White Paper

Intel® Centrino® with vPro™ Technology

Intel® Core™2 Processor with vPro™ Technology

Understanding Alternative Compute Models

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

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

Rarely is a single compute 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 optimal solution will likely involve a combination of models.

Although thin terminals may be suitable hardware for certain purely server-based models, rich PC platforms based on technology such as Intel® Core™2 processor with vPro™ technology for desktops and Intel® Centrino® with vPro™ technology for notebooks offer highly manageable and secure platforms on which to deploy a wide range of solutions.

Table of Contents

Issues to Consider

Here are some considerations to keep in mind in evaluating centralized models for your organization.

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 install locally.

It's helpful to segment users based on tasks performed and applications required. Individual users may well 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 compute 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.

Table 1. Compute Model Comparison

	Terminal Services	Virtual Hosted Desktop	Blade PCs	OS and Application Streaming	Application Streaming or Application Virtualization	Rich Distributed Computing
Application Execution	Server	Server	Server	Client	Client	Client
Application Data Storage	Server	Server	Server	Server	Client or server	Client or server
Local Device Connect and Synch (Bar Code Reader, PDA, Phone, etc.)	No	Limited	Yes	Yes	Yes	Yes
Mobility (Off-Network Operation)	No	No	No	No	Yes, with client caching option	Yes
Clients	Terminal, desktop PC, laptop PC	Terminal, desktop PC, laptop PC	Terminal, desktop PC, laptop PC	Desktop PC, laptop PC	Desktop PC, laptop PC	Desktop PC, laptop PC
Major Vendors	Citrix, Microsoft	VMware, Citrix	HP, ClearCube	Citrix	Microsoft, Citrix, Symantec, AppStream	Traditional application and management software providers

Terminal Services

Summary

Terminal services are 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 are 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. 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

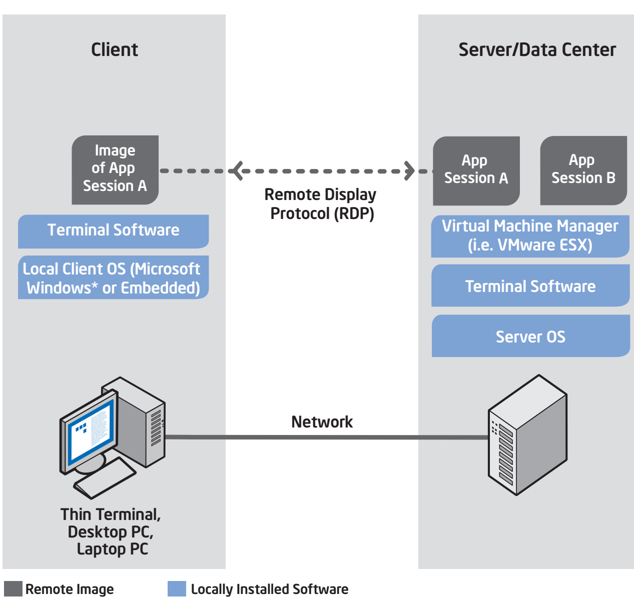

Table 2. Advantages of Terminal Services

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 are a mature and well-understood technology. Implementation is fairly straightforward. ▪ 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.
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.
Reduced Client Power Consumption	▪ With thin client terminals, client power consumption is lower than most desktop PCs. However, client power savings may be offset by greater storage requirements and power consumption in the data center.

Table 3. 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, Adobe Flash,* or animated GIFs, do not perform effectively over RDP. ▪ 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 virtual hosted desktop, 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 high user dissatisfaction and complaints.

Virtual Hosted Desktop

Summary

Virtual Hosted Desktop (VHD) is a new model that is attracting a great deal of industry buzz, although it's not yet widely implemented or well understood.

VHD is designed to offer the responsive and customizable user experience of rich 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. Video, Adobe Flash,* Voice over Internet Protocol (VoIP), and other compute- or graphics-rich applications are not well suited to this model. Also, VHD requires a persistent network connection, so it's not appropriate where mobility is required.

