT managers deploying the solution:

- Average workloads have more multimedia, impacting negatively on server density.
- Peak workloads become more difficult to assess, forcing IT managers to put aside an increasing amount of CPU buffer for “just in case” situations.

The Teradici APEX 2800 server offload card addresses these two pain points by unburdening the PCoIP encoding tasks from the server CPU of the most active 64 displays on the server. By constantly monitoring the activities of all the virtual machine (VM) displays on the server, and dynamically and instantaneously offloading the most active 64 of them (for displays with resolution up to 1920 x 1200), the APEX 2800 is able to:

- Reduce the impact of multimedia workloads on the server CPU, improving its density and decreasing the cost per VM.

1 <http://www.slideshare.net/btdrew/vdi-adoption-indications-2012>

2 <http://virsto.com/images/site/page-content/infographic960.png>

- Lower the amount of CPU buffer put aside for “just in case” peaks outside of the initial base worst-case scenario. Businesses can regard adoption of the Teradici APEX card as an insurance policy, since it takes the risk out of deploying VDI to a new, more demanding user base.

A Brief History of VDI Roadblocks

Since the idea of virtual desktops was first postulated, it seems like not a year has gone by without the industry heralding the “Year of the Desktop.” The reality, however, is that virtual desktop projects have often stalled after successful proofs of concept. This is often due to businesses scaling up their deployment and running up against bottlenecks that inhibit their ability to maintain a cost-effective VDI implementation. The challenges most often come when attempting to add more users while meeting increasing demands.

The **main bottlenecks** for full-scale VDI deployments have so far been around storage (IOPS) and memory. Memory was a major bottleneck as organizations tried to leverage their virtual desktop commitment as part of their strategy for migrating users from Windows XP to Windows 7. As part of a future-proofing process, these organizations often decided to make the switch from 32-bit computing, which was first introduced with Windows 95, to a true 64-bit operating system. However, many were somewhat taken aback by the increased memory requirements that 64-bit operating systems introduced. Fortunately, the amount of physical RAM that a hypervisor like VMware ESX can address has massively and simultaneously increased. In addition, the amount of physical RAM customers can afford has increased, while new servers and hardware physically support more and more memory slots. For instance, the new servers shipping with the Intel Sandy Bridge chipset can now address up to 1TB of memory. Enhancements at the virtualization layer coupled with technologies such as VMware’s transparent page sharing and memory compression features, the memory bottleneck is slowly disappearing, as most new servers can now easily support up to 2GB per user, at full density.

Similarly, storage used to be one of the principle bottlenecks, both in terms of the cost of storage capacity and the IOPS generated by activities such as boot storms and login and logout. Before the introduction of VMware “linked clones” or other storage vendor-assisted clones, storage was one of the biggest cost implications of VDI. The sheer amount of disk space required to hold the virtual desktop images was daunting. However, with the use of thin-provisioned virtual disks, efficient cloning technologies have almost eliminated this issue. There does remain the issue of disk IOPS caused by events such as logoff and logout, but most customers have successfully addressed this issue by removing the disk spindle out of the equation altogether, using technologies such as array-based flash-based caches or through tactical deployment of solid-state drives (SSDs). Recent research from Virsto.com shows that most proofs of concept fail because of unplanned costs (31%) and poor end-user performance (29%), whereas the storage bottleneck has been relegated a relatively minor concern (19%).³

As ever in the world of IT, solving one bottleneck only helps identify the next bottleneck in the stack. The long-term trends point to a new type of user that expects to have access to social networks and websites like YouTube.com during office hours. Studies in 2008 by Nielsen Online⁴ showed that 65% of online users watch videos between the hours of 9 a.m. and 5 p.m. Similar Nielsen, Burst Media and eMarketer⁵ studies reveal that the average employee spends one to two hours each day using the Internet for personal reasons.

Businesses that allow end users to access these sites only outside official office hours should expect to see higher activity during the early morning and lunch breaks, pointing to peaks in CPU activity for video encoding.

