



Total cost of ownership for various computing models

Executive summary

Information technology (IT) managers can choose from a wide variety of computing models that each tout benefits beyond those of traditional rich clients. Proponents of these models claim lower power utilization, greater ease of manageability, better security, and lower cost of ownership than traditional clients. IT must evaluate these factors and weigh the pluses and minuses of each of the models across their user base.

Intel Corporation (Intel) commissioned Principled Technologies (PT) to evaluate the costs and benefits of six computing models including typically managed desktops. Figure 1 lists the models.

Computing models we evaluated

- Server-side computing models
- Terminal/Presentation server
 - Virtual hosted desktop
- Client-side computing models
- Typically managed rich desktop
 - Well-managed application streaming desktop using Intel vPro technology
 - Well-managed operating system (OS) streaming desktop using Intel vPro technology
- Other computing models
- Blade PC desktop computing model (a hybrid server-side approach where each user session has a dedicated PC blade)

Figure 1: Computing models we evaluated. The Computing models section describes these in more detail.

We measured and analyzed the competing models on a wide variety of characteristics including performance, power utilization, and capabilities. We also conducted a somewhat more subjective but still fact-based analysis of such factors as desktop environment and future-proofing and estimated the total cost of ownership (TCO) of the models. Figure 2 presents our key findings for each of the categories we analyzed.

Key findings		
TCO	Deployment	The Blade PC desktop and server-side computing models have significantly higher deployment costs than the rich client computing models.
	Power	The Terminal/Presentation server model has the lowest power costs. Power costs for all models however, account for only a small fraction of the overall TCO.
	Manageability	Manageability costs are lowest with server-side and Blade PC desktop computing models.
	Productivity/ User experience	The shared nature of server-side platforms and the slow nature of Blade PCs hinder user experience and productivity, particularly in the case of knowledge or power users. Combining well-managed rich clients with application streaming and/or OS streaming can provide all the benefits of server-side computing models without significantly affecting productivity. Because lost user productivity can easily be higher than any of the other costs in this analysis, enterprises need to consider those costs carefully.
	Total cost	TCO is the sum of acquisition costs and sustaining costs. TCO for server-side and Blade PC desktop computing models is higher than for client-side computing models primarily due to deployment costs and productivity losses.
	Security	All other platforms we examined offer considerable security improvements over Typically managed rich desktops.
	Future proofing/ rich application support	Client-side computing models based on rich desktops and notebooks offer significant future-proofing benefits over server-side and Blade PC desktop models. Notably, server-side models offer limited multimedia and rich collaboration support as well as limited Flash-based Internet usage.
	Desk-side environmental	Smaller and usually cooler and quieter than rich clients, thin clients used in server-side and Blade PC desktop models have less of an impact on the desk-side environment.
	Compliance	All other computing models we examined offer significant compliance benefits over Typically managed rich desktops.
	Performance	Client-side computing models showed impressive performance gains over server-side and Blade PC desktop models in our tests. Productivity loss can be significant for knowledge or power users on server-side and Blade PC desktop models.
	Mobility	True mobility support is available only with the application streaming computing model.

Figure 2: Our key findings for each of the categories we analyzed.

The first five analysis categories relate to total cost of ownership (TCO). The remaining categories are subjective or objective analysis categories that influence model selection. The Analysis categories section later in this report describes these in more detail.

We found that server-side models may be an appropriate solution for task workers or in places where security or centralized management requirements vastly dominate other factors. However, productivity and mobility considerations can quickly outweigh these issues where knowledge or power users are concerned. Well-managed rich clients supported by third-party manageability software, provide greater benefit for lower costs. The additional management and security capabilities of Intel vPro technology extend that advantage. On top of that, combining well-managed rich clients with application streaming and/or OS streaming can provide the benefits of server-side computing models without significant loss of end-user productivity and result in the lowest cost of ownership. Figure 3 summarizes our model comparison.

		Terminal/ Presentation server	Typically managed rich desktop	Virtual hosted desktop ¹	Well-managed OS streaming/ vPro	Blade PC desktop	Well-managed application streaming/vPro
T C O	Deployment (\$ = more cost = bad)	\$\$\$\$	\$	\$\$\$\$	\$\$\$	\$\$\$\$\$\$	\$\$\$
	Power (\$ = more cost = bad)	\$	\$\$\$	\$\$	\$\$\$	\$\$\$\$	\$\$\$
	Manageability (\$ = more cost = bad)	\$\$	\$\$\$\$\$\$	\$\$	\$\$	\$	\$\$\$
	Productivity/ user experience (\$ = more cost = bad)	\$\$\$	-	\$\$\$	\$	\$\$\$\$\$\$	-
	Total cost (\$ = more cost = bad)	\$\$\$\$	\$\$\$	\$\$\$\$	\$\$	\$\$\$\$\$\$	\$\$
Security (more * = more secure = good)		****	*	****	*****	***	****
Future proofing/ rich application support (more * = more flexibility = good)		**	****	***	****	*	*****
Desk-side environmental (more * = more comfort = good)		****	**	****	**	****	**
Compliance (more * = more compliant = good)		*****	*	****	****	****	****
Performance (more * = good)		***	*****	Not tested	****	***	*****
Mobility (more * = good)		*	****	*	*	*	*****
Platforms we tested		Citrix Presentation Server to thin client	Typically managed rich client	VMware- based VDI to thin client	Ardence to Intel vPro/rich client	HP CCI to thin client	AppStream to Intel vPro/ Centrino Pro rich client

Figure 3: Compute model TCO and benefits comparison for the six computing models we examined.

¹We estimated productivity and deployment costs for the Virtual hosted desktop model.