New deployments of VHD to large numbers of users may be costly. 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. IT departments need to weigh these infrastructure costs against potential TCO savings in manageability.

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). The major difference is that VHD offers each user their own complete virtual machine and customized desktop, including the OS, applications, and settings.

Figure 2. Virtual Hosted Desktop Architecture

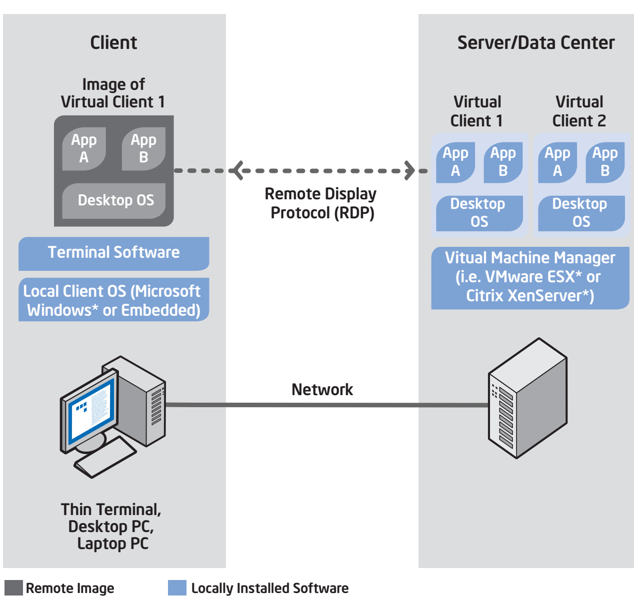

Table 4. Advantages of Virtual Hosted Desktop

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.
Reduced Client Power Consumption	▪ With thin client terminals, client power consumption is lower than most desktop PCs. However, client power savings may be offset by greater storage requirements and power consumption in the data center.

Table 5. Limitations of Virtual Hosted Desktop

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 Virtual Hosted Desktop (VHD) users. ▪ Remote Display Protocol (RDP) creates a serious graphics bottleneck. Video, Adobe Flash,* and animated GIFs do not perform effectively over RDP. ▪ Network bandwidth and loading are major factors in client system performance. Networks must have sufficient peak capacity to support the number of VHD users. ▪ Users expect rapid response to mouse clicks and keystrokes. Terminal servers must be physically located close enough to their clients to meet responsiveness requirements.
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 Compatibility	▪ Software such as VoIP, streaming media, and compute- or graphics-intensive applications may be poorly suited to server-based computing.
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.
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 high user dissatisfaction and complaints.

Blade PCs

Summary

Blade PCs bring client computing into a central location, promising higher manageability and security than rich 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 rich 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 compute models may require a complete "forklift upgrade." In addition, one-to-many blade PCs may demonstrate some of the same draw-backs as VHD: performance and graphics issues, as well as high infrastructure costs. Be sure to weigh these factors in any ROI analysis.

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 is 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

