SOLUTIONS
WHITE PAPER

Insights Into
Alternative Client Compute Models

Comparing Approaches and Sharing Best Practices

gtsi®

Contents

Executive Summary

Organizations face daily challenges ensuring that end users have a reliable and secure computing environment. A large portion of an IT department's time is currently spent in reactive mode, supporting numerous end-user applications and patching desktops to secure against the latest system vulnerabilities.

In addition, IT departments allocate over one-third of their hardware budget to desktop and laptop computers, given that 75% of organizations replace or upgrade their PCs every three years or less (Info-Tech, 2009). Furthermore, the yearly total cost of ownership (TCO) for a notebook computer can be as much as six to seven times the cost of the device.

TCO is not the only factor to consider when exploring alternative client architectures. Security and data protection should also be major considerations for IT managers. Data on personal computers is harder to secure and more susceptible to loss. A recent research study showed that over 12,000 laptops are misplaced or lost in U.S. airports each week. Research showed that 53 percent of surveyed mobile professionals carry confidential company information and 65 percent of those don't take steps to protect that data (Ponemon Institute, 2008). As the use of mobile notebook computers continue to increase relative to the number of desktop computers, organizations need to look at ways to better secure the mobile end-user computing environment. Data encryption is one approach for securing the data on mobile devices, but IT departments can also look at various client architectures to further enhance security.

The objective of this white paper is to discuss the different alternative compute models available. It evaluates and compares various approaches and brings insight into the challenges that an organization may face when incorporating technologies into their environment. Finally, it shares best practices on how these technologies can be incorporated into an organization's end-user computing environment – leading to improved security, reduced costs, and increased IT efficiency.

Origin and Types of Client Compute Models

With rising costs associated with computing requirements and concerns about efficiency and data security, there has been a rapid adoption of consolidation and centralization of resources and data. This consolidation is expanding to the end-user community through alternative approaches to client computing. One of the most common architectures is server-based computing and the use of thin clients. This concept of having the end user access computing resources and data from a "dumb terminal" goes back to the days of mainframe computers.

What are now called thin clients were originally called "graphical terminals" because they were a natural development of the text terminals that had gone before them. X terminals were a relatively popular form of graphical terminal in the 1990s. Citrix Systems approached Microsoft with an idea for a multi-user version of Windows similar to what had been done with Unix. Microsoft agreed to license the Windows NT 3.51 source code which Citrix then turned into a product called WinFrame – a version of NT 3.51 that allowed multiple users to run on the same server. Microsoft later licensed the technology back from Citrix and incorporated it into a special version of NT 4.0 (known as NT 4.0 TSE, or Terminal Server Edition) and then into all subsequent versions of their server operating systems.

Terminal Services allows the operation of standard Windows software in a centralized computing "mainframe model" versus a distributed computing model. Users log onto the server using devices such as thin clients and the server creates a session in memory dedicated to that user. Any command that would normally be executed on the end user's local device is instead compressed and sent to the client through an efficient delivery protocol.

As the processing power of the personal computer (PC) increased, it created an environment where applications could be run locally on the PC. This enabled the end user to work independently and control the data they generated. This has led to a distributed computing environment where users run robust applications through fast and ubiquitous networks. The larger network bandwidth allows easy downloading of applications and the proliferation of data. As the technology has matured, it has created issues with managing the applications and the data that is generated. Add to this the issues with security and operating system patching and upgrades, and even a small computing environment can become daunting to an IT department. There are several approaches that IT departments can take to better manage these issues.

Considering Alternative Client Compute Models

The current drive to implement green computing and operate more efficiently has triggered a surge in the deployment of alternative client compute models. These approaches provide alternatives for managing those environments as well as having centralized control of data. Data security in general has been an ongoing challenge for IT departments. To further complicate the situation, regulations such as Sarbanes Oxley, HIPPA, SEC Rule 17a, and the Personal Information Protection and Electronic Documents Act (PIPEDA) have placed new requirements on IT departments to not only maintain but also to protect the data they manage, resulting in additional work.

Top Government Drivers:

1. Costs
2. Security
3. Manageability
4. Flexible Remote Access
5. Decreasing Energy Consumption Costs
6. Reducing Environmental Impact

In order to obtain the advantages inherent in an alternative client compute model and minimize the complications posed, several drivers such as standardization, consolidation, management, hardware costs, security, and business continuity need to be considered.