We assigned one or more dollar signs (\$) for TCO categories and one or more stars (*) for the other analysis categories. For TCO categories, we estimated the dollar signs using a relative cost per dollar sign, with more dollar signs indicating higher cost. For the other analysis categories, we assigned a relative estimate using one to five stars, with more stars indicating better quality. We based our calculations of productivity/user experience costs, performance, and access infrastructure number of clients per server for all models except the Virtual hosted desktop model on performance tests we conducted. We did not test the Virtual hosted desktop model's performance. We used outside sources to estimate the Access infrastructure and lost productivity for this model.

We assumed the cost and performance of common application and resource servers, such as those providing file, email, database, network services (DNS, Active Directory), and Web services would be the same across all platforms and thus excluded those costs from our analysis.

Computing models

We analyzed six computing models. Our models included a range of computing options. Three of these models included thin clients (Terminal/Presentation server, Virtual hosted desktop, and Blade PC desktop) and three used rich clients. Two models used server-side computing (Terminal/Presentation server and Virtual hosted desktop), three used client-based computing, and one (Blade PC desktop) used a hybrid server-side approach where each user session has a dedicated PC blade. For all models except for the Typically managed rich desktop model, we assumed a well-managed infrastructure where IT uses advanced automation tools to support and manage clients.

Terminal/Presentation server computing model	
<p>The Terminal/Presentation server model is the traditional thin client, server-side computing model. The client device does little more than accept keystrokes and mouse clicks for input and render the response from the server to a display. Many of its strengths derive from the client device's lack of features. For example, there is no local storage to hold an unauthorized copy of an application or non-compliant data. We assume this model exists, as it typically does, within a well-managed IT infrastructure.</p>	<p>Clients: Our test bed for Terminal/Presentation server Included 10 Wyse Winterm V50 thin clients.</p> <p>Access server: Our server (two 3GHz Intel Xeon 5160 dual-core processors, 16GB DDR2 SRAM, a Supermicro X7DBE+ motherboard, and four 73GB Seagate Cheetah ST373455SS disks) used Citrix Presentation Server 4.5 to administer clients.</p>
Typically managed rich desktop computing model	
<p>The Typically managed rich desktop model is the most common model in use. Each user has a full-featured PC, gaining maximum power and flexibility and, to a degree, maximum risk as well. IT policies require considerable manual effort to implement and maintain. IT has not automated all management functions, such as patch management and asset inventories. Factors such as machines being powered off limit the effectiveness of automated functions.</p>	<p>Clients: Our test bed for Typically managed rich desktop included 10 hand-built rich clients (Intel Desktop Board DQ965GF, 2.13 GHz Intel Core 2 Duo E6400, 1GB DDR2 SRAM running Windows XP Professional version 2002 with Service Pack 2).</p> <p>Access server: None</p>
Virtual hosted desktop computing model	
<p>The Virtual hosted desktop is a hybrid that attempts to gain some of the advantages of both the Terminal/Presentation server and the Typically managed rich desktop model. User processing occurs on the server, as in the Terminal/Presentation server model, but each user runs an independent session in a virtual system, allowing users to see different operating systems and different versions of applications. We assume this model exists within a well-managed IT infrastructure.</p>	<p>Clients: We did not test this model, but instead estimated performance for an implementation of 10 Wyse Winterm V50 thin clients to estimate the costs.</p> <p>Access server: Our hypothetical server (two 3GHz Intel Xeon 5160 dual-core processors, 16GB DDR2 SRAM, a Supermicro X7DBE+ motherboard, and four 73GB Seagate Cheetah ST373455SS disks) administered clients using VMware: ESX Server 3.0.1.</p>

Well-managed OS streaming desktop/vPro	
<p>OS streaming is a variation on the rich client model. At boot time, the client boots from a virtual disk on the server. Thus, the storage and OS image are on the server, but the actual work happens on the client. Although it is possible to use OS streaming on clients that also have a local disk, here we only consider the case of diskless clients. We assume this model exists within a well-managed infrastructure.</p>	<p>Clients: Our test bed for OS streaming desktop/vPro included 10 hand-built rich clients with Intel vPro technology (Intel Desktop Board DQ965GF, 2.13 GHz Intel Core 2 Duo E6400, 1GB DDR2 SRAM running Windows XP Professional version 2002 with Service Pack 2 streamed from the server) and Ardence Evaluation Client 4.1.</p> <p>Access server: Our server (two 3GHz Intel Xeon 5160 dual-core processors, 16GB DDR2 SRAM, a Supermicro X7DBE+ motherboard, and four 73GB Seagate Cheetah ST373455SS disks) administered clients using Ardence Evaluation Server 4.1 OS.</p>
Blade PC desktop	
<p>The Blade PC desktop model is a hybrid that attempts to gain some of the advantages of both the Terminal/Presentation server and the rich desktop. The users have thin client devices at the desktop. However, each user runs an independent session on a blade PC. As with Virtual hosted desktops, users can see different operating systems and different versions of applications. However, in the case of the Blade PC desktop model, performance is a function of the speed of the blade PC, not the server. We assume this model exists, as it typically does, within a well-managed infrastructure.</p>	<p>Clients: Our test bed for Blade PC desktop included 10 Wyse Winterm V50 thin clients.</p> <p>Blade PCs: HP BladeSystem bc2000 Blade PCs (AMD Athlon 64 2100+ 1.20GHz, 1GB DDR2 RAM running Windows XP Professional version 2002 with Service Pack 2).</p> <p>Access server: None</p>
Well-managed application streaming desktop/vPro	
<p>Well-managed application streaming desktop/vPro is a rich client model. The client has local storage where the operating system is located. However, the applications live on a server that streams them to the client as needed. The work happens on the client, and users can store data locally as well. We assume this model exists within a well-managed infrastructure.</p>	<p>Clients: Our test bed for Well-managed application streaming desktop/vPro included hand-built rich clients (Intel Desktop Board DQ965GF, 2.13 GHz Intel Core 2 Duo E6400, 1GB DDR2 SRAM running Windows XP Professional version 2002 with Service Pack 2) running AppStream Technology Windows Edition 5.2.1 client software.</p> <p>Access server: Our server (two 3GHz Intel Xeon 5160 dual-core processors, 16GB DDR2 SRAM, a Supermicro X7DBE+ motherboard, and four 73GB Seagate Cheetah ST373455SS disks) streamed applications using AppStream Technology Windows Edition 5.2.1 server software.</p>

Figure 4: Test beds for the six computing models we examined.