3 <http://virsto.com/images/site/page-content/infographic960.png>

4 http://www.readwriteweb.com/archives/most_online_video_is_watched_at_work.php

5 <http://www.wisegeek.com/how-do-employers-monitor-internet-usage-at-work.htm>

It is becoming apparent that the CPU will increasingly be regarded as the limiting factor in increasing or even maintaining an acceptable server density while meeting the requirements of these new types of VDI users. All VDI projects attempt to lower the cost per VM as much as possible, but they face increasing challenges to meet that goal. The fact that CPU capacity has been recently upgraded with the introduction of the Intel Sandy Bridge chip only highlights further the need to use the CPU for running applications versus PCoIP encoding tasks. VDI represents significantly different workload-to-server consolidation – and this fact is often overlooked by virtualization experts with little or no background in supporting desktop environments.

The default protocol in a VMware View environment is the PCoIP protocol. This protocol works by using the server CPU coupled with a sophisticated set of codecs to render pixels and graphics at the server side and remotely display them across the network at the user's screen. The PCoIP protocol uses a sophisticated multi-codec compression protocol that delivers the best image quality for the available network bandwidth. The more graphically intense the user's PCoIP session, the more workload is placed on the server CPUs for the rendering and protocol encoding process. The Teradici APEX card does not render pixels – that's still the job of the software layer – but rather offloads the PCoIP protocol encoding tasks from the main server CPU, freeing up valuable CPU cycles for the applications. It also means that the virtual desktops themselves could be configured with less vCPUs, translating to less work for the ESX "scheduler" in scheduling those vCPUs down to the physical CPU. Additionally, it could result in extending the shelf life of older servers with a lower clock speed.

In the past, VDI projects concentrated on the low-hanging fruit represented by task workers. These users' desktops were the easiest to virtualize because often the users were focused on a small number of applications that weren't especially graphically intensive. With the relative success of these VDI projects, both business managers and other end users are demanding virtual desktops over desktop PCs. The problem is that these new "structured task" and power user varieties often have

higher performance and experience expectations than traditional users, putting an increased workload on the on-board CPUs in charge of delivering the users' environment.

The world of VDI is not standing still. Virtualization vendors, such as VMware with its View product, are seeking to add native support for the advanced features of GPUs into their hypervisors such as VMware's ESX, with intelligence up to the broker and virtual desktop to utilize them. The net result is that the rate of pixel processing is not going to decrease or stay static – it will only increase over the next couple of years as a result of this improved user experience. So what may scale well today may not in the future, and businesses that wish to take advantage of the new GPU functionality might find themselves inhibited by the bottleneck of the CPU on the server motherboard. In this respect, the Teradici APEX is complementary to these new GPU developments – it does not seek to compete against the GPU approach but rather enables customers to more rapidly adopt the new GPU capabilities.

The new ability for virtual desktops to leverage GPUs for improved user experience is not free and actually consumes a lot of CPU cycles to encode the additional pixels being generated. In some cases, GPU-enabled virtual desktops can have CPU utilization that is more than 15% higher, thus decreasing virtual desktop ratios by as much as 20%. For example, turning on 3D/Aero on Windows 7 VMs on the server, the APEX 2800 will reduce the impact on server density by 75% compared to not having the card. Additionally, it cannot be ignored that with improved graphics support from both the virtualization vendors and the guest operating systems, customers and end users will eventually expect to enable graphically intensive features such as 3-D and Aero Glass support in Windows 7 and the Aero Glass experience is the only choice available in Windows 8. The Teradici APEX card empowers the VDI administrator to improve the user experience without negatively impacting perceived performance and/or server density.

Legacy Methods Will Not Scale

The typical approach to issues of scalability has been to throw hardware at the problem. That has meant either buying more powerful servers to absorb the increased CPU load and increased requirement for a “just in case” CPU buffer or, alternatively, scaling out the solution by buying more servers. Neither approach is very slick or serviceable, given that most organizations would prefer not to embark on an expensive and time-consuming upgrade of factory-fitted CPUs. Nor would most consider buying more servers a cost-effective solution, as this would undo much of the good work in the past decade to reduce the physical server footprint via server virtualization.