Standardization
By centralizing and standardizing the applications, profiles, and data on servers, IT managers can provide a more predictable user experience while lowering administration costs. One of the most common approaches today is server-based computing. This may be the ideal approach for organizations that primarily have task workers because it provides the greatest number of users per server, resulting in a greater return on investment. A recent scalability analysis conducted by Citrix Systems showed that a Windows Server 2003 running XenApp 5 (64-bit edition) with two dual-core Intel Xeon 5150 processors and 16GB of RAM can support up to 240 concurrent users while still maintaining acceptable performance for Office applications (Citrix, 2008). This approach is often coupled with the use of thin client devices, leading to more substantial predictability and a reduction in desktop management costs.

Consolidation
With a fully centralized server and desktop environment there are savings in relation to economies of scale and system utilization. The average server utilization rate in today's data centers is only 15 percent. This is due to the old adage of one server for each application. By consolidating servers and using them to virtualize desktop services, they are utilized more efficiently and fewer are needed for the same amount of services. This leads to savings in the physical space and a reduction in power and cooling needs.

Management
Unlike basic Terminal Services, where many users share the same working environment within the server, desktop virtualization has enabled a more personalized user experience. This technology has created an environment where the desktop operating systems and applications can be centrally managed and maintained by the IT department while maintaining the same end user experience. These alternative compute models provide efficiencies around configuration management and provisioning. For example, a single OS image that is shared by multiple users can be patched and maintained. This is in contrast to managing an OS on each individual desktop. Since fewer resources are needed to perform tasks such as software updates or changes, IT departments can now focus their resources on proactive versus reactive management of the desktop environment.

Hardware Costs
With the rising costs and environmental concerns around power, the U.S. Environmental Protection Agency (EPA) is working diligently with Energy Star to create new standards for power and cooling in the data center as well as the desktop. Organizations are looking for alternatives to traditional PCs, and one such alternative is what is commonly referred to as thin clients. With thin clients there are direct savings in acquisition costs. They are also less costly to maintain since there are no moving parts (i.e., hard drives) that are susceptible to failure. Thin clients also have an average life span of seven years, which is longer than the typical three- to four-year life span of the average PC. From a power and cooling perspective, shifting from PCs to thin clients can save an organization upwards of 25 percent in power savings, according to IT analysts (Infoworld, 2008). This potential energy savings is driving IT executives to reconsider trading in users' thick clients for thin ones.

Security
Thin clients are inherently secure since they do not have disk drives that can be used to store data. Some thin client implementations are even stateless, meaning that there is no underlying thin OS is resident on the device. Organizations that need to protect confidential information benefit the most from the use of thin clients. Data can be centrally stored and managed on secured servers, allowing IT administrators to have additional control over where the data is stored, who can access the data, and how the data is accessed. Policies can then be defined to enforce security parameters such as which drives are mapped to users and devices (i.e., network and local drives such as flash drives and USB drives) and encryption.

Business Continuity and Remote Access
Another benefit derived from a centralized compute model is the ability to back up desktop images, profiles, and applications so that IT departments can institute a business continuity and disaster recovery plan around the organization's desktop environment. In a traditional PC model, it is often cost-prohibitive and difficult to develop a backup strategy for desktops; therefore, it's not uncommon to leave existing user desktops out of a business continuity plan. In a virtualized and centrally managed desktop environment, since user desktops are now just a set of files that are stored in a shared storage array, these can be replicated and easily backed up in a secondary site for disaster recovery purposes. In other words, the same virtualized desktop that a user connects to at the main office can be backed up, restored, and available at a secondary site. And from the user standpoint, it's irrelevant

Shifting from PCs to thin clients can save an organization upwards of 25 percent in power savings.

what device is used to connect to the virtual desktop. This could be a stateless thin client or even a home computer, provided it has the right client software to connect to the virtual desktop. This model can also foster a more secure and predictable remote access strategy for the organization. Users and management would be more amenable to teleworking if the experience is secure, predictable, convenient, and familiar.

Challenges with Alternative Compute Models

There are various technical and cultural challenges that need to be addressed when considering an alternative approach to traditional computing. Two main challenges include quantifying a return on investment (ROI) in order to get management buy-in, and designing the right architecture to deploy and manage successfully.

Return on Investment and Justifying Costs

Since there is an upfront investment required, one of the challenges of alternative compute models is identifying the return on your investment and justifying the costs needed to get started. This initial investment often includes enterprise-class systems, networks, and storage devices that have substantial capital and operating costs associated with them. Another factor revolves around operating system costs and software licensing. The hard cost of software may outweigh the soft cost savings from improved IT administration. Therefore, it is important to truly understand the organizational needs and identify which, if any, alternative compute models would most benefit the organization. IT departments need to understand both the acquisition and ongoing maintenance and support costs of the desktop environment in order to articulate the benefits derived from an alternative approach to the traditional compute model.