We looked at TCO as well as qualitative and quantitative categories. We created and ran performance tests on test beds for all but the Virtual hosted desktop model and used those results to develop estimates for the Productivity/user experience and Performance categories. Figure 4 identifies the models and the test beds.

All thin clients were Wyse Winterm V50. All rich clients were Intel DQ965GF vPro systems. In our performance tests, all of the rich clients benefited from the Core 2 Duo processor's performance. We considered the Intel vPro-specific capabilities, such as those related to management and security, of the DQ965GF system only in the Well-managed/vPro cases. For instance, Active Management Technology (AMT) on PCs powered by Intel vPro

technology would allow out-of-band desktop management, Trusted Platform Module (TPM) support would allow the devices to take advantage of third-party encryption tools, and the Intel Stable Image Platform Program (SIPP) would help standardize desktop and notebook components and drivers, lowering TCO and simplifying system updates.

Each of our test networks included a file server and 10 client systems. We used a pair of identical file servers to allow us to test two networks at a time. We used a 100Mbps network infrastructure for the clients, and a 1Gbps network infrastructure for the servers to reflect typical scenarios in enterprises today. We also set up access infrastructure servers for the models that required them.

Example enterprise

Our example enterprise is evaluating a change from its current typically managed rich computing model. In this model, IT has automated only some management functions—such as patch management and asset inventories—and factors such as machines being powered off fundamentally limit the effectiveness of automated functions. IT wants to conduct a TCO analysis of various models to compare deployment costs, manageability costs (including user downtime), power costs, and possibly costs of lost user productivity due to platform slowness. Factors such as worker location and type and platform refresh cycle among others influence the analysis and results. Different enterprises would of course see different results. We list here our main assumptions about factors that influence TCO for the example enterprise used in this analysis:

- Enterprise is considering whether to change from current typically managed rich computing model for 10,000 users.
- All users are in one location or campus.
- Only one user uses each client.
- All workers are knowledge workers.
- All clients are at the end of their current refresh cycle and due to be replaced.
- IT anticipates a four-year refresh cycle for the selected computing model.
- All users require a desktop solution, although the enterprise wants a solution that can expand to include mobile users as well.
- The 10,000 users run 120 applications including standard office applications. The users are well trained in the current applications, and IT hopes to transfer all applications to the selected computing model.
- Average burdened worker hourly rate is \$40.82; average burdened IT hourly rate is \$63.27.
- The enterprise is in the United States and values represent US dollars.
- TCO analysis should consider deployment costs, manageability costs including user downtime, power costs, and possibly costs of lost user productivity due to platform slowness.
- Servers and Blade PCs (for Blade PC Desktop solution) and support staff are in single location separate from the 10,000 users' location.
- IT wants to retain the current client and server operating systems.

Analysis categories

Details of our testing of the various platforms is available.

Performance

We discuss performance first because we include those results in calculations for other categories. To compare model performance, we ran three different applications scenarios on each type of client:

- **Acrobat compress.** This single-task scenario tested how quickly the test system was able to open Acrobat and compress a 4.01MB PDF file (located on the file server) from within the Acrobat application.
- **Excel subtotals.** This single-task scenario tested how quickly Excel could perform the subtotal function on a 1.79MB Excel worksheet (located on the file server).
- **Explorer compress and PPT change view.** This multitasking scenario tested how long it took to compress a 265MB folder (located on the file server for the thin clients; local for the rich clients) while changing views within a 30.4MB Microsoft PowerPoint presentation (located on the file server).

We ran the three test scripts on each test network with four client configurations: a single client running the script and then 2, 5, and 10 clients simultaneously running the script.

Figure 5 shows the performance comparison for the computing models. We omit the Virtual hosted desktop model because we did not test it in our lab.

The rich desktop-based computing models showed the best performance in our tests. Compute-intensive and graphics-intensive tasks on server-side and Blade PC desktop computing models result in poor end-user experience and lost productivity.




Figure 5: Performance comparison – the sum of script execution times for the three test scenarios in seconds. Lower results are better. We did not test the Virtual hosted desktop model.

In addition to using the test results to evaluate performance, we also used the results of the 10-client test to estimate lost productivity and clients per access infrastructure server. To get our clients-per-server estimate for the Virtual hosted desktop model, we averaged the values from VMware Infrastructure 3 VDI Server Sizing and Scaling (http://www.vmware.com/pdf/vdi_sizing_vi3.pdf). Based on experience, we estimated the lost productivity of the Virtual hosted desktop model to be between 2 and 5 minutes per user per day; we used the average, 3.5 minutes per day, for our calculations.

Total cost of ownership analysis

We calculated TCO for an example enterprise considering converting 10,000 clients to a new platform. We assumed a ratio of one user to one client system and a four-year refresh cycle. Figure 6 shows our estimate of the annual TCO for six computing models for this enterprise, including costs of deployment, power, and manageability, as well as the cost of lost user productivity due to access server congestion or Blade PC slowness. We also looked at TCO scaling, analyzing costs for different sized enterprises. We saw no significant cost differences in TCO within any model at enterprise-level client counts.

Client-side computing models showed impressive performance gains over server-side and Blade PC desktop models in our tests. Productivity loss can be significant for knowledge or power users on server-side and Blade PC desktop models.

Models	Annual TCO with lost productivity	Annual TCO without lost productivity
Terminal/Presentation server	\$1,370	\$826
Typically managed rich desktop	\$1,025	\$1,025
Virtual hosted desktop	\$1,419	\$836
Well-managed OS streaming/vPro	\$898	\$771
Blade PC desktop	\$2,254	\$975
Well-managed application streaming/vPro	\$774	\$774

Figure 6: Annual per client TCO with and without lost user productivity costs for 10,000 clients in an example enterprise.

