Principled Technologies®

WHITE PAPER
February 2007

THIN CLIENTS:
Exploring rollout costs

Contents

Executive summary

Intel Corporation commissioned Principled Technologies, Inc. (PT) to examine the many types of changes and costs organizations should consider when they are contemplating deploying thin clients.

As we discussed in a previous white paper (www.principledtechnologies.com/clients/reports/Intel/ThinClientWP.pdf), thin clients serve well some market niches, such as call centers. Thin clients also appear initially to be far less expensive to deploy than rich clients. It's tempting to make that calculation simply by comparing the purchase prices of the two types of devices.

The real picture, however, is predictably much more complicated, because many of the costs of adopting thin clients are not initially visible. Instead, those costs emerge only as organizations put in place the infrastructure necessary to support thin clients. In addition, each organization's existing hardware infrastructure, network capacity, and level of experience with thin clients can significantly affect the true cost of deploying them.

In this white paper, we discuss the many types of costs of adopting thin clients. In particular, we will examine the following cost areas that the adoption of thin clients may affect:

- Application servers
- Network infrastructure
- Planning
- Application migration
- Training
- Software licensing
- Productivity
- Intangibles
- The future

Figure 1 details the costs for the first six of those areas for a sample company with 32,000 desktop PCs to deploy thin clients across 8,000 of those seats in multiple locations. (More on this sample company below.)

Cost summary*	
Category	**Cost**
Application servers and equipment room	$1,447,530
Network infrastructure	$60,000
Planning	$236,994
Application migration	$3,231,340
Training	$1,002,265
Software licensing	$4,541,800
Total (not including thin client devices)	**$10,511,989**

Figure 1: Thin client adoption cost summary.
**Assumes servers are in a central data center and does not include the costs of the thin clients themselves.*

As you can see, the actual costs of deploying thin clients can be extremely large and reach far beyond the expense of purchasing the devices. (Because our focus is on additional rollout costs, we did not include the expense of the devices themselves in Figure 1.)

Any organization considering deploying thin clients should carefully examine these potential areas.

In the remainder of this white paper we will break down the cost of each of these sections in more detail.

In addition to these costs, organizations considering thin clients may face potential additional costs in the last three areas--productivity, intangibles, and the future--in our previous list. We discuss these three areas later in this paper, but we did not attempt to associate hard costs with any of them due to the high degree to which they depend on such hard to quantify factors as an organization's culture and goals. Nonetheless, these areas are worthy of careful evaluation.

We also limit the scope of this paper to rollout costs. Longer-term operational costs are beyond the scope of this report.

To be able to give example costs, we had to postulate a sample company, which we refer to as our "stalking horse" company. Each organization's real costs will, of course, be different and reflect its size, infrastructure, level of thin client deployment, and so on.

We chose for our stalking horse company the average enterprise in the August, 2006 Wipro Ltd. Study, *Measuring the Value of Intel® vPro™ Technology in the Enterprise* (see Appendix B for more information on it). This company has 32,000 desktop computing devices. For this paper, we assumed that it was considering a move to thin clients for one quarter (8,000) of its desktops. We assumed the 8,000 end users would each get a thin client and were distributed equally across 10 locations. We also assumed an existing well-managed, homogenous Microsoft Windows desktop environment that the company owns (rather than leasing). Because thin clients must be connected to the network to work, we assumed that users would be continuously connected during their business use of the systems.

Figure 2 provides some key information about our stalking horse company.

Basic stalking horse company information	
Total number of PCs to replace with thin clients	8,000
Total number of users for these devices	8,000
Highest percent of thin clients that will connect simultaneously to the servers	100%
Total number of simultaneous connections	8,000
Total number of sites to convert to thin client	10
Average number of thin clients per site	800
Average number of thin client users per site	800
Number of applications run by thin client users	60
Types of users converting to thin clients	**% of total users by type**
Light	50%
Medium	50%
Heavy	0%

Figure 2: Basic company information for our stalking horse.

The types of works the thin client users would be doing will, of course, affect the rollout costs. After studying multiple sources, including http://www.citrix.com, http://www.dell.com, and http://www.hp.com, we settled on a mix of 50 percent light users, such as call center users, and 50 percent medium users, such as general office staff.

To make the thin client purchase expenses concrete, we assumed our stalking horse would be using a $199 thin client, such as the new $199 HP Compaq t5135 (source: http://www.hp.com) that would connect to an access infrastructure containing Citrix Access Suite 4.0 servers.

The following sections explore each of the cost areas in more detail.