Ensuring a Robust Infrastructure for Centralized Computing

Unlike a traditional desktop compute model where the processing is distributed to each end-point device, a centralized compute model consolidates all the processing needed by multiple users into a few servers. As such, these servers need to be highly available due to the number of users that can potentially be affected by an outage. In addition, these server resources need to be sized adequately to handle sudden increases in utilization during peak periods or withstand a server failure. To properly size the server resources needed to support the organization, it's important to assess the user population in order to categorize the type of workloads present in the environment. Some users might be standard task workers that use just a handful of applications and therefore do not need a lot of processing power. Others might be power users that require graphic intensive applications and dedicated resources. It's also critical to know if any

Common Classifications of Client Compute Architectures

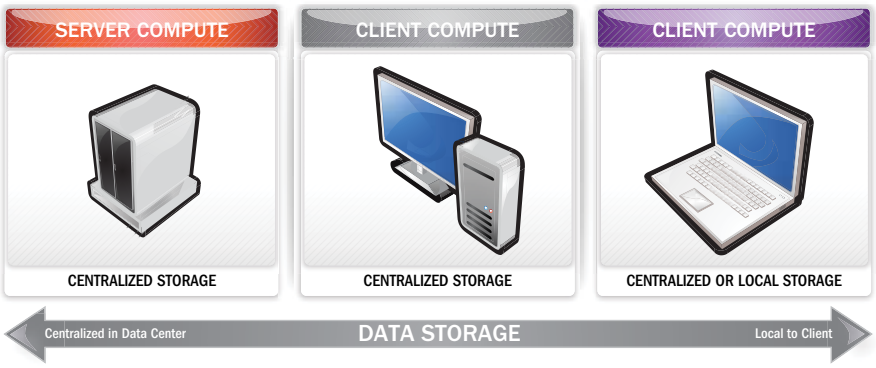

of the users are remote, mobile users, and/or require offline access to applications. By understanding the categories of users in the environment, the number of users that need to be supported, and the type of users (mobile versus non-mobile users), an organization can begin to map the right client architecture for its users.

After an organization has identified the right client architecture for the different user categories and designed the supporting infrastructure needed to support the technology in that environment, the next consideration revolves around ensuring end users get a comparable, if not better, end-user experience than they had before. This can be achieved by monitoring the end-user experience using application performance monitoring (APM) tools, continually monitoring the systems, and checking in with end users after implementing the technology. Organizations need to continuously evaluate performance and make adjustments accordingly.

Selecting the Right Model: Common Classifications of Client Compute Architectures

We can categorize client compute models into three major categories: server compute, centralized storage; client compute, centralized storage; and client compute, centralized and/or local storage.

A central theme in each of these compute models is data. What ultimately differentiates one compute model versus the other is where the data resides in the infrastructure. There are various architecture options around the virtualization of the desktop environment. Each has its advantages and disadvantages and there is not a one-size-fits-all solution. The most prudent approach is to assess the needs and typical workload of the environment and based on the findings implement the proper technologies that best meet the organizational needs.

Server-based Computing

The most common approach in the industry today is the shared desktop model also known as server-based computing. In this model, multiple users share a common desktop and access resources (i.e., CPU, memory, disks) on servers in the data center. Commonly referred to as presentation virtualization, all the processing is centralized on the server and video, mouse movements, and keystrokes are encapsulated and sent to the user via a highly efficient delivery protocol such as Remote Desktop Protocol (RDP) or Independent Computing Architecture (ICA). Citrix XenApp and Microsoft Terminal Services are the most common shared desktop technologies in the market today.

Server-based computing is perhaps the most cost effective in terms of the number of users it can support per server, but a downside is that sessions are not isolated from each other. Performance for all the users on a server can be adversely affected by a rogue process