Moreover, this approach comes with a double burden. Organizations may be wasting money by spending on scaled-up or scaled-out server infrastructures as an insurance policy against excessive CPU spikes. Indeed, it could be regarded as a classic case of over engineering, where IT overprovisions resources at the server back end out of fear of being overwhelmed with resource requests only to find those request don’t surface in the short to medium term. This leaves corporate IT in the unwelcomed situation of trying to explain spending on resources that cannot be justified by the amount of actual workload. In short, financial resources could be spent on an insurance policy rather than buying more servers. The reality is that what constitutes peak CPU load is increasingly harder to predict when all the new changes are placed end to end. The traditional methods of dealing with these uncertainties necessitates an unwanted spend on additional server resources to absorb the potential workload that may be infrequently called upon. Unfortunately, the new generation of power users has a low tolerance of even the slightest delay, which can be attributed to both their high expectations generally and the perception that virtual desktops should outperform or at least match the experience delivered by a conventional PC. End users don’t “see” the wider business advantages of VDI. What they “see” is their screen and their application, and if that is perceived to be suboptimal compared with their PC, they will invariably complain.

Teradici APEX 2800 Business Benefits

The Teradici APEX 2800 is a simple-to-deploy PCIe card that is fitted to the VMware ESX host. It requires a small, lightweight driver that is installed using the ESXCLI command leveraging the VMware Install Bundle format. The driver can be incorporated into the ESX image media and slipstreamed into any existing deployment process. At the virtual desktop side, a small driver can be installed into the parent VM or template and thus becomes part of the standard build for any virtual desktop pool. The whole process can be carried out without impacting end users by using VMware's maintenance mode feature automated by vMotion. Once installed, the APEX card immediately starts the process of offloading the PCoIP protocol encoding. Scalability is delivered by installing the card into all ESX servers in a dedicated VDI cluster. The system supports not just one offload card, but two – all PCoIP sessions are offloaded by default, and the system will actively offload the most active 64 PCoIP displays if the ESX server has more. Support for up to 128 offloaded displays per ESX server can be added by merely incorporating a second Teradici APEX card into the server. Although the card is local to the ESX host, it does not stop other popular features of vSphere from functioning. So features that VMware admins know and love – such as vMotion, DRS and HA – function with complete compatibility.

Teradici APEX 2800 At A Glance

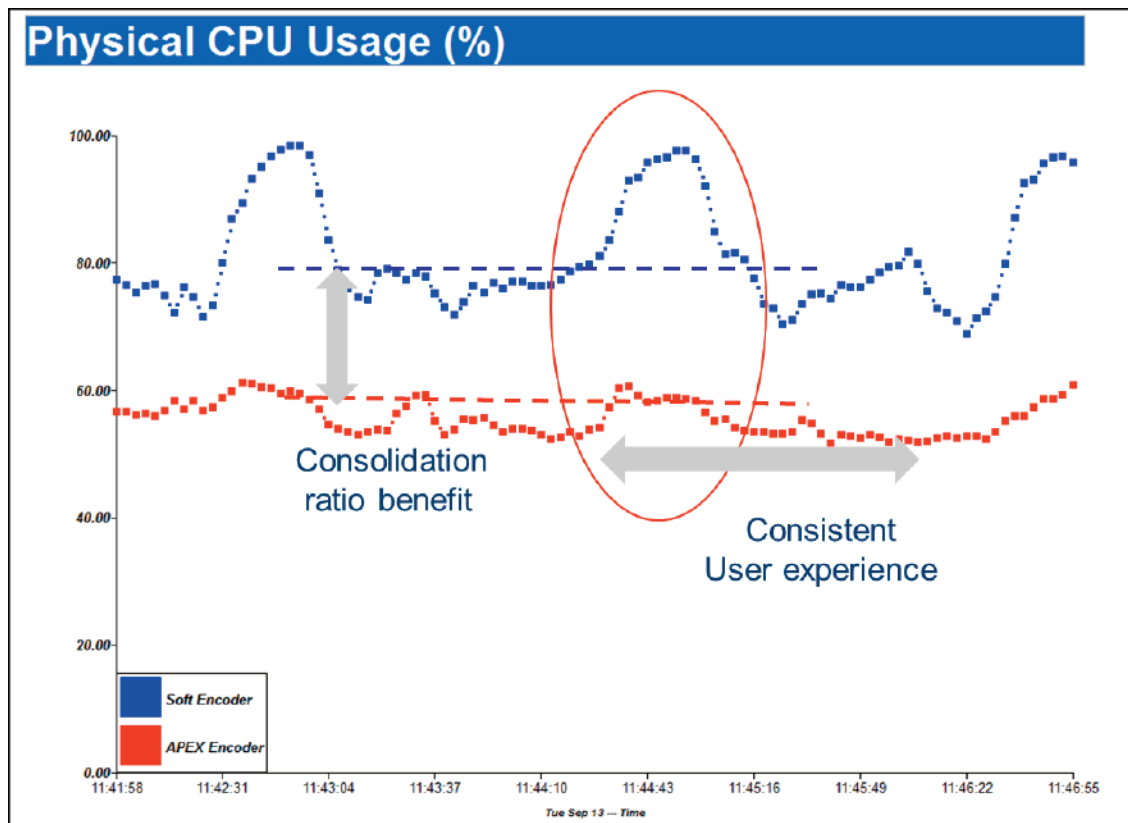
The Teradici APEX 2800 server offload card ensures the success of VMware View deployments by offloading PCoIP image encoding tasks and reducing server CPU utilization. This empowers IT managers to protect and ensure a consistent user experience, while enabling increased VDI consolidation ratios. The APEX 2800 is a performance enhancement to VMware View deployments that's simple to install and set up.

