White Paper

Intel® Core™ vPro™ Processors

Emerging Compute Models

Understanding Desktop Virtualization

Driven by ever-increasing pressure on a multitude of issues – including cost control, manageability, security, regulatory compliance, and business continuity – some IT managers are considering desktop virtualization models as an alternative to traditional distributed software deployment.

Thanks to recent advances in technology, IT has a range of centrally managed models from which to choose. From terminal services to VHD, blade PCs through application and OS streaming, each model has its advantages and trade-offs.

Rarely is a single model adequate to meet the needs of all users and all applications. Instead, the needs of each group of users, as well as IT requirements and existing infrastructure, must be carefully considered. In most cases, an enterprise solution will likely involve a combination of models.

Although thin terminals may be suitable hardware for certain purely server-based models, intelligent PC platforms based on technology such as the Intel® Core™ vPro™ processor family for desktops and laptops offer highly manageable and secure platforms on which to deploy a wide range of solutions.

This white paper describes the variety of desktop virtualization models available, how each works, and the advantages and limitations of each model for key vectors such as security, manageability, and power consumption. Also identified are the major providers of software solutions and the end-point devices to deploy each model.

Table of Contents

Evaluating Desktop Virtualization Models

Choosing the appropriate model requires balancing several interrelated factors. Excessive focus on only one of these factors is likely to result in a less than optimal solution. For example, focusing exclusively on IT requirements may result in an easy-to-manage system that is difficult for users to use.

Be sure to consider the application and business needs of the entire user base to understand which applications and data it makes sense to centralize versus installing locally.

It's helpful to segment users based on tasks performed and applications required. Individual users may need to access different applications using a mix of models. Rarely will a single model meet all the needs of a given group of users, let alone the needs of every user in the enterprise.

Decisions about which model to use often get intertwined with the client device on which it will be deployed. It's important to consider these topics separately.

For example, your business scenario may dictate server-based computing for a certain application. However, this server-based model, often referred to as "thin client," does not necessarily have to be deployed on a thin terminal. A desktop or laptop PC may actually be a more appropriate device, depending on the user's total application and mobility needs. The user may also require locally executed applications. The client device should accommodate the most demanding requirements of target users.

For cost comparisons, it is suggested that the total cost of the complete solutions are compared under different models. Costs may just shift from the clients to the data center and vice versa, without actually reducing the total delivery cost. Also, infrastructure costs for equivalent processing and storage capacity can be significantly different, depending on whether these are located in an endpoint or delivered from a data center.

Table 1. Factors to Consider

IT Requirements Standard IT concerns, such as security, manageability, business continuity, etc.	▪ Security ▪ Image management ▪ License management ▪ Support structure ▪ Disaster recovery ▪ Lifespan of investment
Infrastructure Available hardware, connectivity, and bandwidth. If the infrastructure is not in place to support a particular model, either it needs to be purchased and installed or another model must be chosen.	▪ Servers ▪ Storage ▪ Data center space, power, and cooling ▪ Network bandwidth ▪ Budget priorities
User Experience The workflow needs of the system's users, including the need for mobility and performance. In many cases, either by custom or policy, users expect that the device may be used for some personal tasks, which may result in personally identifiable information being on the system. These issues of "ownership" and privacy should be considered.	▪ Mobility ▪ Responsiveness ▪ Customization ▪ Connectivity ▪ "Ownership" and privacy
Application Workload The computing and graphics demands of the applications the user will be running. Some applications may be intolerant of network delays, such as VoIP or streamed video. Headroom for future application growth should also be considered.	▪ Compute load ▪ Graphics load ▪ Delay sensitivity (e.g., video, motion graphics, VoIP) ▪ Headroom ▪ Web server load

Table 2. Desktop Virtualization Model Comparison.

	Terminal Services	Virtual Hosted Desktop (VHD)	Blade PCs	OS Image Streaming	Remote OS Boot	Application Streaming or Application Virtualization	Virtual Containers
Application Execution	Server	Server	Server	Client	Client	Client	Client
Application Data Storage	Server	Server	Server	Server	Server	Client or Server	Client or Server
Local Device Connect and Synch (Bar Code Reader, PDA, Phone, etc.)	Partial (vendor-specific)[a]	Partial (vendor-specific)[a]	Partial (vendor-specific)[a]	Yes	Yes	Yes	Limited
Full Windows* Application Support (Including VoIP and Rich Media)	Partial (vendor-specific)[a]	Partial (vendor-specific)[a]	Partial (vendor-specific)[a]	Yes	Yes	Yes	Yes
Full Support for Microsoft Windows XP,* Windows Vista,* and Windows 7*	Partial (vendor-specific)[a]	Partial (vendor-specific)[a]	Partial (vendor-specific)[a]	Yes	Yes	Yes	Yes
Off-Network Mobile Option	No	No	No	No	No	Yes	Yes
Typical Clients	Terminal, desktop PC, laptop PC	Terminal, desktop PC, laptop PC	Terminal, desktop PC	Desktop PC	Desktop PC	Desktop PC, laptop PC	Desktop PC, laptop PC
Major Solution Providers	Citrix, Microsoft	Citrix, Microsoft, Red Hat, VMware	ClearCube, HP, Dell, Devon IT	Citrix, Dell, Lenovo	Lenovo	Citrix, Microsoft, Symantec, VMware	Citrix, Microsoft, VMware
Major Solutions*	Citrix XenApp, XenDesktop (HDX[a]) Microsoft Remote Desktop Services	Citrix XenDesktop (HDX[a]) Microsoft VDI (RemoteFX[a]) VMware View (PCoIP[a])	ClearCube (four blade workstation models) HP Consolidated Client Infrastructure Dell Dedicated Remote Workstation DevonIT HC12 Remote Workstation	Citrix XenDesktop w/ Provisioning Server for Desktops Dell On-Demand Desktop Streaming (ODDS)	Lenovo Secure Managed Client (SMC)	Citrix XenApp (XenDesktop) Microsoft System Center Config Manager + App-V Symantec Workspace Virtualization VMware ThinApp	XenClient Type 1 Microsoft MED-V and VPC7 VMware View

[a] Several vendors offer enhanced proprietary remoting protocols which provide differing levels of media, graphics, and peripheral support. In some cases customized HW is also required.