Server-based Computing Model

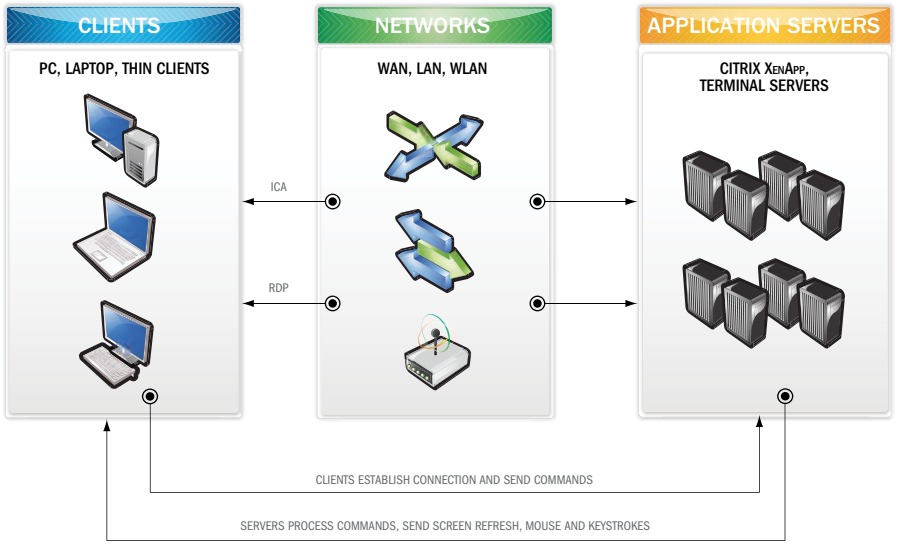

Application Streaming Model

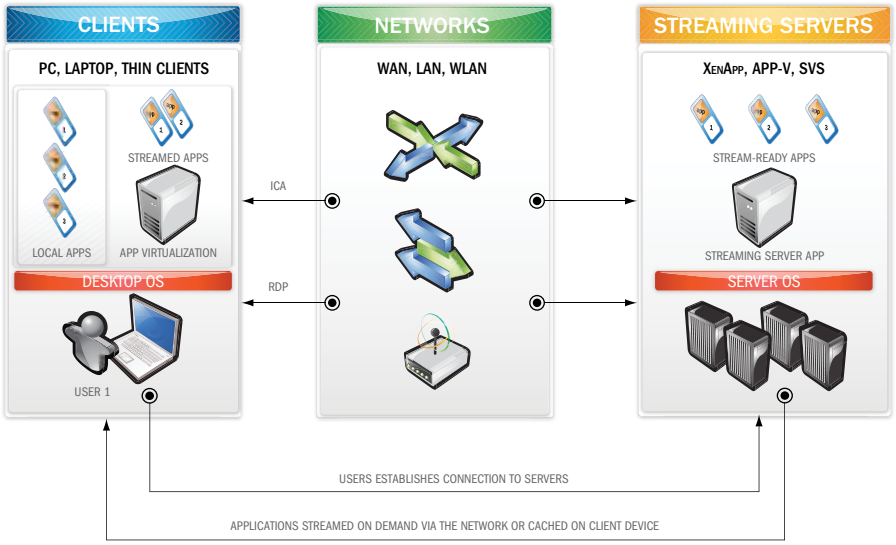

or application in one user's session. There have been some technology advancements in recent years around isolation of applications to minimize the risk of application conflicts which often result in rogue applications. In these instances, the binaries that comprise an application execute in their own virtualized environment, or "sandbox." These isolated applications interact with their own instances of the system files, DLLs, user, and system registries, thus reducing the risk of unstable applications due to modification of system files. There are certain limitations to be aware of with application isolation such as the inability to isolate device and kernel drivers, or isolating Windows services.

Application Streaming

A technology that is becoming more relevant today due to the increased need for support of mobile computing is application streaming. When an application is streamed, a server sends application bits to the user device relying on the device's local computing resources to run the application. Not only does application streaming allow application isolation for users, it also offloads computing resources from the servers and distributes this workload to multiple users, allowing the server to scale up with the number of supported users.

An important use of application streaming is the ability to allow users access to applications while their systems are offline – such as in the case of notebook computers. While users gain the ability to take applications offline with them, IT departments can still maintain control of these applications, configuring how long users can use the application offline, and whether applications can be taken offline in the first place.