The APEX 2800 solution:

- Protects and ensures a consistent user experience
- Enables VDI consolidation ratio improvements
- Allows for simple setup and installation
- Requires no special management tools or software to function
- Is integrated into VMware View, allowing for specific virtual desktop pools

Close integration with the VMware View product allows the administrator granular control over APEX functionality. It makes it easy to focus the full power of the card toward the users that will benefit the most, namely the demanding power users, who represent the most valuable constituents to a business that is using virtual desktops. Additionally, command-line tools are installed into the ESX host to enable advanced troubleshooting and diagnostic work to be carried out remotely if need be.

Once installed, the card immediately begins the process of offloading the PCoIP encoding tasks from the CPU without any additional tweaking. It does this by picking up the default settings enabled in VMware View.

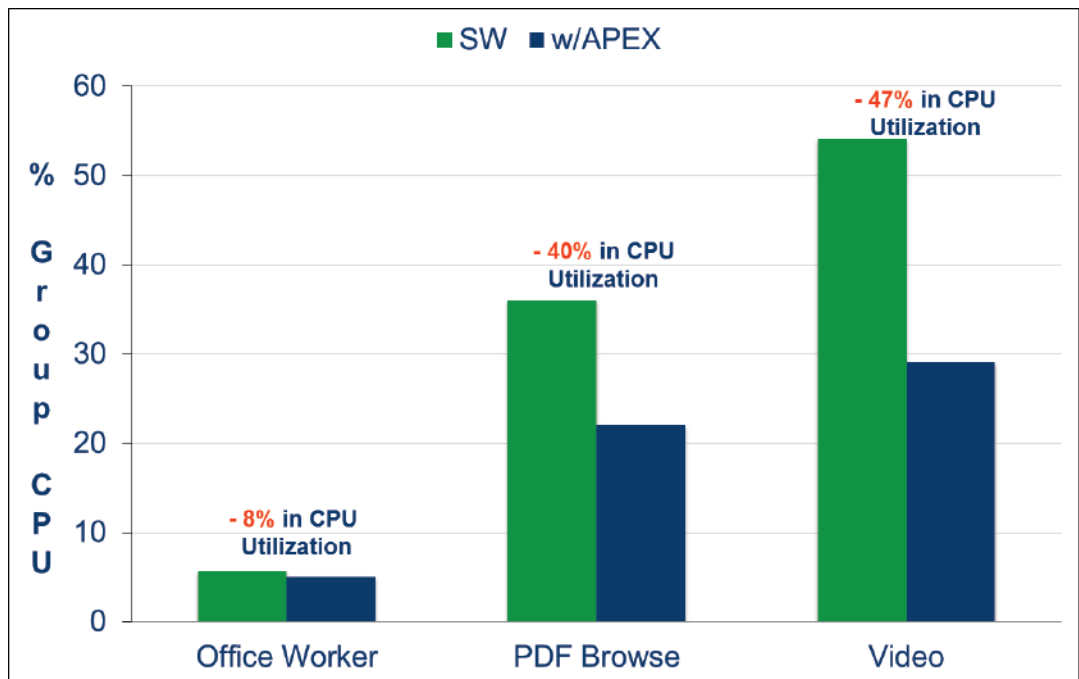


As you can see in the chart above, the two core benefits are being delivered. Initially, there is a 20% reduction in CPU activity on the ESX host that allows for a greater consolidation ratio of VMs to hosts, but plenty of free CPU cycles are left to absorb unexpected CPU spikes. Notice the two CPU spikes that raise activity to almost 100% for the server without an APEX barely register once the APEX encoder is engaged. This contributes to a second benefit: consistency. When a server becomes heavily loaded by higher than expected demand for rendering pixels and engaging codecs, the increased activity is offloaded to the APEX card, significantly reducing the impact to the on-board CPUs and leaving the CPUs free to service the requests of the user applications. One important aspect of all desktop delivery systems is being able to guarantee a high quality of service in all cases. APEX has been shown to reduce these peak loads by as much as 50%.

By removing the difficulty of planning for CPU spikes that are only due to pixel changes, businesses can design VMware ESX clusters based on average load rather than having to spend money on scaling the system to cope with the peak load activity. APEX delivers this without the need for an unnecessary scale out of the ESX host environment, contributing to a smaller server footprint and quantifiable savings in power management and cooling costs.