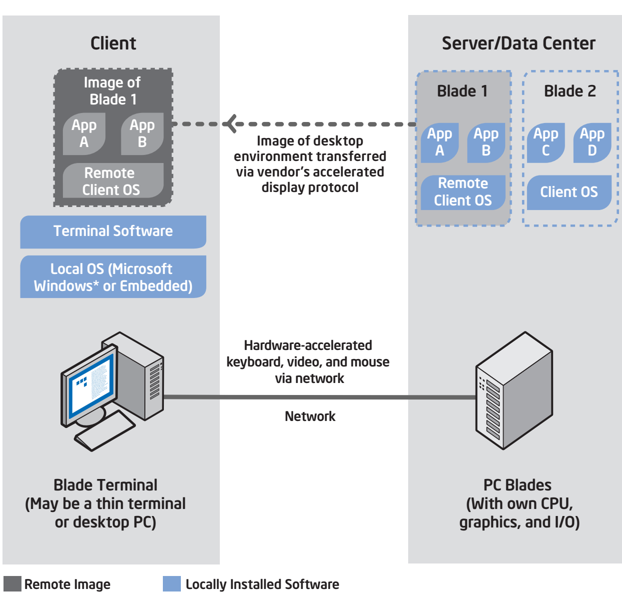

Figure 4. One-to-Many Blade Architecture

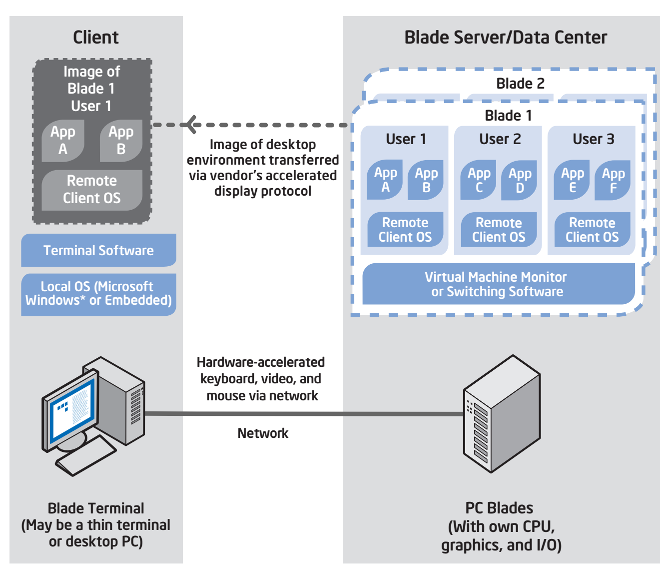

Table 6. Advantages of Blade PCs

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 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 data center 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 7. Limitations of Blade PCs

Performance	▪ In 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 compute models requires a complete hardware upgrade in most cases. ▪ Available management tools may be limited, and IT may be dependent on their 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 and Application Streaming

Summary

Like VHD, streaming OS and applications promises centralized management of the entire desktop image. This model 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 an OS streaming product with 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 or applications installed. The OS and applications are streamed to the client over the network, where they execute locally using the client's own CPU and graphics. Application data is stored in a data center. The client can even be a PC with no hard drive, which uses RAM exclusively.

With streaming technology, the OS and application software does not stream to the clients in the same form as it comes from the software vendor. The software first goes through a process called sequencing, 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. Sequencing allows the application 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.

A small minority of applications cannot successfully complete the sequencing process because of the way are designed or linked to other software. In these cases, another means of delivery must be used.

Figure 5. OS and Application Streaming Architecture

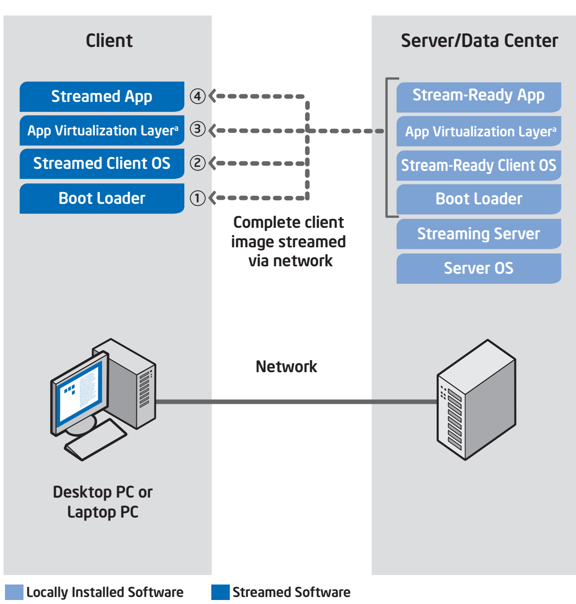

[a]Streaming ISV's agent. Degree of virtualization and isolation varies by vendor and policy.