Application servers

Thin clients, by their very mode of operation, demand more computing power from servers than PCs and thus lead to higher server hardware costs. The servers that support thin clients must run at least two types of software: the access programs necessary to communicate with the thin clients, and the end-user applications those clients are running. The level of demand each thin client places on the server depends largely on the demands of the applications the thin clients are running. In some cases the existing servers will be able to handle a portion of this new load. Because the servers are doing more work, each one will typically be able to support far fewer thin clients than rich clients.

Consequently, we can expect that our stalking horse will be forced to augment existing servers so it can support the new thin clients. Having to buy more or more powerful servers leads to additional costs, such as more power, more cooling, and; if the solution is to buy more servers; more space to put those servers.

As a starting point for our estimates, we assumed a server could support 60 to 80 simultaneous light users or 40 to 60 simultaneous medium users. In all our server load calculations we use the average of these user counts, or 60 simultaneous connections per server.

Consequently, the company would need 134 servers if it centralizes its servers in a data center, or it would require 140 servers if it deploys them in various locations.

In some cases the servers that make up the thin client access infrastructure do not run the actual end user applications. Instead, those servers only present the thin client desktop and run a web browser or terminal emulator. In such cases, the number of simultaneous connections a server can support will be higher. Each organization should, of course, do its own analysis to determine the number of additional servers it will need.

Some thin client implementations may also require additional supporting servers for the access infrastructure. Most server farm technologies, for example, require some sort of back-end database to maintain server and connection information.

One of the first issues you will face during the transition to thin clients is where to actually place your access infrastructure servers. One option is

to put all of the servers that support the remote sites in the organization's existing central data center. Doing so is likely to increase the demands on the data center's bandwidth. While bandwidth estimates per user vary widely, after checking multiple sources we decided on an estimate of 100 Kbps per active thin client user. We will discuss this network topic further in the next section.

Figure 3 shows the costs of setting up thin clients with centralized servers.

Application servers and equipment rooms cost summary	
Item	**Cost**
Application servers and racks*	$1,060,690
Equipment rooms**	$386,840
Total cost	**$1,447,530**

Figure 3: Thin client costs with centralized servers.
**Assumes 134 servers at a central data center. Includes servers, racks, and other supplies as well as backup and antivirus software. Does not include backup hardware.*
***Includes floor space (Annual Base Cost per Square Foot), UPS, HVAC, network, electrical and monitoring software costs.*

For this analysis, we assumed 60 thin client users per server, which yields a total of 134 servers to support the 8,000 users at our stalking horse. In the centralized model, these servers would be in 7 racks inside the corporate data center.

To provide cost estimates, we selected the dual-CPU HP ProLiant DL360 G5 as our stalking horse's standard server. We then added KVM switches, racks, PDUs, and other supplies, as well as backup and antivirus software, which led to a price of approximately $7,900 per server (source: http://www.hp.com). For server monitoring software, we assumed the company had an existing system management infrastructure and that it would have to buy only additional software or operating system modules. When we considered these items and added the costs of data center real estate, additional HVAC, and redundant power, the average price per server increased to over $10,800.

This centralized server approach may not be ideal, of course, for highly distributed organizations, because in such groups many user communities may not be near a corporate data center. The main factor to consider when deciding where to place the thin-client access infrastructure is the location of the majority of the servers to which the clients must connect. Other key factors to weigh include the cost of increased WAN bandwidth vs. the labor costs of supporting remote servers.

To examine the cost of local servers, we considered the option of placing 14 servers in a single rack at each of the 10 remote locations. To place one rack at each location we would need to buy more racks than in the centralized model. These racks and rack parts would increase the per server cost to $8,000 and change our overall per-server cost (including HVAC, real estate, and other items) to $11,249.

Figure 4 provides a summary of thin client costs with localized servers.

Application servers and equipment rooms cost summary	
Item	**Cost**
Application servers and racks*	$1,121,840
Equipment rooms**	$453,020
Total cost	**$1,574,860**

Figure 4: Thin client costs with localized servers.
**Assumes 140 servers at 10 remote locations. Includes servers, racks, and other supplies as well as backup and antivirus software. Does not include backup hardware.*
***Includes floor space (Annual Base Cost per Square Foot), UPS, HVAC, network, electrical, monitoring software and key-card access costs.*

For those groups requiring localized servers, the local equipment room, which is commonly known as the "wiring closet," may be the only viable solution. Often, phone and network cables terminate there, so the room has some power and contains one or more network devices, such as a switch or router. Such devices do not, however, typically have the power or cooling requirements servers need.