Server-Based Models

Terminal Services

Summary

Terminal services is a time-tested and reliable server-based model dating back to mainframe computing. For supported software applications, this model offers strong security and manageability. However, users accustomed to the PC experience may find this model unsatisfying in terms of performance, customization, flexibility, and mobility.

Most large enterprises use terminal services for some applications, especially where security is essential and users are in a fixed location with constant network access. Bank tellers accessing the transaction system, call center workers entering orders, and healthcare professionals working with patient records are examples where terminal services may be a good solution.

PCs are a good platform when terminal services is used only for a few key applications, while the rest are locally installed. Thin clients are appropriate only for an environment consisting of 100 percent terminal services. And even then, replacing large numbers of existing PCs with new terminals can be costly. These costs should be closely evaluated to determine true return on investment (ROI). Many IT organizations choose to use existing or waterfalled PCs as terminal services clients.

How it Works

The client is merely a display and input device. All computation is done centrally on the server, and all data is stored in a data center. Nothing is executed or persistent on the client. Usually, Remote Display Protocol (RDP) or Independent Computing Architecture* (ICA*) is used to push an image of the server-based application to a terminal viewer on the client.

Figure 1. Terminal Services Architecture

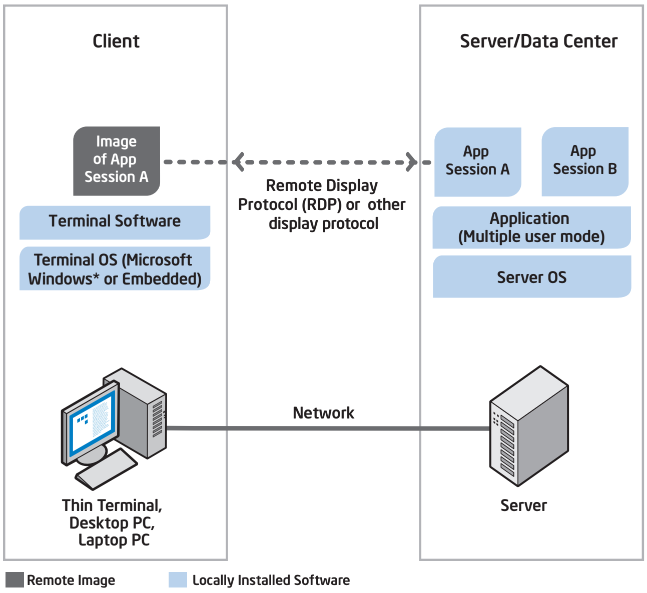
Client
Server/Data Center
Image of App Session A
App Session A
App Session B
Remote Display Protocol (RDP) or other display protocol
Terminal Software
Application (Multiple user mode)
Terminal OS (Microsoft Windows* or Embedded)
Server OS
Network
Thin Terminal, Desktop PC, Laptop PC
Server
Remote Image
Locally Installed Software

Table 3. Advantages of Terminal Services

Enhanced Security	▪ With the OS, applications, and data locked down in the data center, this model has lower risk of a security breach or data loss via the client than client-based models.
Simplified Manageability	▪ Terminal services is a mature and well-understood technology. ▪ Application and data management are centralized, allowing simpler administration and more reliable backup. ▪ Software image management, validation, and support are simplified. Driver and dynamic link library (DLL) conflicts are reduced. Adding, moving, and changing users is simple.
Lower Cost of Incremental Software Deployment	▪ Most enterprises already have some terminal services, so adding new applications may not require much in the way of new infrastructure or software. ▪ More users can be supported per terminal server than other server-based models, such as VDI.
Remote Access	▪ Centralized computing allows access from any network-connected client. Users do not need to be at "their" workstation.
Disaster Recovery and Business Continuity	▪ In the event of a data center or worksite disaster, work can shift to other sites relatively easily, assuming redundant servers and data storage are in place.
Reduced Client Power Consumption	▪ With thin client terminals, client power consumption is lower than most desktop PCs. However, total power consumption may remain equivalent due to increase in power consumption in the data center.

Table 4. Limitations of Terminal Services

Performance and Responsiveness	▪ With even moderately compute-intensive applications, system responsiveness degrades as the number of users increases. Network bandwidth and loading become major factors in client system performance. Networks must have sufficient peak capacity to support the number of terminal users. ▪ Remote Display Protocol (RDP) creates a serious graphics bottleneck. Motion graphics such as video or Adobe Flash* do not perform effectively over RDP. A higher performance third-party display protocol may be required. ▪ Users expect rapid response to mouse clicks and keystrokes. Terminal servers must be physically located close enough to their clients to meet responsiveness requirements.
Software Compatibility	▪ Not all software runs under terminal services. Also, Voice over Internet Protocol (VoIP), streaming media, and compute- or graphics-intensive applications may be poorly suited to server-based computing.
Lack of Mobility	▪ Application delivery via terminal services requires a persistent network connection with adequate bandwidth. Wireless laptops or tablets can be used with terminal services, but the session ends if the device disconnects from the network.
Cost of New User Deployment	▪ New deployment to numerous users may involve significant expense, such as new servers, software, network storage, and network infrastructure. ▪ Acquisition cost of thin client terminals is similar to many desktop PCs. Initial cost of terminal services is about the same as VDI, but higher than streaming or well-managed PCs.[1]
Single Point of Failure	▪ Loss of data center or network function takes all users offline unless backed by a redundant system.
Lower User Satisfaction	▪ Users are conditioned to the PC experience in terms of performance, customization, flexibility, and mobility. Inappropriate or overly strict implementation of terminal services can result in user dissatisfaction and complaints.