Deployment costs

Deployment costs are the one-time costs needed to configure the data center and workspaces to use the client type in question. We looked at the following costs:

- Per-client seat costs, including client hardware and licenses
- Server costs, including racks, hardware, and software
- Access infrastructure costs, including additional storage and management software
- Physical costs, such as additional wiring
- Implementation and planning costs for both internal staff and consultants
- Training for both end users and IT staff
- The cost of porting or replacing applications

The Blade PC desktop and server-side computing models have significantly higher deployment costs than the rich client computing models.

As part of this analysis, we used the performance test results as well as processor utilization results we collected during the tests to calculate the number of clients per server for servers used to host the desktop and running applications. Client-side models and the Blade PC desktop model don't use these servers because they perform these functions either locally or on a dedicated blade. The number of these servers for other models ranged from a high of one per 36 clients for the Virtual hosted desktop model down to one per 399 clients for Well-managed application streaming/vPro.

Figure 7 shows our estimate of the deployment costs for our example 10,000-client conversion.

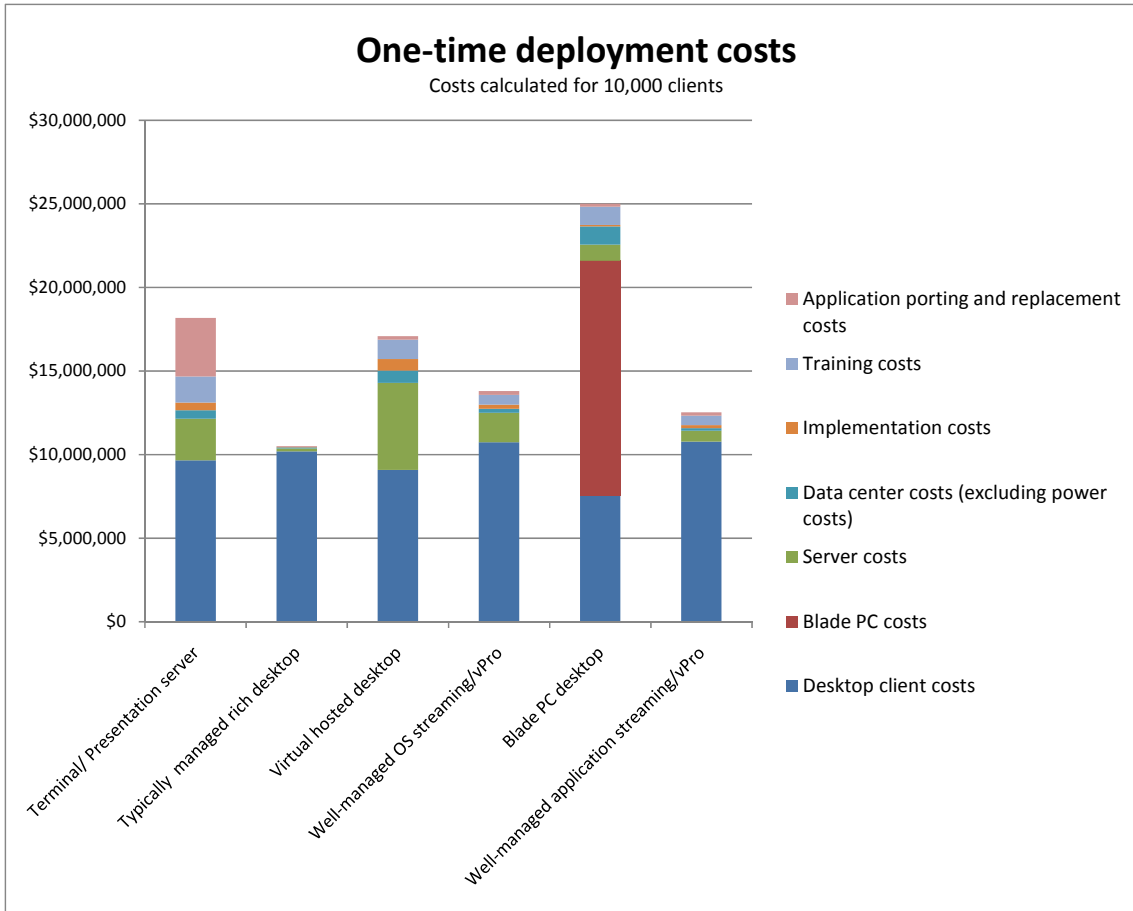


Figure 7: One-time deployment costs of six computing models for 10,000 clients.

Power costs

The power cost covers the electricity needed to run the equipment and to keep it cool. We looked at the costs of the following:

- Power the client devices themselves consume (In the case of Blade PC desktop, we also looked at the cost of power the blades and any supporting hardware required.)
- Power the monitors consume
- Power the servers consume
- Power to cool the equipment

Figure 8 shows the annual power costs of the six computing models for 10,000 clients.

The Terminal/Presentation server model has the lowest power costs. Power costs for all models however, account for only a small fraction of the overall TCO.

In our analysis, server-side computing models had the lowest power costs because the thin client devices consumed less power than rich clients. The Blade PC desktops had the highest power costs because of the need to power both the thin clients and the blade PCs.

Organizations can lower their power consumption by turning off desktops and displays when not in use. For this analysis, we assumed that users turn off desktops and monitors on non-work days, leave them powered on during the work day, and leave them in standby mode overnight on work days.




Figure 8: Annual power costs of the six computing models for 10,000 clients.

Manageability costs

The manageability costs are the ongoing costs to maintain and run the client infrastructure. We looked at the costs of the following:

- Maintaining an accurate inventory
- Patch management
- Support to resolve hardware and software problems
- Adding, moving, or deleting clients
- Security incidents, such as virus outbreaks
- Complying with laws and standards
- Managing the additional access and management servers required by each model

Manageability costs are lowest with server-side and Blade PC desktop computing models

Figure 9 shows the annual manageability costs of the six computing models for 10,000 clients.