A related problem is that for many organizations the office spaces have guaranteed cooling only during normal business hours. Though that level of cooling is typically enough for the network devices, it may be inadequate for servers running in cramped and often over-heated closets. Addressing these issues raises costs: the expense of additional real estate, the cost of more power and cooling for the servers, and so on. IT organizations may also want to consider adding environmental monitoring gear or software to the closets to make sure the servers stay within their operational guidelines.

To address these issues for our stalking horse, we added a set of $2,400 2U rack mounted UPSs (source: http://www.hp.com) for each rack to handle the power continuity requirements. We also added an auxiliary cooling unit from http://www.portablecooling.com to handle the

around-the-clock HVAC needs for each site. Finally, we added a product (we used the NetIQ monitoring software, http://www.netiq.com) for remote server management.

A final potential cost issue involves securing the wiring closet. Although any room with networking infrastructure components should be secure, this is not always the reality. Organizations may be lax in this area when the only wiring room contents are cables and switches. When the rooms contain servers that hold potentially sensitive corporate and/or client data, security must become a paramount concern. Increasing security would add the cost of locks and key-card access, which can run to over $1,000 per wiring closet (source: http://www.kaba-ilco.com).

Network infrastructure

The data that must flow between thin clients and their supporting servers is relatively small and consists primarily of the information necessary to update the thin client displays. The resulting additional bandwidth requirements can, nonetheless, be significant in a large deployment.

Figure 5 shows the network infrastructure additional costs our stalking horse would face.

Network infrastructure cost summary	
Item	**Cost**
Upgrade to 100 Mbps (cost per port)	$0
Remote network or WAN improvements	$60,000
Network optimization technologies	$0
Total cost	**$60,000**

Figure 5: Thin client network infrastructure cost summary.

The amount of bandwidth each thin client requires varies a great deal with the types of applications the clients are running. A 2000 study by The Tolly Group found running such applications such as Microsoft Office on thin clients can dramatically increase the amount of bandwidth those clients require versus the amount necessary to support simpler applications exchanging screen image data between the clients and the server.

It is vital to note that this low cost assumes, as Figure 2 shows, that our stalking horse already has all of its users on a 100 Mbps network. As we mentioned earlier, we are assuming the thin client users at our stalking horse will each consume around 100 Kbps of network bandwidth. If we calculate that load across the 800 users per site, we would see a utilization of nearly 80 Mbps, easily saturating any 10Mbps network. A company not yet at 100 Mbps can expect to pay $124 per port for enterprise-grade switches (source: http://www.networkworld.com/news/2006/110306-lan-questions-cisco.html). This seemingly low cost would translate to approximately $990,000 for our 8,000 users if the remote sites were still at 10 Mbps. Organizations might also face additional costs for rewiring, depending on the quality and type of existing cabling.

As serious as the bandwidth issue may be for users in the same facility as their servers, it can be even more of a problem for users at remote locations. Those users often connect to their main networks with fractional T1, DSL, or cable-modem connections. Such connections are unlikely to provide adequate performance for significant numbers of thin clients, which means either moving servers closer to the clients or adding more bandwidth.

According to the Citrix Ace Cost Calculator, (http://www.acecostanalyzer.com) a company may have to upgrade its WAN if its bandwidth is at 50 percent or greater saturation. Adding the necessary WAN bandwidth to support the stalking horse's 800 remote users in each of 10 different facilities can add from hundreds to thousands of dollars per month per site. A monthly increase of just $500 per site would add over $60,000 to the stalking horse enterprise's annual communications budget. Figure 2 reflects those costs.

One way to help alleviate some of this WAN bottleneck for companies considering the centralized server model is to add network optimization technologies, such as Citrix WANScaler. According to Citrix, this hardware product can accelerate application performance by an average of 5 to 30 times across wide area networks. The devices retail at $8,000 for remote sites and $40,000 for the data center model (source: http://www.citrix.com). We did not include optimizer costs in this analysis.

In any thin client implementation, the network becomes a single point of failure, because thin clients cannot work without a supporting server and a connection to that server. Though rich-client users can keep working even when their network connections or servers are temporarily down or inaccessible, thin-client users cannot. Even if the network does not stop entirely and instead simply starts delivering degraded performance, thin-client users may experience poor response times and decreased productivity. Any organization considering a move to thin clients should include

the cost of avoiding these network problems in its rollout cost analysis.

Planning

Another cost of any IT initiative is planning and pilot testing. In our research, we found references for planning and installation costs ranging from 2 months to over a year. These tasks require internal staff involvement, and often companies may also choose to involve outside consultants.