Virtual Hosted Desktop (VHD)

Summary

Virtual hosted desktop (VHD, previously called virtual desktop infrastructure or VDI) is a newer model that many organizations are evaluating. VHD is designed to offer the responsive and customizable user experience of intelligent distributed computing along with the management and security advantages of server-based models. It promises centralized management of the entire desktop image.

As with other server-based models, performance and responsiveness vary depending on number of users, physical distance, and type of application. Performance at the client endpoint is especially important. Otherwise, video, Adobe Flash,* Voice over Internet Protocol (VoIP), and other compute- or graphics-rich applications are not well suited to this model without additional media acceleration. However, recent advancements in remote desktop protocols, such as HDX, RemoteFX, and PC over IP (PCoIP), which utilize the local resources of intelligent client endpoints, make this a more desirable solution. (HDX and PCoIP improve the user experience over WAN, while RemoteFX provides improvements for LAN/non-WAN users and requires a server GPU card to encode media across the RDP channel.)

VHD requires a persistent network connection, so it's not appropriate where off-network mobility is required.

All client computation, graphics, and memory resources must be built into the data center, and the storage system must accommodate OS, applications, and data for each user. Cost of this infrastructure needs to be considered against potential TCO savings in manageability.

VHD can be effectively combined with other desktop virtualization models such as application streaming. Application data can be used in conjunction with OS streaming into a virtual machine on the server.

How it Works

As with terminal services, all computation and storage are centralized, with application images pushed over the network to the client via Remote Display Protocol (RDP) or other display protocols. The major difference is that VHD can offer each user his or her own complete virtual machine and customized desktop, including the OS, applications, and settings.

Figure 2. VHD Architecture

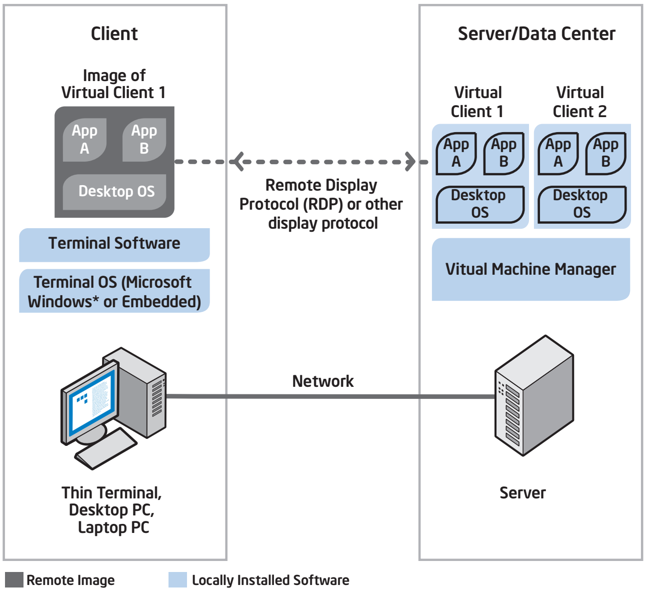
Client
Image of Virtual Client 1
App A
App B
Desktop OS
Terminal Software
Terminal OS (Microsoft Windows* or Embedded)
Thin Terminal, Desktop PC, Laptop PC
Remote Display Protocol (RDP) or other display protocol
Network
Server/Data Center
Virtual Client 1
Virtual Client 2
App A
App B
Desktop OS
App A
App B
Desktop OS
Vitual Machine Manager
Server
■ Remote Image
□ Locally Installed Software

Table 5. Advantages of VHD

Performance	▪ VHD can take advantage of client computing resources to improve performance at the endpoint. ▪ Recent advancements in remote desktop protocols, such as HDX and PC over IP (PCoIP), can improve local performance of video, Adobe Flash,* Voice over Internet Protocol (VoIP) and other compute-intensive applications, with possible increased requirements on the client.
Enhanced Security	▪ With the OS, applications, and data locked down in the data center, this model has lower risk of a security breach or data loss via the client than client-based models.
Simplified Manageability	▪ Application and data management are centralized, allowing simpler administration and more reliable backup. ▪ Software image management, validation, and support are simplified. Driver and dynamic link library (DLL) conflicts are reduced.
User Customization	▪ Each user owns a full virtual machine on the server, allowing PC-like personalization of preferences and settings.
Remote Access	▪ Centralized computing allows access from any network-connected client. Users do not need to be at "their" workstation.
Disaster Recovery and Business Continuity	▪ In the event of a data center or worksite disaster, work can shift to other sites relatively easily, assuming redundant servers and data storage are in place.
Reduced Client Power Consumption	▪ With thin client terminals, client power consumption is lower than most desktop PCs. However, total power consumption may remain equivalent due to increase in power consumption in the data center.