One-to-Many PC Blade Model

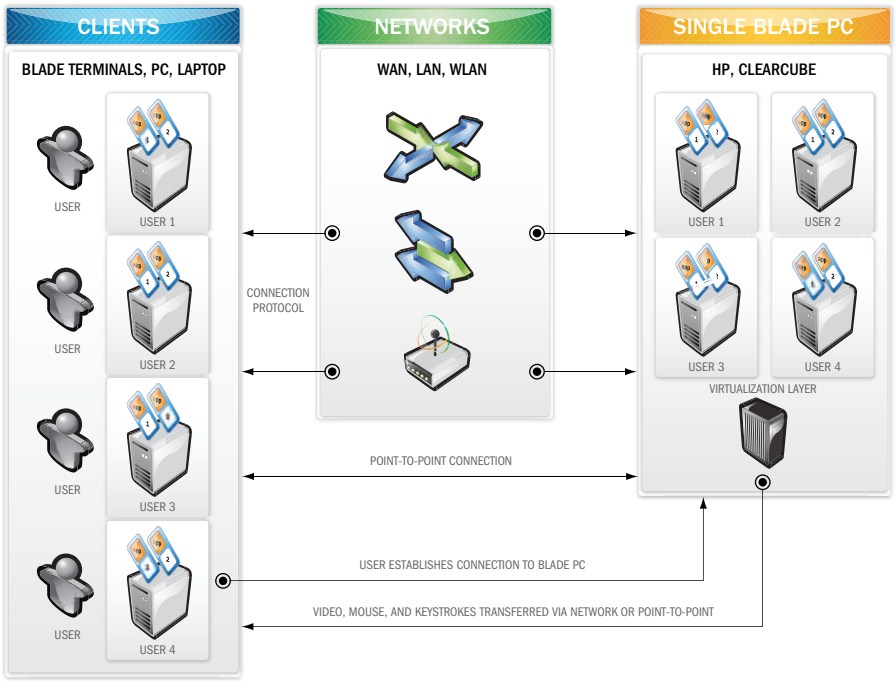

PC Blade Models

Another approach is dedicating a PC-like, blade-class device (PC blades) in the data center for each of the users. In this approach, a 1:1 mapping of user to physical device is feasible. An advantage of this model is that each user gets a dedicated resource allowing full use of all the processing capabilities of the device. Organizations that have strict security requirements can leverage a direct connection to a blade PC through fiber, ensuring a secure point-to-point connection that eliminates comingling of network traffic. IT administrators can still benefit by having these PC blades centrally located and managed in the data center. This approach is ideal for users who require dedicated computing resources or work with graphic-intensive applications such as CAD or GIS-type applications.

PC blade models tend to be the more costly among all the centralized computing models. This has led to the utilization of virtualization software to increase the ratio of users per device, hence reducing the number of PC blades needed and lowering the acquisition cost of physical devices. With the increased capacity and processing capabilities of the latest generation PC blades, it is not uncommon for organizations to have multiple users connect to a single PC blade (one-to-many PC blade model) and leverage virtualization software to isolate these user sessions. Virtualization hypervisors such as VMware Server, VMware ESX, or Hyper-V are supported on PC blades allowing multiple instances of desktops to run on a single physical device.

Virtual Desktop Infrastructure (VDI)

Another approach that is quickly gaining market acceptance is the use of server hardware and a multi-user environment framework (similar to Citrix XenApp and Terminal Services), while utilizing virtualization software on the servers to provide isolation and individual desktops for users. Commonly referred to as "hosted" VDI, or desktop virtualization, this approach is similar to the PC blade environment since each user gets a personalized desktop. The advantages of this model are high utilization of computing resources and individual desktops

VDI Model

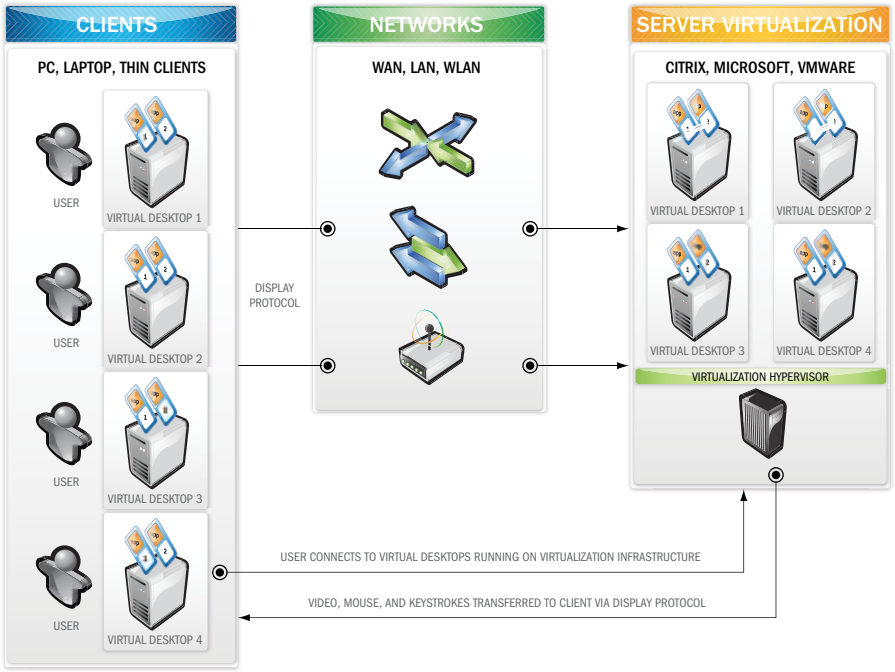

for users. A challenge is the difficulty of quantifying actual cost savings, thus resulting in the lack of organizational buy-in and support. There is also a perception in the market that VDI is not a comprehensive and proven end-to-end solution because of its relative infancy as compared to other server-based computing architectures and its lack of support for offline VDI access. While there are technologies underway that would allow virtual desktops to be taken offline through the use of client-side hypervisors, this has only been experimental and has not been widely used or deployed in production. There is a variant of the VDI architecture where virtualization software, such as VMware Workstation or Microsoft Virtual PC, is loaded onto a computer that has an operating system. This approach is commonly referred to as Type-2 hypervisors, and although this allows for virtual desktop functionality offline, this is not a preferred approach to replacing the traditional desktop model. The common use case for this type of virtualization has been for developers and system testers who have a need to view multiple operating system environments from a single machine.