Based on several Forrester studies on The Total Economic Impact of Deploying Sun Ray Clients (SunRay_FinalTEI_040504.pdf), we estimated that the project would take 2.7 full-time equivalent (FTE) staff a total of 9 months. These FTEs would be comprised of 30 percent senior administrator time and 70 percent junior administrator time, which results in a cost of nearly $237,000.

We also found that many companies use outside consultants to supplement internal thin client knowledge. Assuming a conservative engagement length of 16 days at a rate of $2,000 per day adds another $32,000 in cost to our project.

Figure 6 illustrates the cost summary for planning.

Planning cost summary	
Item	**Cost**
System administrator	$201,150
Outside consultant	$32,000
Other staff	$3,844
Total cost	**$236,994**

Figure 6: Thin client planning cost summary.

We also include some planning time for staff other than the system administrators. These tasks include attending planning meetings and supporting the project administrators. Assuming a weekly salary of $1,025 for such personnel adds approximately $3,800.

These costs for other staff members will be higher in implementations where less consultant time or system administrator time is available for the rollout. For example, if a staff member with an annual salary of $120,000 must spend 4 months on the effort, salary costs alone would go up by more than $40,000.

Application migration

Another potentially large cost of moving to thin clients is application migration. Most applications today target the PC platform. Running those applications on thin clients often means running multiple copies of them on the application servers.

Custom legacy line-of-business applications written for PCs bring their own special problems. Such applications frequently assume that they have access to the full resources of the PC on which they are running. In a multi-user environment, where multiple thin clients are sharing the server, this is not the case. The problem may be as simple as a temporary file with a hard coded path name, or an application may not work at all and may require architectural and programming changes to enable it to run correctly. In either case, fixing the problem will require either figuring out how to trick the application into working in the new environment or modifying the application's code.

Users at our stalking horse have over 100 applications on their servers and PCs. Part of the planning phase of the thin client deployment must include reducing that list to a smaller number of applications the thin-client environment will support. We estimate that 60 of our over 100 applications will survive these cuts.

The stalking horse enterprise will use Citrix software to publish its applications and desktop environments to the thin clients. Citrix will help to eliminate some of the application migration costs by allowing many of the applications to migrate directly.

To avoid encountering problems during normal business operations, companies considering thin clients must also install and test all such applications in a server-based computing environment before they deploy the thin clients. The stalking horse enterprise's analysts will need to install and test each of these applications on the servers, test and tune them, and then recode as necessary.

As Figure 7 shows, the total of all of these types of migration costs can be quite large.

Application migration cost summary	
Item	**Cost**
Application installation and testing	$56,400
Application recoding	$550,000
Application replacement	$2,625,000
Total cost	**$3,231,400**

Figure 7: Thin client application migration cost summary.

We will assume that 11, or roughly 18 percent of the 60 applications at the stalking horse, will require some level of recoding. We further estimate that 5 will require simple porting or tuning, at a cost of $10,000 per application; 5 will require medium porting, at a cost of $40,000 per application; and 1 will require a more complicated revision, at a cost of $300,000. The combined total costs for these recoding efforts would be approximately $550,000.

Depending on the application, it is entirely possible that it is cheaper to buy rather than build a thin client-friendly application. Purchasing such applications for a portion of our 8,000 users can cost from hundreds of thousands up to millions of dollars. For instance, our stalking horse may decide to replace its legacy accounting system with Microsoft Dynamics GP (Great Plains) Software. If we assume we will need to buy 100 copies of an application at a cost of $2,250 per concurrent user, we will spend another $225,000 (source: http://www.internetnews.com/ent-news/article.php/3619411). If it is a call center application that 2,000 users need at $1,200 per seat, then we have yet another $2,400,000 of cost (source: http://www.buyerzone.com/software/call_center_software/buyers_guide9.html).

The enterprise must test all of the applications in the new environment. Testing the 60 applications at our stalking horse will cost an estimated $56,340, assuming 2.5 testing days per application and a daily system administrator labor cost of $375.60 (Source: National average for system administrators at http://www.salary.com).

Training

Because thin clients are ideally drop-in replacements for PCs, you may be tempted to imagine that no training costs are necessary. Most organizations should plan, however, on some training costs for users, support staff, and IT staff. For the purposes of this paper we assembled rough averages of pertinent salaries from a variety of sources, including http://www.salary.com.

Figure 8 provides a breakdown of the various areas that contribute to the total cost of thin client training.

Thin client training costs summary	
Item	**Cost**
User training	$664,000
System administrator training	$299,490
Help desk training	$38,775
Other training	$0
Total cost	**$1,002,265**

Figure 8: Thin client training costs summary.