Table 6. Limitations of VHD

Performance and Responsiveness	▪ With even moderately compute-intensive applications, system responsiveness degrades as the number of users increases. Network bandwidth and loading become major factors in client system performance. Networks must have sufficient peak capacity to support the number of VHD users. ▪ Remote Display Protocol (RDP) creates a serious graphics bottleneck. Motion graphics such as video or Adobe Flash* do not perform effectively over RDP. Newer display protocols, such as HDX and PC over IP (PCoIP), may be required. ▪ Users expect rapid response to mouse clicks and keystrokes. Terminal servers must be physically located close enough to their clients to meet responsiveness requirements.
Network Performance	▪ RDP network traffic may fluctuate greatly with media and increased screen image movement.
Manageability	▪ Although software images don't reside on the client, IT must still manage, update, and patch all the virtual desktop images now stored in the data center.
Software and Device Compatibility	▪ Applications, devices, and their associated drivers are intermediated with the hardware by the virtual machine manager (VMM). Devices that are highly specialized or software that requires direct hardware interaction may experience compatibility issues.
Lack of Mobility	▪ Desktop delivery via VHD requires a persistent network connection with adequate bandwidth. Wireless laptops or tablets can be used with VHD, but the application session ends if the device disconnects from the network.
Cost of New User Deployment	▪ New deployments of VHD to large numbers of users may involve significant expense. All client computation, graphics, and memory resources must be built into the data center, and the storage system must accommodate OS, application, and data for each user. New hardware, networking, and building space may be required. ▪ The number of users supported per VHD server is lower than other server-based models. ▪ Acquisition cost of thin client terminals is similar to many desktop PCs. ▪ New media redirection technologies such as RemoteFX require the addition of a server GPU card, limiting the number of virtual machines or displays that can be supported.
Single Point of Failure	▪ Loss of data center or network function takes all users offline unless backed by a redundant system.
Lower User Satisfaction	▪ Users are conditioned to the PC experience in terms of performance, flexibility, and mobility. Inappropriate or overly strict implementation of VHD can result in user dissatisfaction and complaints.

Blade PCs

Summary

Blade PCs bring client computing into a central location, promising higher manageability and security than intelligent distributed computing through restricted physical access, software imaging policies, and limits on user activities.

If each user is assigned a single PC blade (one-to-one), the model most closely resembles intelligent distributed computing, except it can only be used in a fixed location and users must be constantly connected to the network. If individual PC blades service multiple users simultaneously (one-to-many), this model more closely resembles VHD.

The proprietary nature of blade computing makes it relatively costly to implement and requires IT organizations to make a long-term commitment to a vendor-specific architecture. Switching vendors or models may require a complete "forklift upgrade." In addition, one-to-many blade PCs may demonstrate some of the same drawbacks as VHD: performance issues and high infrastructure costs.

How it Works

Blade PCs repartition the PC, leaving basic display, keyboard, and mouse functions on the client, and putting the processor, chipset, and graphics silicon on a small card (blade) mounted in a rack on a central unit. OS, application, and data storage are centralized in a storage array.

Unlike server blades, PC blades are built from standard desktop or mobile processors and chipsets. The central unit, which supports many individual blades, is secured in a data center or other IT-controlled space. In some cases, remote display and I/O are handled by dedicated, proprietary connections rather than using RDP over the data network.

Blade PC vendors initially targeted a user-to-blade ratio of one-to-one, where each user was dynamically assigned a blade and had exclusive use of it. As blade solutions and virtualization software have advanced, most vendors are now offering one-to-many capabilities.

Figure 3. One-to-One Blade Architecture

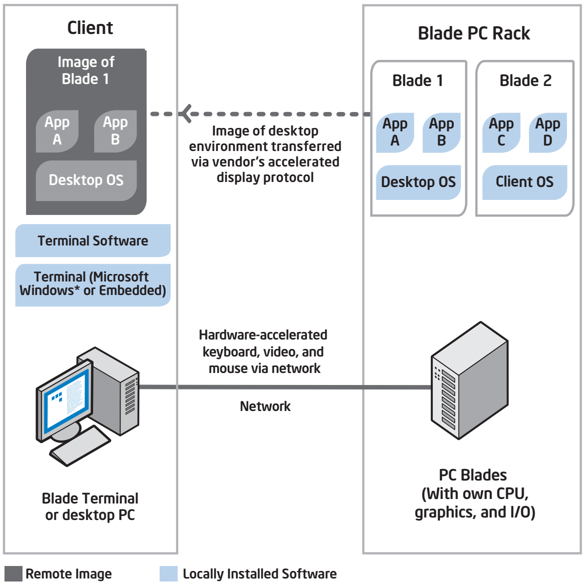
Client
Image of Blade 1
App A
App B
Desktop OS
Terminal Software
Terminal (Microsoft Windows* or Embedded)
Blade Terminal or desktop PC
Image of desktop environment transferred via vendor's accelerated display protocol
Hardware-accelerated keyboard, video, and mouse via network
Network
Blade PC Rack
Blade 1
App A
App B
Desktop OS
Blade 2
App C
App D
Client OS
PC Blades (With own CPU, graphics, and I/O)
Remote Image
Locally Installed Software

Figure 4. One-to-Many Blade Architecture

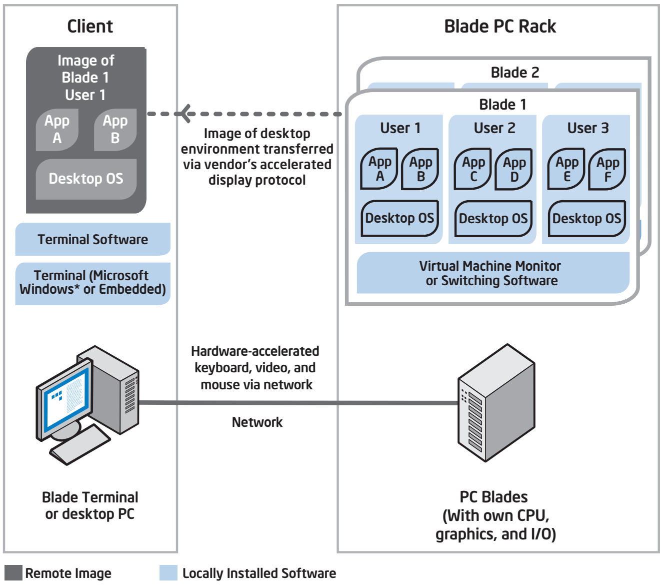
Client
Image of Blade 1 User 1
App A
App B
Desktop OS
Terminal Software
Terminal (Microsoft Windows* or Embedded)
Blade Terminal or desktop PC
Image of desktop environment transferred via vendor's accelerated display protocol
Hardware-accelerated keyboard, video, and mouse via network
Network
Blade PC Rack
Blade 2
Blade 1
User 1
User 2
User 3
App A
App B
App C
App D
App E
App F
Desktop OS
Desktop OS
Desktop OS
Virtual Machine Monitor or Switching Software
PC Blades (With own CPU, graphics, and I/O)
Remote Image
Locally Installed Software