VDI has evolved quite rapidly in the last few years due to the use of enterprise-wide server virtualization and vendor competition, which is driving the rapid pace of feature enhancements and capabilities. Since VDI leverages the underlying technology available in server virtualization, these enhancements support VDI and its continued growth. New enhancements in server virtualization hypervisors give IT departments the ability to use a single image and allow multiple users to leverage that same image. This eases the management burden in terms of patching and securing operating systems. In addition, this allows more efficient use of storage through drastic reduction of disk requirements, sometimes resulting in up to a 90 percent reduction (SearchStorage.com, 2009). Virtualization vendors are also making improvements in display protocol technologies. For example, VMware has teamed with Teradici to utilize PCoIP (PC-over-IP) in the latest version of their VDI platform. As organizations reap the benefits of server virtualization technologies through

Client Virtualization Model

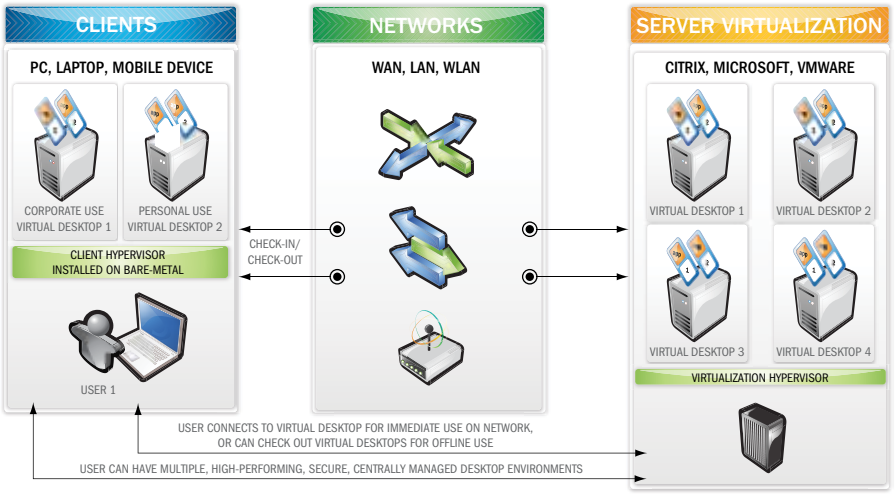

consolidation of physical servers in the data center, IT managers now have the ability to increase the return on investment through the virtualization of desktops.

Client Virtualization

Client virtualization addresses the limitation of the hosted VDI model for mobile users who need offline access to virtual desktops. Client virtualization software installs on bare-metal of end-user devices, such as laptops, similar to how type-1 virtualization hypervisors are installed on server hardware. The software that is installed locally on the client works in conjunction with the management server and host servers in the data center to allow check-in/check-out functionality for virtual desktops. By having centralized management of the desktops in the data center, IT departments can maintain a single OS image to provision to end users. In addition, IT security is enhanced since data is stored centrally, and the desktops that are provisioned to client devices are encrypted, requiring multi-factor authentication and/or controlled via policies by the IT department. The benefits to the end users are a better desktop experience by leveraging local resources such as memory, CPU, and graphic processors on the device; offline access to the desktop; and the ability to run multiple environments (i.e., personal and corporate desktop images) securely on the end-user laptop, thus eliminating the need to carry or manage numerous devices.

There are other innovative approaches to improving the end-user computing experience while allowing organizations to reduce the cost associated with enterprise desktop management. Often, these are narrower in use and address the unique requirements of a smaller subset of users. An example of such technology is OS streaming, similar to Citrix Provisioning Server. OS streaming technology applies well to diskless PC environments found in training labs or call centers.

In summary, client compute models found in organizations today range in scope from the traditional personal computer (PC), to operating system (OS) and application streaming that can be used with a variety of endpoint devices such as stateful and stateless thin clients, traditional PCs/laptops, and diskless workstations. The following chart depicts certain capabilities for the various client compute models. Depending on the categorization of the user population, one or a combination of any of these models could be a viable technology approach for an organization.