Manageability costs are highest for Typically managed rich clients, with the majority of those costs going to desk-side help desk support and services. PCs enabled with Intel vPro technology save over Typically managed rich desktops principally because, aided by third-party management tools, help desk staff can solve more problems remotely, cutting down on desk-side visits and reducing user downtime. Application streaming on those same desktop systems incurs slightly lower help desk costs, saving on software support. Well-managed OS streaming/vPro shows lower costs because they lack disk drives which are frequently the reason for desk-side visits in other client-side models. Server-side and Blade PC desktop computing models have the lowest support costs with few desk-side management requirements.

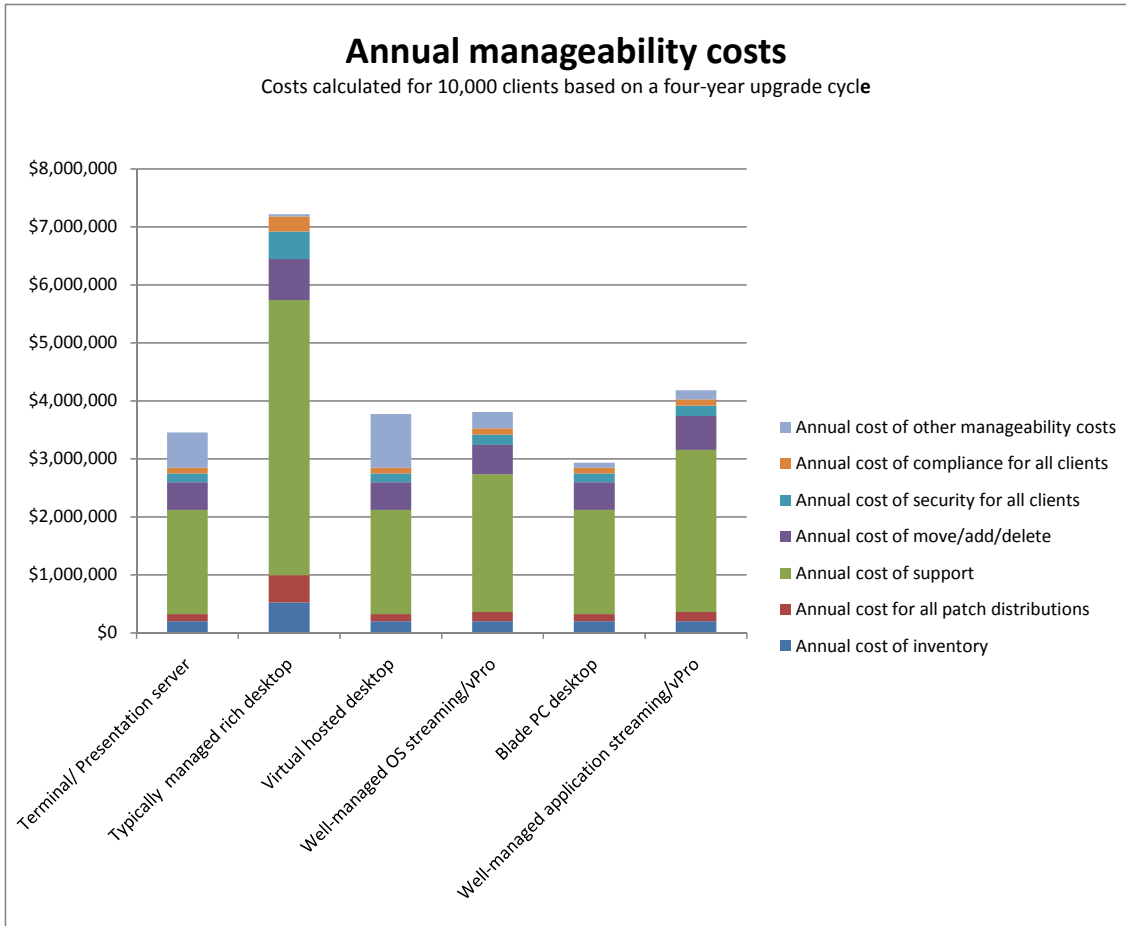


Figure 9: Annual manageability costs of six computing models for 10,000 clients.

Productivity/user experience costs

Changes to user experience and productivity significantly affect compute model choice. With server-side computing models, users typically lose productive time due to server congestion. When multiple users simultaneously execute server-intensive tasks, some or all users may have to wait while the server processes their work. To quantify the amount of time users lose to waiting during these instances, we followed this process:

1. Sum the total number of seconds it took each client to complete our three tests on the model in question when 10 clients were running at once.
2. Subtract the sum of the times the rich clients required to execute the same tasks.

The result is the shared server penalty in productivity.

We estimated that the typical eight-hour workday contains at least the following four periods of peak usage when more than 10 clients are simultaneously executing server-intensive tasks:

- at the beginning of the workday
- before lunch

The shared nature of server-side platforms and the slow nature of Blade PCs hinder user experience and productivity, particularly in the case of knowledge or power users. Combining well-managed rich clients with application streaming and/or OS streaming can provide all the benefits of server-side computing models without significantly affecting productivity.

Because lost user productivity can easily be higher than any of the other costs in this analysis, enterprises need to consider those costs carefully.

- after lunch
- at the end of the workday

To calculate the total daily penalty per user per eight-hour workday, therefore, we multiplied the shared server penalty by four.

In the case of the Blade PC desktop, the script execution times are significantly slower than for rich clients with even one user. Therefore, we estimated that 10 times during the eight-hour workday, a typical user is attempting compute-intensive tasks.

To quantify the amount of time users lose to waiting during these instances, we followed this process:

1. Sum the total number of seconds it took each client to complete our three tests on the Blade PC desktops when 10 clients were running at once.
2. Subtract the sum of the times the rich clients required to execute the same tasks.