Table 8. Advantages of OS and Application Streaming

Security	▪ Critical application data is stored in the data center. ▪ Any application or OS corruption is eliminated, and patches are automatically applied when the system is reimaged at startup. ▪ Applications can be isolated in protected virtual "containers," limiting data exposure to other applications or the OS.
Manageability	▪ OS, application, and data management are centralized, allowing simpler administration, easy software migration, and more reliable backup. ▪ OS and application streaming significantly reduces the image management and storage challenges found with VHD, as well as the desktop software support issues found with local rich client. ▪ 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. ▪ Caching options can be set to accelerate initial startup without local storage of data. ▪ After initial boot-up, network loading takes less time 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 are required compared to server-based models, with significantly lower network storage requirements than VHD. ▪ Streaming technology has the lowest initial deployment costs 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. ▪ Users can continue to work on local clients even if the data center or network is offline.

Table 9. Limitations of OS and Application Streaming

Security	▪ Data and applications reside on the client at runtime and are therefore more susceptible to client-side attacks than server-side models.
Performance	▪ Booting a streamed OS over the network can be somewhat slower than booting locally. ▪ 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 software cannot successfully complete the sequencing process and thus cannot be streamed. ▪ The initial setup and debug of the software sequencing process can be time- and labor-intensive.
Lack of Mobility	▪ As of this writing, no commercially available product enables off-network or mobile use of a streamed OS.

Application Streaming

Summary

Like terminal services, application streaming enables centralized application management, but 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 this the only centralized 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 data center 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 it can be 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, virtually eliminating 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, but it may be difficult to predict all combinations users will require.

Unlike other centralized computing methods, 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 6. Application Streaming Architecture

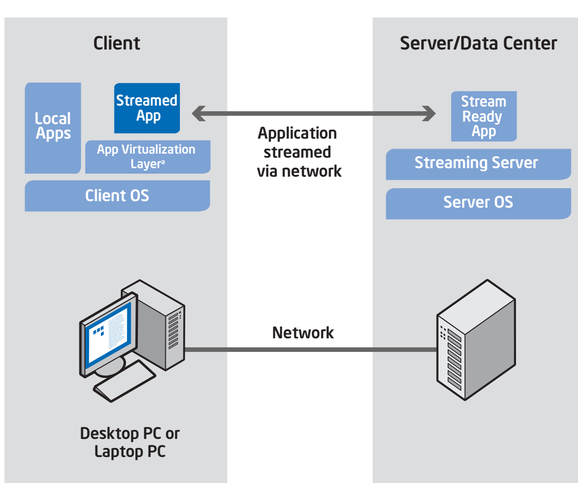

[a]Streaming ISV's agent. Degree of virtualization and isolation varies by vendor and policy.

Table 10. Advantages of Application Streaming

Security	▪ Critical application data is stored in the protected data center. Application corruption is eliminated, and patches are automatically applied when the application is reloaded at startup. ▪ Virtualized applications can be isolated in protected virtual "containers," limiting data exposure to other applications or the OS.
Manageability	▪ Applications and data are centralized, allowing simpler administration, easier software migration, and more reliable backup. ▪ Application streaming significantly reduces the image management and storage challenges of VHD, as well as the desktop software support issues found with local rich client. ▪ 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 loading takes less time than server-based solutions, because Remote Display Protocol (RDP) is not used to push application screen images over the network. ▪ Streaming only the applications reduces the network load compared to streaming both the OS and applications. ▪ User experience for OS boot is the same as local rich client.
Infrastructure Cost Savings	▪ Less server and network infrastructure are required compared to server-based models, with significantly lower network storage requirements than VHD. ▪ Streaming technology has the lowest initial deployment costs of the centralized compute models[1].
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 methods, streamed applications can be cached for off-network use on mobile clients.

Table 11. Limitations of Application Streaming

Security	▪ Data and applications reside on the client at runtime and are therefore 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 software cannot successfully complete the sequencing process and thus cannot be streamed. ▪ The initial setup and debug of the software sequencing 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.