Energizing National Nuclear Security Administration's (NNSA) Cyber Security Measures with Client Computing

With responsibility for the Department of Energy's (DOE) classified networks, NNSA is vigilant in the application of cyber protection measures. The use of classified removable electronic media (CREM) to store information presented a persistent security challenge to NNSA. When a federal mandate for increased cyber security was issued, DOE's chief

Table 1. Compute Model Comparison

	Traditional PC Client	Terminal Services	Virtual Desktop Infrastructure	Blade PCs	OS Streaming	Application Streaming	Client Hypervisor
Typical End-Point Devices	Desktops, Laptops	Thin Clients, Desktops, Laptops	Thin Clients, Desktops, Laptops	Thin Clients, Desktops	Desktops, Laptops	Desktops, Laptops	Desktops, Laptops
Application Execution	Client or Server	Server	Server	Server	Client or Server	Client or Server	Client
Mobility	Yes	No	No	No	No	Yes	Yes
Application Storage	Client or Server	Server	Server	Server	Server	Client or Server	Client or Server
Leading OEM Providers	Dell, HP, IBM	Citrix, Microsoft, Sun	Citrix, Microsoft, VMware	ClearCube, HP	Citrix, Lenovo, Wyse	Citrix, Microsoft, Symantec, VMware	Citrix, Neocleus, VMware

information officer, in partnership with NNSA, proposed conducting a thorough evaluation of commercially available thin client (diskless) technology to determine its viability in reducing either the intentional or unintentional mishandling of classified information. The agency's objective was to design and implement a diskless technology solution that would permit desktop IT functions to be performed without risk.

One of the key requirements for the solution was that it had to easily integrate into the existing infrastructure. As such, it was decided that commercially available technologies around hardware, software, power, cooling, and network components from multiple vendors were to be evaluated. The selection of products, including those from Ardence, Cisco, Citrix, ClearCube, Decru, Dell, HP, NetApp, LG, RSA, and Symbio, among several others, ensured that the recommended diskless technology solution met NNSA's requirements for security, reliability, deployability, certifiability, interoperability, and scalability.

Of primary importance was how easily the components could be integrated into existing NNSA backend infrastructure and their ability to meet specifications for:

- Transparent deployment and integration with NAS, DAS, IP-SAN, FC-SAN, and tape
- Wire speed encryption of data at rest for stored data protection
- Strong access controls, authentication, and tamper-proof auditing
- No application/database changes or downtime
- Native support for NFS, CIFS, iSCSI, Fibre Channel, and SCSI
- Operating system agnostic; no software agents required
- Secure, enterprise-wide, and lifetime key management

Since the architecture was built to be modular, NNSA has the ability to add or decrease layers of security according to mission requirements. After the pilot program was successful, this turnkey, stateless thin client solution has since been implemented at other DOE and NNSA facilities.

Careful Planning Enables a Better Managed Client Environment

Successful initiatives begin with sound planning strategies. To enable a successful client compute approach, we recommend focusing on the planning phase with strong emphasis on exploring alternatives, assessing the current client environment, understanding future end-user computing needs, monitoring the current desktop environment, and developing a proof of concept (POC) before implementing throughout the organization.

Assess

The first step in enabling alternative client compute models is to start with a thorough assessment of the current

IT environment. In this initial phase, engineers analyze relevant aspects of the environment that affect the design and deployment of client compute technologies, and engineers identify business requirements and define the metrics to be used for measuring results. Gaps are identified and documented, and an initial baseline of the environment is developed. As part of the assessment, an end-user classification is established. This addresses common-base requirements for all end-users, as well as the computing environment needed to support them. Information such as operating systems, network topology, interfaces, and the required applications should be gathered. Then end-user profiles for data and applications can be established.

The combination of this information can then be used to define the user categories and profiles for the organization. It is also necessary to categorize the end users into areas of responsibility and functionality. This will require an understanding of the computing requirements of the end users and the creation of categories that will be used as functional templates according to job requirements. A survey of the requirements needs to be performed. It is imperative that the information be gathered in a comprehensive fashion in order to create the highest amount of accuracy in the final design. An assessment of the end-user computing requirements will result in a design of a client compute model with the capabilities needed to support the functionality of the mission.

Monitor

By monitoring the actual use of resources from the current distributed desktops for a full month, IT departments can get a sampling and true baseline of data. A 31-day evaluation captures a full monthly cycle of usage patterns. This provides the data required for determining the amount and type of resources (servers, storage, network, and client devices) needed to support the desktop services for end users. This information can then be used to further determine configurations for the infrastructure, such as the number of users per server and LUN sizing for the storage. It is important to consider all application use cases – especially those requiring graphics, like Microsoft PowerPoint, Adobe Photoshop, or Autodesk AutoCAD, to determine the type of client compute model suitable for the workload.