The result is the blade penalty in productivity.

We then multiplied the blade penalty for Blade PC desktop by 10 to get the total daily penalty for an eight-hour workday.

Ours is a conservative estimate of costs of lost productivity due to server congestion or blade slowness. We ignore other causes of lost productivity on these models including boot time delays.

Figure 10 shows the annual costs of lost productivity for the six computing models for 10,000 clients.




Figure 10: Annual cost of lost user productivity of six computing models for 10,000 clients.

Total costs

TCO includes the costs of deployment, power, manageability, productivity, and data center space costs. Traditionally, TCO calculations omit costs due to lost productivity. Figure 11 shows TCO excluding productivity for 10,000 clients.

This analysis shows Typically managed rich clients incur the greatest TCO. This model's high costs are due primarily to high manageability costs. The Blade PC desktop model is the next most expensive model because of its high deployment costs. The other models all have similar costs with different balances between deployment and manageability costs. The other TCO cost factor we considered, power costs, had little overall impact on TCO.

TCO is the sum of acquisition costs and sustaining costs. TCO for server-side and Blade PC desktop computing models is higher than for client-side computing models primarily due to deployment costs and productivity losses.

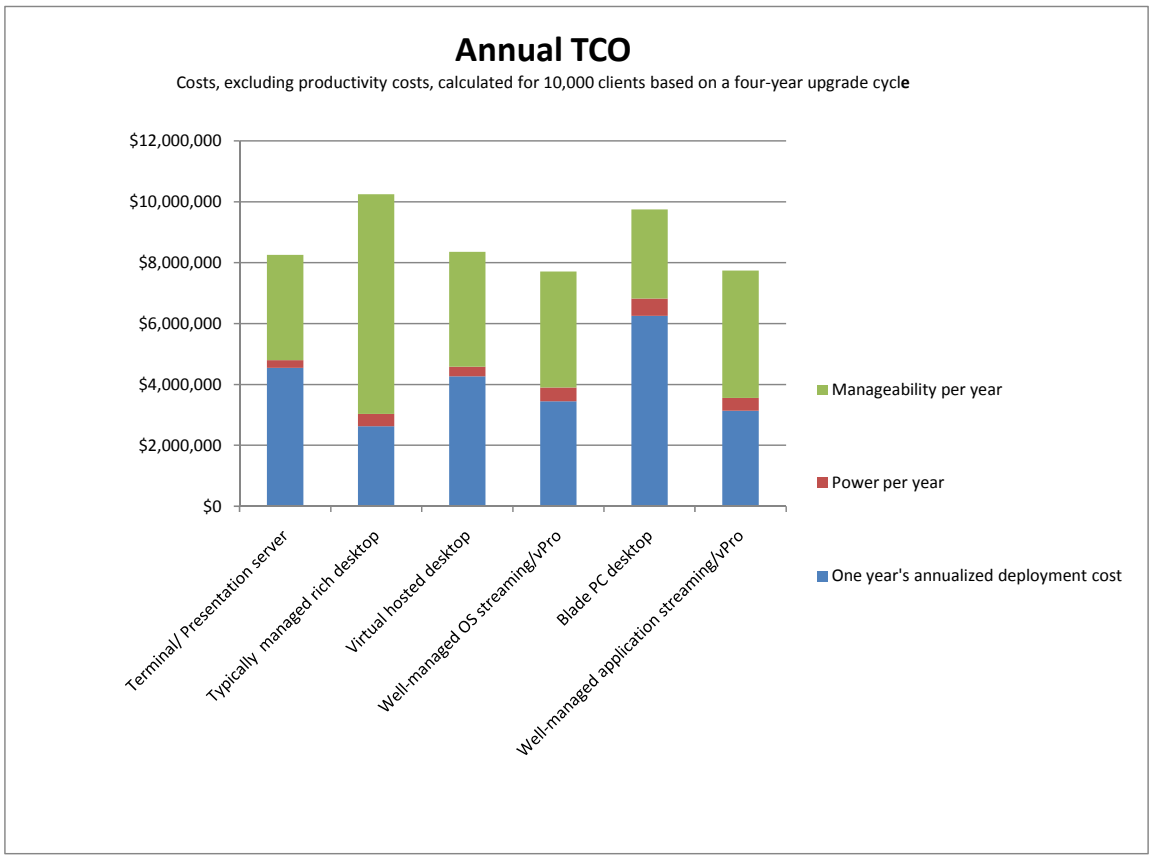


Figure 11: Annual TCO, excluding lost user productivity costs, of six computing models for 10,000 clients based on a four-year upgrade cycle.

Figure 11 excludes costs of lost user productivity, which can be considerable and should be a factor in cost evaluations. Figure 12 shows Annual TCO including productivity costs for 10,000 clients.

The annualized capital and expense costs of the six computing models differ significantly when you consider the costs of lost user productivity. We looked at the productive time users lost due to server congestion or Blade PC slowness. The higher costs of lost productivity for Blade PC desktop, coupled with their higher deployment costs, made that model the most expensive. The two well-managed rich desktop models have the lowest annual TCO

due to low manageability costs and incurring fewer costs for lost user productivity from server congestion or platform slowness.




Figure 12: Annual TCO, including lost user productivity costs, of six computing models for 10,000 clients based on a four-year upgrade cycle.

Security

For security, we examined the vulnerabilities of each model and the effort required to protect against those vulnerabilities. The vulnerabilities we looked at included the following:

- Virus contamination/malware
- Unauthorized access to information
- Theft of proprietary information
- Denial of service (DoS) attacks
- Hacking-related attacks

All other platforms we examined offer considerable security improvements over Typically managed rich desktops.

Typically managed traditional clients are vulnerable to attack. All models except Typically managed rich with Well-managed OS streaming/vPro can provide strong security. A Well-managed rich desktop based on Intel vPro technologies like SIPP, AMT, and Agent Presence matches the security of the best server-side computing models and, when accompanied by OS streaming, exceeds it. OS streaming improves security by supporting diskless clients with no state retention and by allowing disk activity to be redirected to the server.