Rich Distributed Computing

In this model, applications are stored and executed on client PCs or laptops. This model is the de facto standard in the industry today and well understood and will not be discussed in detail here.

Its main advantages are versatility, responsiveness, mobility, and availability of a wide range of software and management tools. On the other hand, as is well known, it requires the most work to manage and secure.

Factors to Consider

Choosing the appropriate compute 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 users can't stand.

IT Requirements Standard IT concerns, such as security, manageability, business continuity, etc.	▪ Security ▪ Image management ▪ License management ▪ Support structure ▪ Disaster recovery
Infrastructure The hardware, connectivity, and bandwidth available to deliver the compute model. If the infrastructure is not in place to support a particular compute 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 compute 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

User Scenarios

The following scenarios demonstrate how these factors play out for three different types of users.

Scenario: Call Center Worker

Tasks: Take customer orders and inquiries on the phone and enter them in a central database.

IT Requirements	Infrastructure	User Experience	Application Workloads
▪ Protect database ▪ No unauthorized software ▪ Onsite support ▪ No service interruptions	▪ Data center not constrained by space, power, cooling, or budget ▪ Gigabit network ▪ Full redundant systems	▪ Fixed location ▪ Fast database response ▪ Little customization ▪ No personal information or applications on client	▪ Not compute- or graphics-intensive ▪ Applications tolerant of minor network delays ▪ Workload not expected to change in future

Scenario: Construction Estimator

Tasks: Visit project sites, schedule jobs, analyze data and drawings for estimates, and present findings to customers and management.

IT Requirements	Infrastructure	User Experience	Application Workloads
▪ Protect network/database ▪ Protect confidential data ▪ "Anywhere" access ▪ Serviceable in the field	▪ Data center not constrained by space, power, cooling or budget ▪ Wired and cellular network ▪ Variable bandwidth	▪ Highly mobile ▪ Fast response required ▪ Heavy customization ▪ Some personal information or applications on PC	▪ Occasionally compute- and graphics-heavy ▪ Workload may change in future

Scenario: University Students

Tasks: Access a standard department image on lab workstations or from remote locations, use course-specific applications, create their own documents, and use computer for personal tasks and entertainment.

IT Requirements	Infrastructure	User Experience	Application Workloads
▪ Standard image delivery ▪ Malware protection ▪ Licensing compliance ▪ Off-campus access	▪ Data center expansion constrained by space and budget ▪ Gigabit on-campus wired network ▪ Variable wireless and off-campus bandwidth	▪ Highly mobile ▪ Responsive applications ▪ Heavy customization ▪ Personal info and applications on PC	▪ Frequently compute- and graphics-intensive ▪ Video and graphics applications intolerant of network delays ▪ Workloads may change with coursework

Comparison of Solutions

For all but the call center representatives, a mix of compute models works best to deliver the user requirements, provide the management and security IT needs, and work within the infrastructure constraints.

When a mix of server-based and client-based models is needed, a PC is the best device for the job. Thin terminals are only appropriate in cases when a 100% server-side model makes sense.

	Call Center Representatives	Construction Estimators	University Students
Terminal Services	Order database	Scheduling database	
Virtual Hosted Desktop			Remote access to standard department image
Application Streaming		CAD package	Course-specific applications
Local Installation		Office applications, browser	Office applications, browser, media players
Appropriate Client Platform	Desktop PC or thin terminal	Laptop PC	Laptop or desktop PC

Conclusion

As software delivery models have evolved, IT departments have more choices than ever. Choosing the appropriate compute 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.

Desktop PCs with the Intel Core 2 processor with vPro technology and Intel Centrino Pro with vPro technology for notebooks offer highly manageable, secure platforms on which to deliver all compute 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." See Additional Resources for reference.

*Other names and brands may be claimed as the property of others.

Printed in USA 0108/WW/OCG/XX/PDF Please Recycle 319175-001US