Businesses that wish to take on the next generation of VDI power users can do so with confidence that their existing environment can on-board these new users without affecting existing users of the VMware View infrastructure. It liberates VDI project stakeholders to focus their time and energy on the users' applications rather than needlessly fretting over scalability issues. Without the APEX, customers will find that scalability begins to flat-line when resorting to server CPU horsepower alone. With the APEX card, customers can continue to increase the ratio of virtual desktops to servers without hitting a brick wall of decreased performance. By deploying the Teradici APEX card, it's possible that your 2.66-GHz CPU will act like a 3.2-GHz CPU.

Clearly, the reduction of CPU activity is not uniform for all environments and all workloads. Much depends on the profile of the users and the kind of work that makes up their daily routine. For this reason, the Teradici APEX card brings many benefits to organizations that render and process many pixel updates and leverage Teradici's intelligent use of codecs to progressively build the users' screen based on the content presented. For users dealing with dense Adobe PDF documents and video, the benefits are significant, as the chart below demonstrates.



Note: Office worker workload included Word, Excel, PowerPoint, PDF reading and web browsing as the sample applications.

Conclusion

While VDI adoption continues to grow in popularity, so are the needs of today's more demanding user base. The more structured type of power user has higher performance and experience expectations. To take the risk out of deploying VDI to this new user base, businesses should consider technologies like the PCoIP protocol in VMware View, which allows desktops to be centrally located and managed in the data center while still providing a high-performing experience to remote users. Teradici's APEX 2800 server offload card empowers IT to protect and ensure a cost-effective, consistent user experience, while enabling increased VDI consolidation ratios. To learn more about Teradici's APEX solution, go to www.teradici.com/APEX.

About the Author

Mike Laverick is a former VMware instructor with 17 years of experience in technologies such as Novell, Windows, Citrix and VMware. Since 2003, has been involved with the VMware community. Laverick is a VMware forum moderator and member of the London VMware User Group. He is also the man behind the virtualization website and blog [RTFM Education](#), where he publishes free guides and utilities for VMware customers. Laverick received the VMware vExpert award in 2009, 2010 and 2011.