Table 7. Advantages of Blade PCs

Security	▪ With the OS, applications and data locked down in the data center, this model has lower risk of a security breach or data loss via the client than client-based models.
Manageability	▪ OS, application, and data management are centralized, allowing simpler administration and more reliable backup. ▪ PC blades offer a uniform hardware platform, which simplifies validation, image management, and support. ▪ Users can be dynamically assigned any PC blade that is available. Adding, moving, and changing users is relatively simple.
User Customization	▪ Each user can have a unique OS and application image, allowing PC-like personalization of preferences and settings.
Remote Networked Access	▪ Centralized architecture allows users access from any blade-connected client, or in some cases, any Internet-connected client. Users do not need to be at "their" workstation for access.
Disaster Recovery and Business Continuity	▪ In the event of a blade rack or worksite disaster, work can shift to other sites where redundant infrastructure is available. ▪ A user can easily be migrated to another blade if their assigned blade fails.
Lower Client Power Consumption	▪ Power consumption, heat, and fan noise at the client device are lower than most desktop PCs, although total power consumption of clients, blades, and associated storage may be comparable.

Table 8. Limitations of Blade PCs

Performance	▪ In a one-to-many deployment, application performance may degrade depending on the number of users and their workloads.
Manageability	▪ Although software images don't reside on the client, IT still must manage, update, and patch all of the centralized desktop images stored in the data center.
Vendor Lock-in	▪ Blade PCs are not standardized, and each vendor has a proprietary implementation. Once IT has selected a vendor's blade architecture, the costs of switching may be extremely high. Switching vendors or models requires a complete hardware upgrade in most cases. ▪ Available management tools may be limited, and IT may be dependent on their blade vendor's tools and development schedule.
Lack of Mobility	▪ No mobile option exists for blade PCs. They are only suitable for users with a persistent network connection.
Higher Cost Per User	▪ Due to their non-standard architecture, blade PC acquisition costs per user are higher than other models. ▪ The storage system must accommodate the OS, application, and data for new users. New hardware, networking, and building space may be required.
Single Point of Failure	▪ Loss of a blade server, network access, or data center takes all users offline unless backed by a redundant system.

OS Image Streaming

Summary

Like VHD, OS image streaming promises centralized management of the entire desktop image. This models offers users the responsiveness and performance advantages of local execution, while IT gets the manageability and security benefits of centralization without the larger infrastructure build-out often required by VHD.

Currently, no vendor offers OS image streaming with persistent local caching, so this model cannot be used where mobility is required. Because data resides on the client during execution, it may not be suitable for applications that require the highest data security.

How it Works

At startup, the client is essentially "bare metal," with no OS image installed locally. The OS image is streamed to the client over the network, where it executes locally using the client's own CPU and graphics. Application data is stored in a data center. The client is usually a PC with no hard drive, which uses RAM exclusively.

Figure 5. OS Image Streaming Architecture

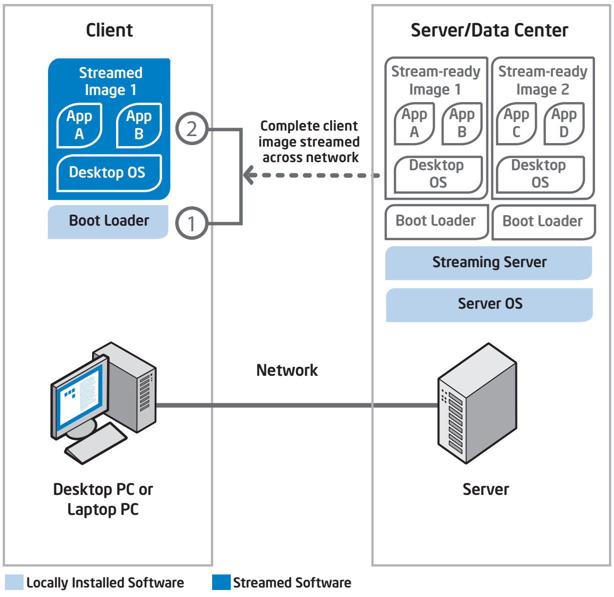

With streaming technology, the OS image software does not stream to the clients in the same form as it comes from the software vendor. The software first goes through a preparation process, where it is divided up into prioritized blocks and placed in a specific order for streaming to the client. The basic launch and initiation software go first, followed by high-demand services and capabilities. These optimizations allow the OS to launch and begin operations, even before all the code is streamed to the client. In order to reduce network traffic, some less frequently used capabilities may remain in the data center until requested.

Table 9. Advantages of OS Image Streaming

Security	▪ Critical application data is stored in the data center. ▪ Any application or OS corruption is eliminated. Patches are automatically applied when the system is reimaged at startup.
Manageability	▪ OS, application, and data management are centralized, allowing simpler administration, easy software migration, and more reliable backup. ▪ OS image streaming significantly reduces the image management and software support issues found with locally installed software. ▪ Software licensing can be centrally managed. Streaming provides better insight into actual application usage, enabling greater licensing optimization. ▪ Stateless clients make adding, moving, and changing users very simple.
Performance	▪ Performance is virtually identical to traditional locally installed applications. Compute- and graphics-intensive applications, as well as video, Adobe Flash,* and streaming media, all perform well. ▪ After initial boot-up, network load is lower than server-based solutions because Remote Display Protocol (RDP) is not used to push application screen images over the network.
Infrastructure Cost Savings	▪ Less server and network infrastructure is required compared to server-based models. ▪ OS image streaming technology has the lowest TCO of all centralized compute models.[1]
Disaster Recovery and Business Continuity	▪ In the event of a data center or worksite disaster, work can shift to other sites relatively easily, assuming redundant servers or a data center is in place.