Future proofing/rich application support

Future proofing looks at the ability of each model to deal with the demands of emerging applications, tools, content, and needs. The factors we looked at included the following:

- The features of the client, such as type and availability of ports
- The demands of rich applications
- The demands of rich collaboration tools such as Live Meeting
- The demands of rich Internet content such as Macromedia Flash animation
- The life cycle of the clients
- The demands of Microsoft Vista (e.g., the Aero interface)
- The expectations of current applications
- The upgrade path for the client type

Client-side computing models based on rich desktops and notebooks offer significant future-proofing benefits over server-side and Blade PC desktop models.

Notably, server-side and Blade PC models offer limited multimedia and rich collaboration support as well as limited Flash-based Internet usage.

Client-side alternative computing models offer total flexibility for future applications. Server-side models are less versatile and less able to support significant processor demands that new applications may require. Multimedia support is limited on many server-side models.

Desk-side environmental

Desk-side environmental refers to those factors that affect users' physical comfort, and thus their productivity. The factors we looked at included the following:

- Noise
- Heat
- Footprint
- Stability

Smaller and usually cooler and quieter than rich clients, thin clients used in server-side and Blade PC desktop models have less of an impact on the desk-side environment.

Client-side computing models that use good case design and power-saving technologies, such as Enhanced Intel SpeedStep technology (EIST), offer much of the same desk-side environment benefits as server-side solutions.

Compliance

Compliance deals with the relative ease or difficulty of complying with license restrictions, laws such as the Sarbanes-Oxley Act (Sarbox) and standards such as the Payment Card Industry Data Security Standard (PCI DSS). Factors we looked at include the following:

- Availability of data for audit or examination
- Susceptibility to sensitive data being modified
- Safety of data from unauthorized access, including erasure of prohibited data (e.g., merchants who do not erase customer credit card information after a transaction, as required by PCI DSS)

All other platforms we examined offer significant compliance benefits over Typically managed rich desktops.

IT is often unable to monitor and enforce compliance with Typically managed rich clients. Server-side models and OS streaming solve and simplify compliance problems by moving the user environment to the server. Desktops powered by Intel vPro technology are highly compliant when used in conjunction with third-party management tools.

Mobility

Mobility looks at the suitability of replacing desktop clients in each model with a notebook. The factors we looked at included the following:

- Ability to work off line
- Compactness
- Complexity
- Licensing impact, as in application streaming

True mobility support is available only with the application streaming computing model.

Because thin client notebooks used with server-side and blade technologies require a network connection, they have limited utility on airplanes or outside of Wi-Fi hotspots. Application streaming solutions can run without network access. Notebooks powered by Intel Centrino Pro technology, with similar functionality to Intel vPro powered desktops, offer mobile users a consistent rich client on-line or off-line experience both on the road and in the office, while improving on the security of traditional rich notebooks.

We could have given the Well-managed OS streaming desktop/vPro platform a rating of Not applicable (NA) on the mobility category because OS streaming rarely goes mobile. However, because the rich client notebooks that support OS streaming in the office can go mobile when running a second local disk profile and because mobile OS streaming is technically even if not practically feasible, we gave the solution a single star.

Microsoft Windows Vista Migration

Enterprises often cite application incompatibility and advanced hardware requirements as the reasons why they have chosen not to upgrade to Microsoft Windows Vista. In this analysis we analyzed costs based on Windows XP Professional version 2002 with Service Pack 2 for all models except Terminal/presentation server. A Microsoft Windows Vista migration would have different costs considerations for the various computing models in this analysis. This section lists some of those differences.

- **Microsoft Windows Vista license costs.** We assumed the enterprise already had Windows XP Professional licenses for existing rich clients and would transfer those to the clients for the selected computing model. Microsoft Windows Vista licenses would add costs. For the Virtual hosted desktop model, Microsoft offers a Vista Enterprise Centralized Desktop license, which allows customers to access Vista enterprise sessions remotely running on server hardware. For the Well-managed OS streaming/vPro computing model, the Vista Enterprise license agreement allows customers to use Vista Enterprise on a diskless PC.
- **Operating system deployment costs:** Deployment costs would be high for the typically managed rich platform due to its high touch requirements for all manageability tasks. The centralized nature of the Virtual hosted desktop, Well-managed OS streaming/vPro, and Blade PC desktop computing models makes Vista deployments less expensive. Vista deployment for the Well-managed application streaming/vPro computing model using enhanced vPro management technologies is less expensive than for the Typically managed rich platform.
- **Costs of incompatible applications:** Commercial off-the-shelf (COTS), custom, or homegrown applications may be incompatible. IT would need to test these applications and replace or port them as necessary. The simplest solution is to delay deployment to users of incompatible applications until after IT has resolved the incompatibilities. However, for enterprises that can't wait, the computing models offer different solutions or workarounds for deploying incompatible applications. Each of these solutions can add significant costs to the deployment.
 - In the Well-managed OS streaming/vPro and Virtual hosted desktop computing models, IT can deploy Vista-incompatible applications by providing multiple OS streamed images or virtual desktops that pair Vista-compatible applications with Vista OS and non-Vista-compatible applications with the original OS.
 - In the Typically managed rich platform, the Well-managed application streaming/vPro computing models, and the Blade PC desktop model, Virtual PC 2007 or dual-boot installation can support Vista-incompatible applications. The client could also connect to legacy applications via RDP or ICA running on Presentation servers, blades or VM servers.
 - A second option for the Blade PC desktop has IT resolve application incompatibilities by providing multiple blade images with Vista-compatible applications paired with the Vista OS and non-Vista-compatible applications with the original OS. This could require considerably more blades to deploy.
- **Microsoft Windows Vista graphics features:** Graphics-intensive technologies such as Windows Aero or streaming video are not well supported in the Virtual hosted desktop and Blade PC desktop computing models.