Proof of Concept

Once all the data and a high-level design have been assembled, the next step is to implement a proof of concept to determine what components of the client architectures best meet the needs of the end user community. Using all the information from the assessment, a categorization schema can be built to enable all levels of functionality – from the most basic to the most advanced users. This schema can also be used to design the solution. Most of the desktop configurations presently deployed will be standardized, and these similarities can be used to create a functional and scalable desktop environment for all users. One of the goals of the proof of concept is to ensure minimal impact and training for the end users.

Fully Managed Client Environment

Assess, monitor, and proof of concept are the three key steps that will enable a successful implementation of alternative compute models. The end state should be an IT environment that has standards for managing the different types of users and the capabilities they require. It should allow for users to change profiles as necessary and give them access to newer functionality while maintaining security requirements and data integrity throughout the environment.

However, this process doesn't end with implementation. Ongoing monitoring and assessment of the technology is recommended. The client compute models must remain current in order to continue to support the organization. A support and refresh plan and schedule should be set up with certified engineers to monitor the technology components and ensure they remain functional. As it becomes necessary to replace the technology, an asset disposal plan should also be executed. This will track the final disposal of the assets while adhering to the organization's security policies.

Conclusion

Desktop virtualization offers many advantages for application and desktop delivery. GTSI recommends IT organizations take a strategic view of the opportunities presented by this new technology and evaluate the different client compute models available. By doing this, an organization will be able to select the models that work best within their environment, delivering the most benefits and a strong ROI.

For more information on GTSI and our virtualization solutions, visit GTSI.com or GTSI.com/virtualization, or call us at 800.999.GTSI.

References

Citrix, *XenApp 5 Scalability Analysis v. 1.0*, 2008

Info-Tech Research Group, *PC Management for Cost Savings and Budget Bliss,* 2009

Ponemon Institute, *Airport Insecurity: The Case of Missing and Lost Laptops,* 2008

Samson, Ted, *"Lower Energy Bills, Longer Lifecycle Boost Thin Client's Green Appeal,"* InfoWorld.com, March 12, 2008

Silwa, Carol, *"Virtual Desktop Infrastructure Tutorial: Part 2",* SearchStorage.com, February 24, 2009

Supplemental Material

Citrix Systems, Inc., *Accelerate Business Through a Cost-efficient Virtual Workforce, 2009*

GTSI Corp., *Thin Client Solution Energizes NNSA Cyber Security Measures, 2009*

VMware, Inc. & Server Centric Consulting, *Virtual Desktop Infrastructure: Deployment Considerations, 2008*

VMware & NetApp, *Comprehensive Virtual Desktop Deployment with VMware & NetApp, 2008*

About the Authors

Thomas Barna
Senior Enterprise Solution Consultant, GTSI

Thomas Barna has 16 years of experience in Information Technology designing and implementing IT solutions for both public and private sector customers. Currently at GTSI, Mr. Barna is a senior enterprise solutions consultant focused on helping federal civilian agencies solve their IT challenges through the design and architecture of data center and networking solutions. Prior to joining GTSI in 2004, Mr. Barna designed the backup and archive solution that is now used for all magazines and books produced by National Geographic. He holds technical certifications from Sun, NetApp, VMware, EMC, Symantec, and IBM, and is ITIL Foundation v3 certified.

Miguel Sian
Senior Enterprise Solution Consultant, GTSI

Miguel Sian has 12 years of experience in Information Technology, which includes system architecture, design, implementation, and administration in both the public and private sectors. Currently at GTSI, Mr. Sian is a senior enterprise solutions consultant focused on helping federal civilian agencies solve their IT challenges through the design and architecture of data center solutions, and by leveraging technologies such as server and desktop virtualization. Prior to joining GTSI in January 2006, Mr. Sian led a team of systems engineers that supported the largest-ever launch of new, low-fare service in the history of the airline industry. He holds a Bachelor of Science degree in Computer Science from the University of Maryland University College, and holds technical certifications from Citrix, CompTIA, Microsoft, NetApp and VMware, and is ITIL Foundation v3 certified.

About GTSI

GTSI Corp. is the first information technology solutions provider offering a Technology Lifecycle Management (TLM) approach to IT infrastructure solutions delivered through industry-leading professional and financial services. GTSI employs a proactive, strategic methodology that streamlines technology lifecycle management, from initial assessment to acquisition, implementation, refresh, and disposal. TLM allows government agencies to implement solutions of national and local significance quickly and cost-effectively. GTSI's certified engineers and project managers leverage strategic partnerships with technology innovators. These experts use proven, repeatable processes to design, deploy, manage, and support simple to complex solutions, to meet governments' current and future requirements and business objectives. GTSI is headquartered in Northern Virginia, outside of Washington, D.C.

2553 Dulles View Drive, Suite 100, Herndon, VA 20171-5219 · 800.999.GTSI · GTSI.com