Table 10. Limitations of OS Image Streaming

Security	▪ Data and applications reside on the client at runtime and are therefore somewhat more susceptible to client-side attacks than server-side models.
Network	▪ Booting the OS requires increase network utilization of 2–5 MBps/user during bootup, but traffic trends to lower levels as local cache is filled. ▪ Streaming download speeds are affected by physical distance from the server, network load, and number of users.
Software Compatibility and Implementation Issues	▪ Due to its internal architecture, some legacy or custom application software cannot successfully complete the preparation process and thus cannot be delivered via OS image streaming. ▪ With OS image streaming, the initial setup and debug of the software preparation process can be time- and labor-intensive. ▪ Recent remote OS boot solutions utilize an unmodified "gold image" and hence do not encounter the limitations of preparation process necessary for OS streaming.
Lack of Mobility	▪ As of this writing, no commercially available product allows off-network or mobile use of OS image streaming.

Remote OS Boot

Summary

Remote OS Boot is a new technology that resembles OS image streaming. Remote OS Boot provides end users with full PC fidelity without requiring back-end server infrastructure, though the solution is managed from a central location. This solution delivers the benefits of desktop virtualization while maintaining full PC and Microsoft Windows* fidelity at a cost that is very attractive to IT and business owners.

How it Works

Remote OS Boot includes three main components: an Intel® Core™ vPro™ processor-enabled client, an Intel®-based storage appliance, and a central administrative console. All management, including connection broker capabilities and component provisioning, are managed from a centralized location.

When a client is powered on, it uses built-in iSCSI boot technology to connect to the remote storage appliance. A Type 1 hypervisor, based on Xen,* is then copied from the remote storage and loaded into the client's memory, and a connection to the central console is established. Based on the client's credentials, the console brokers the connection between the client and the remote storage and mounts the remote storage to the client. From the client, the connection looks like a standard SATA hard drive. The client then boots Windows as it would from a standard hard drive.

The simplicity of the solution enables IT and end users to use standard off-the-shelf or home-grown Windows applications. In addition, standard peripherals such as serial, USB, and printing devices work as they do on a standard PC. These characteristics enable IT departments to quickly deploy new users without having to modify existing tools and technologies.

Remote OS boot is similar to OS image streaming in that it delivers a complete OS image to a "stateless" PC whose hard drive has been deactivated or removed. Unlike OS image streaming, clients boot directly from the storage area network (SAN) device. The PC treats the SAN just like its local hard drive. The OS image is unmodified from the "gold" image that would be used on a local disk. Remote OS boot solutions eliminate the dependence on pre-execution environment (PXE) and PXE servers, further simplifying deployment and lowering TCO.

Figure 6. Remote OS Boot Architecture

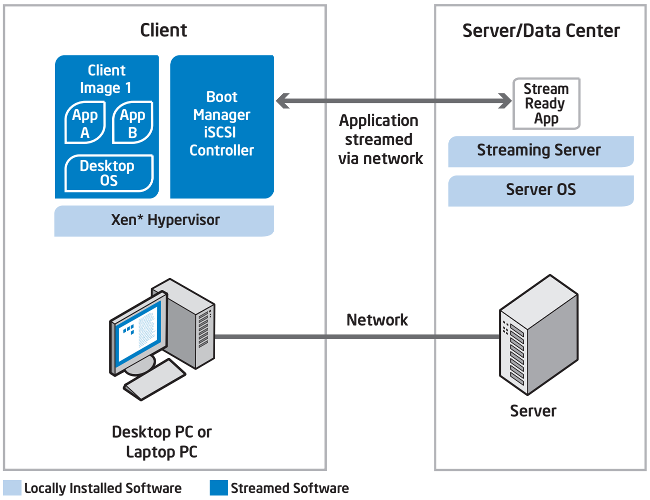

Table 11. Advantages of Remote OS Boot

Security	▪ Critical applications are stored and managed from a central location. ▪ Any application or OS patch can be centrally applied when the client is rebooted. ▪ Windows* standard encryption and authentication devices are supported.
Manageability	▪ Application, data, and end user profiles are centralized for simpler administration and more reliable backups. ▪ Industry-standard IT tools are supported, minimizing the need for specialized tools. ▪ Support for Active Directory* and roaming profiles enables IT organization to use the same processes for PCs, laptops, and secure managed client devices. ▪ Software image management, validation, and support are simplified. Use of standard Windows drivers and dynamic link libraries eliminates conflicts. Adding new images is easy.
Compatibility	▪ Use of unmodified Windows operating system and standard drivers makes Windows applications and standard PC options compatible with this solution.
Performance	▪ Support for VoIP and rich media delivers an end user experience that is similar to or better than a regular PC.
Cost of New User Deployment	▪ Cost of a standard deployment is similar to the cost of a PC. ▪ Adding additional clients is lowest cost among desktop virtualization solutions.
Disaster Recovery and Business Continuity	▪ In a blackout or disaster, access can be shifted to another site. ▪ In the event of local hardware problems, users can easily log onto a new client.
Client Power Consumption	▪ Power consumption and heat are lower than a standard PC. ▪ Total power consumption is lower than other desktop virtualization solutions. ▪ Standard PC power modes are supported as S3 and S4.
Bandwidth	▪ This model is very network-friendly once the client has booted. Network traffic is only required when the client needs to access its virtual hard disk.

Table 12. Limitations of Remote OS Boot

Compatibility	▪ Native Linux* support is currently not offered.
Mobility	▪ Mobile computing is not supported.
Bandwidth	▪ Sufficient network bandwidth is required at boot.