- **Hardware incompatibility.** The hardware used for our testing would require a memory upgrade to satisfy the commonly accepted Vista memory requirements. Hardware upgrades for the Blade PC desktop computing model are more limited than the other models due to the proprietary nature of blades.

Most of these considerations are not relevant to the Terminal/presentation server computing model because this model does not host applications on the client operating system.

WAN impact

For our study, we assume all users are at a single site or campus. In cases where users are at multiple locations or are separated by slow or inconsistent WAN links, the cost and overall effectiveness of several of the platforms is severely impacted.

Terminal Presentation server, Virtual hosted desktop and Blade PC desktop models may be a better choice over slow WAN links when running applications against back-end corporate servers such as databases. Users of these models, however, may experience degraded performance when they print files locally, copy files across the WAN, or access high volume content such as voice or video. Not only the requesting user but other users at the site see degraded performance during these operations. The bandwidth requirement per user over any single WAN link limits the maximum number of concurrent sessions.

Typically managed rich desktop or Well-managed application streaming/vPro are better choices when running desktop-based applications. WAN links do not generally support OS streaming models.

WAN optimization products improve application performance and reduce bandwidth requirements by optimizing WAN application traffic. Even with these products, most enterprises with lower WAN bandwidth will need to upgrade to closer to 100Mbps to enable sufficient performance for users in server-based computing and OS streaming computing models.

Conclusions

Six computing models offer different benefits for varying costs. Four of the six models have similar TCO if you exclude the costs of lost user productivity. The exceptions are the more costly Typically managed rich desktop and Blade PC desktop models. In the case of specific niches of task workers, there may be little productivity loss, making server-side computing models a viable solution. However, a poorly performing platform can severely affect the productivity of knowledge and power workers. We believe that enterprises should consider productivity costs in TCO analyses as these costs can outstrip all other costs, particularly if a model cannot meet its users' performance demands.

Along with productivity losses, poor platform performance can cause frustration and low morale in the user community. When we include the costs of lost user productivity caused by server congestion on the access servers or the cost due to blade PC slowness, the two models featuring well-managed rich clients—Well-managed application streaming desktop/vPro and Well-managed OS streaming desktop/vPro—have a lower TCO than Typically managed desktop, server-side, and Blade PC desktop models.

Manageability is one key to low TCO. The Typically managed rich client platform has a high TCO largely because of the frequency and high cost of desk-side visits and the cost of user downtime caused by hardware, software, and security problems. The remaining platforms are all well managed and most have a lower TCO. Well-managed platforms have vigilant IT staff employing sound management practices with the aid of remote management software. These practices and tools provide problem prevention and correction with fewer desk-side visits and less user downtime.

The five well-managed platforms have different mixes of strengths and weaknesses. The strengths of the well-managed rich desktop models are future proofing, performance, security, and, for the Application streaming model, when coupled with Intel Centrino Pro notebooks, mobility. Their weaknesses are worse desk-side environment scores and higher power costs than server-side computing models. The server-side computing models are among the most compliant; are highly secure—particularly from theft—have the best desk-side environmental scores; and have low power costs. However, they lag in future proofing because they are less

flexible than rich clients, have high TCO due to user productivity lost during times of server congestion, and are among the least mobile. The Blade PC desktop platform benefits from low manageability costs but is hurt by high deployment costs, lost productivity costs, and TCO as well as low performance.

We found that server-side models may be an appropriate solution for task workers or in places where security or centralized management requirements vastly dominate other factors. However, productivity and mobility considerations can quickly outweigh these issues where knowledge or power users are concerned. Well-managed rich clients supported by third-party manageability software, provide the greatest benefit for the lowest costs. The additional management and security capabilities of Intel vPro technology extend that advantage. Combining well-managed rich clients with application streaming and/or OS streaming can provide the benefits of server-side computing models without significant loss of end-user productivity and result in a lower cost of ownership.



Principled Technologies, Inc.
1007 Slater Road, Suite 250
Durham, NC 27703
www.principledtechnologies.com
info@principledtechnologies.com

Principled Technologies is a registered trademark of Principled Technologies, Inc.
All other product names are the trademarks of their respective owners

Disclaimer of Warranties; Limitation of Liability:

PRINCIPLED TECHNOLOGIES, INC. HAS MADE REASONABLE EFFORTS TO ENSURE THE ACCURACY AND VALIDITY OF ITS TESTING, HOWEVER, PRINCIPLED TECHNOLOGIES, INC. SPECIFICALLY DISCLAIMS ANY WARRANTY, EXPRESSED OR IMPLIED, RELATING TO THE TEST RESULTS AND ANALYSIS, THEIR ACCURACY, COMPLETENESS OR QUALITY, INCLUDING ANY IMPLIED WARRANTY OF FITNESS FOR ANY PARTICULAR PURPOSE. ALL PERSONS OR ENTITIES RELYING ON THE RESULTS OF ANY TESTING DO SO AT THEIR OWN RISK, AND AGREE THAT PRINCIPLED TECHNOLOGIES, INC., ITS EMPLOYEES AND ITS SUBCONTRACTORS SHALL HAVE NO LIABILITY WHATSOEVER FROM ANY CLAIM OF LOSS OR DAMAGE ON ACCOUNT OF ANY ALLEGED ERROR OR DEFECT IN ANY TESTING PROCEDURE OR RESULT.

IN NO EVENT SHALL PRINCIPLED TECHNOLOGIES, INC. BE LIABLE FOR INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH ITS TESTING, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL PRINCIPLED TECHNOLOGIES, INC.'S LIABILITY, INCLUDING FOR DIRECT DAMAGES, EXCEED THE AMOUNTS PAID IN CONNECTION WITH PRINCIPLED TECHNOLOGIES, INC.'S TESTING. CUSTOMER'S SOLE AND EXCLUSIVE REMEDIES ARE AS SET FORTH HEREIN.