Application Streaming or Application Virtualization

Summary

Like terminal services, application streaming enables centralized application management, but it does so without sacrificing responsiveness or performance. Data can be stored locally or centrally, depending on enterprise policy. Plus, streamed applications can be cached for off-network use, making it a centrally managed model that supports mobile computing.

Because data resides on the client during execution, this model may not be suitable for applications that require the highest levels of security. Application virtualization can also introduce limitations in how applications interact.

How it Works

The client OS is locally installed, but applications are streamed on demand from the server to the client, where they are executed locally.

Although the terms "streaming" and "application virtualization" are often used interchangeably, they are not the same thing. Streaming refers to the delivery model of sending the software over the network for execution on the client. Streamed software can be installed in the client OS locally or (in most cases) virtualized.

With application virtualization, streamed software runs on an abstraction layer and does not install in the OS registry or system files. This simplifies the interactions between the streamed application and the OS, significantly reducing software conflicts and image management problems.

However, some virtualized applications may not interact with other applications as they do when they are both locally installed – for example, cut and paste may not work. This issue can be mitigated by streaming and virtualizing-related applications in bundles, or by using the most recent releases of virtualization software, which address most of these issues.

Unlike other centralized computing models, streamed applications can be cached on a laptop and taken off the network. When the laptop reconnects with the network, the application can resynchronize with the server, check licensing and patch information, and download application data to the data center.

Figure 7. Application Streaming or Application Virtualization Architecture

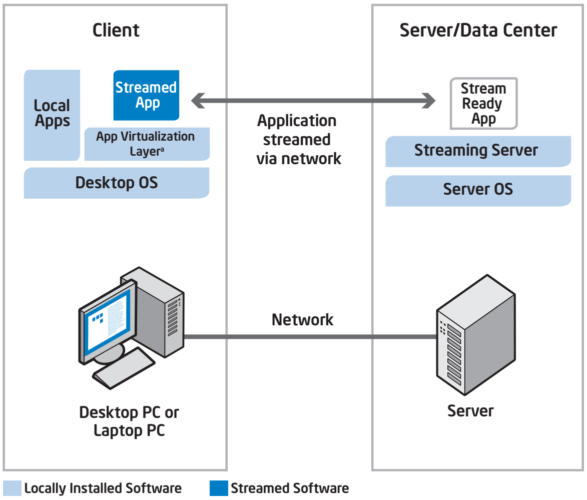

[a]Streaming ISV's agent. Degree of virtualization and isolation varies by vendor and policy.

Table 13. Advantages of Application Streaming or Application Virtualization

Security	▪ Critical application data can be stored in the protected data center based on policy. Local storage can be disallowed. ▪ Application corruption is eliminated. Patches are automatically applied when the application is reloaded at startup. ▪ Virtualized applications can be isolated from each other, limiting data exposure to other applications and the OS.
Manageability	▪ Applications and (in some cases) data are centralized, allowing simpler administration, easier software migration, and more reliable backup. ▪ Application virtualization significantly reduces the image management and software support issues found with locally installed software. ▪ Application virtualization may enable legacy applications to run on a newer OS, even if the application has compatibility problems when locally installed. ▪ Software licensing can be centrally managed. Streaming provides better insight into actual application usage, enabling greater licensing optimization.
Performance	▪ Performance is virtually identical to traditional locally installed applications. Compute- and graphics-intensive applications, as well as video, Adobe Flash,* and streaming media, all perform well. ▪ Caching options can be set to accelerate initial startup and application launch without storing application data locally. ▪ Network load is lower than server-based solutions, because application screen images do not need to be pushed over the network using Remote Display Protocol (RDP) or other protocol. ▪ Streaming only the applications reduces the network load compared to streaming the entire OS image. ▪ User experience is the same as local intelligent client.
Infrastructure Cost Savings	▪ Less server and network infrastructure is required compared to server-based models.
Disaster Recovery and Business Continuity	▪ Users can continue to work on local clients with cached applications even if the network or data center is offline.
Mobility	▪ Unlike other server-based computing models, streamed applications can be cached for off-network use on mobile clients.

Table 14. Limitations of Application Streaming or Application Virtualization

Security	▪ Data and applications reside on the client at runtime and are therefore somewhat more susceptible to client-side attacks or theft than server-based models.
Performance	▪ Streaming download speeds are affected by physical distance from the server, network load, and number of users. ▪ Virtualization may limit interactions between applications (for example, no cut and paste between applications).
Software Compatibility and Implementation Issues	▪ Due to its internal architecture, some legacy or custom application software cannot successfully complete the preparation process and thus cannot be delivered via application streaming. ▪ With application virtualization, the initial setup and debug of the software preparation process can be time- and labor-intensive.
Disaster Recovery and Business Continuity	▪ Although application data is stored centrally, getting users working again in a new site is relatively more complex than other models due to locally installed OS and/or local data.

Virtual Containers

Summary

Virtual containers is a new but rapidly evolving model. It provides centralized management of the complete OS and application image. Unlike a locally installed or streamed OS, the virtual container is abstracted away from the client hardware by a virtual machine manager (VMM). The presence of the VMM eases the validation burden for IT, since various platform hardware differences are hidden from the OS in the container. It also enables new usage models, such as managed contractor images and "bring your own PC" programs. Since client-side execution is used, a large data center build-out is not required. Where server-based execution makes sense, the same virtual containers could be run on a VHD server. For users, virtual containers offer responsive local execution and off-network mobility.

Virtual containers can be general-purpose user environments, such as a standard Windows* OS plus productivity applications; or they can be purpose-built, single-function "virtual appliances" that provide services like compliance monitoring or highly secure applications.

How it Works

In this model, virtual machine images, including the OS and applications, are created and managed centrally by IT. But instead of running the virtual machine on the server (the VHD model), the virtual machine is streamed to the client for local execution on a client-based VMM. Since execution is on the client, even compute- or graphics-intensive applications are responsive, and users can enjoy off-network mobility. There are several technical issues that the industry is striving to solve, such as graphics virtualization, wireless, and power management.

There are two major virtualization approaches offered by software vendors. Advantages related to one approach over the other are identified in the table below.

- Type 1: The virtualization software runs directly on the PC hardware, and the virtual machines or OSs run on top of the hypervisor. Type 1 virtualization software is often referred to as a hypervisor.
- Type 2: The virtualization software runs like an application within an OS such as Windows, and the virtual machines or OSs run on top of the virtualization software. Type 2 virtualization software is often referred to as a virtual machine manager (VMM).

Figure 8. Virtual Containers Architecture

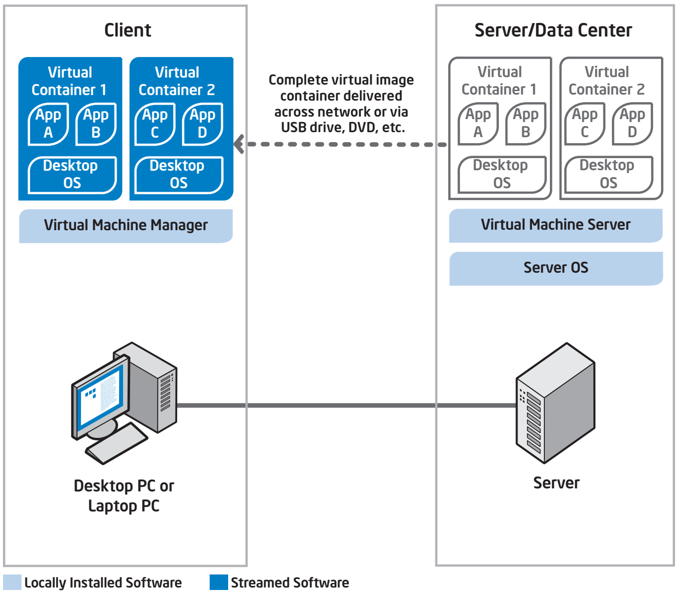

Table 15. Advantages of Virtual Containers

Security	▪ With virtual containers, virus and intrusion threats are contained and applications requiring high security can be easily isolated in separate containers. ▪ Security policies can be set for the needs of the virtual container, allowing IT more control over lock-down, access policy, data storage, and authorization/revocation rights. ▪ Purpose-built virtual appliances can provide valuable security services from outside the user's environment. ▪ Type 1 vs. Type 2: Type 1 is considered more secure because of isolation between virtual machines. With Type 2, if the underlying OS is compromised, the virtual machines would be too.
Manageability	▪ OS images, application, and data management can be centralized, allowing simpler administration, easy software migration, and more reliable backup. ▪ Since virtual containers run on a VMM, not on physical hardware, image validation only needs to be done against the VMM rather than many unique hardware configurations. ▪ Virtual containers are highly portable. Installation may be as simple as streaming a file to the client or inserting a USB flash drive. ▪ Type 1 vs. Type 2: Type 2 can run all native Windows* device drivers and is considered less complex and more proven on PCs than Type 1, which often requires integration of some non-Windows drivers.
Performance	▪ Performance is virtually identical to traditional locally installed applications. Compute- and graphics-intensive applications, as well as video, Adobe Flash,* and streamed media, all perform well. ▪ Type 1 vs. Type 2: Type 1 is considered higher performance because a hypervisor is a smaller code base with less overhead, compared to the OS and VMM required in Type 2.
Mobility	▪ Unlike other server-based computing models, virtual containers can be cached for off-network use on mobile clients. ▪ Users could carry their virtual container image on a USB drive and run it on any PC (home, office, partner site, etc.).
Infrastructure Cost Savings	▪ Less server and network infrastructure is required compared to server-based models.
Disaster Recovery and Business Continuity	▪ Users can continue to work on local clients using cached virtual containers even if the network or data center is offline.

Table 16. Limitations of Virtual Containers

Security	▪ At runtime, data and applications reside on the client and are therefore more susceptible to client-side attacks or theft than server-side models. ▪ VMM is a new layer in the software stack that must also be protected.
Performance	▪ Running multiple virtual machines on a VMM may cause performance degradation. VMM efficiency and hardware-assistance are key to achieving "near native" application performance. ▪ Virtualization may limit interactions between applications. This may be by design in some cases but an unintended consequence in others.
Maturity	▪ The virtual containers model is relatively new. The technology, management tools, and IT processes are not fully mature.
Industry-Wide Technical Challenges	▪ The industry is striving to solve several technical challenges before the full potential of virtual containers can be realized. This includes virtualization of graphics, wireless, power management, docking stations, and peripherals.

Conclusion

As software delivery models have evolved, IT departments have more choices than ever. Choosing the appropriate desktop virtualization model requires a balance of many factors. Finding a solution that meets all relevant user and IT needs is likely to involve a mix of delivery models, even within a single-user segment.

PCs powered by the Intel Core vPro processor family can help businesses cut costs and increase efficiency by taking advantage of intelligent performance and unique hardware-assisted security and manageability features. They offer highly manageable, secure platforms from which to deliver all desktop virtualization models.

Additional Resources

Principled Technologies offers free white papers, which are a great source of third-party quantified data to support the points in this document. They are available on Principled Technologies' web site: www.principledtechnologies.com.

- White Paper: Total Cost of Ownership for Various Computing Models: www.principledtechnologies.com/Clients/Reports/Intel/CompModelsTCO1107.pdf
- Spreadsheet: Total Cost of Ownership Calculator: www.principledtechnologies.com/Clients/Reports/Intel/ComputeModelTCOCalc1107.xls

[1] Principled Technologies, “Total Cost of Ownership for Various Computing Models.”

Printed in USA 0610/WW/OCG/XX/PDF ♻ Please Recycle 319175-005US