Users

Using a thin client is in many ways like using a PC, but it is also somewhat different. Users will need training on both the plusses of thin clients, such as the ability to log into any station and see their data, and the potential minuses, such as the potential performance issues, what to do when something goes wrong, and so on. New hires that are likely to have used PCs are unlikely to have used thin clients and so may require additional orientation. They will also require training on the applications, file management, and procedures.

If we assume the total time for training each of our existing end users is 3 hours (source: http://www.sun.com) and the average call center and office employee makes $205.33 in salary and benefits per day, the costs become quite substantial. For our stalking horse's 8,000 users, this translates to over $600,000 in lost productivity during training alone. Adding in training fees brings the cost to $664,000. A company could choose to skip this training, of course, but it would probably encounter similar or greater productivity issues due to users having to train themselves.

Help desk staff

Training the help desk staff to support the thin client users is likely to be a more substantial cost per person. Support staff must learn new types of troubleshooting and problem-resolution techniques. Some of those techniques, such as assessing available bandwidth to see if it is adequate, will be common across thin clients. Others, however, will vary with the type of thin client, because each vendor's thin client solution is different.

Support staff will need this training because some of the approaches they use with PCs will not work with thin clients. The classic first line of defense, “Power down your system,” does not work for a server in a locked wiring closet. Figuring out what is wrong is also always more complex when

multiple systems (the thin client and the server) are involved.

Each thin client has its own type of administration tools, so support staff members are likely to need specialized training in this area as well. Finding new staff with training in a particular type of thin client administration will generally be harder than finding staff with PC administration experience. This expense can translate to thousands of dollars per new hire per year.

Similarly, the support staff members who install applications will need to work differently than they have in the past. Users will not be able to install their own applications at any time, because they will not have the hardware to do so. Though many organizations already try to avoid having users do this work, most also know that a great many application installations occur via end users. The gain in control of the workstation applications comes with a corresponding increase in the need to manage those installations and the user requests for software not currently on the systems.

For our stalking horse company we calculated the number of help desk staff at one user for each 727 thin client devices (source: http://www.sun.com). If we assume each of our 12 thin client help desk personnel will require 5 days of training, we have a total cost for downtime and fees of over $38,000. This cost could be higher if the company does not plan to create specialized thin client help desk personnel. If all help desk staff serve all needs, then all staff would need training on the thin client environment. For our stalking horse, that could increase the training need from just the 12 FTEs for the thin clients to their total of well over 60 help desk staff and thus increase the cost to over $100,000.

System administrators

The central IT teams that manage the servers and network infrastructure are the groups that will require the most thin-client training.

The administrators deploying thin clients will need training in configuring and upgrading servers appropriately. Most IT staff members already understand the RAM and hard disk needs of their users and the options PCs offer in those areas. With thin clients, those people will need to understand the issues of server capacity and capability management.

Selecting the right thin client will also require some IT training and research. Buyers must decide, for example, which protocols they need to support, whether to have a browser on the thin client, whether to include USB ports, and so on.

Even the operating system on the thin client is not a given. Different thin clients run Windows CE, Windows XP embedded, Linux, and proprietary operating systems. The 2005 IDC thin client market analysis predicts that in 2009 the market will be split fairly evenly, with 25 to 30 percent using each of Windows CE, Windows XP embedded, and Linux operating systems; the remaining 12 to 15 percent will employ other operating systems, including proprietary ones. Choosing the right operating system will require additional analysis, and knowing how to support each thin client operating system the organization uses will require more training.

We assume that each junior administrator at our stalking horse company can support 400 devices and each senior administrator can support 1,200 devices (source: http://www.sun.com). This translates to 7 senior administrators and 20 junior administrators for the 8,000 thin client devices. Each of these staff members will require a variety of thin client training. This training will include at least device management software training, general Citrix or terminal server training, and advanced architecture training. Combined, these types of training will add over $299,000 to the stalking horse enterprise's training budget.

Software licensing

One area of potential complexity that is easy to overlook is software licensing. On the good side, it is possible that by carefully managing usage and taking advantage of licenses that charge for simultaneous and not total usage, an organization can end up needing fewer total licenses for user software than the number of users.

The bad news is that setting up thin clients often will involve procuring new software licenses. Applications that were running on the PCs and that would run on newer PCs may not be compatible with the operating system version on the server, or the organization may not have licenses for the particular versions currently in use.

Consequently, an organization considering thin clients must examine each of the applications its users require and figure out how to make those applications available from the server. Doing so may be as simple as installing the software on the server or as complex and expensive as migrating to new, server-based versions of the software.

Figure 9 summarizes the costs from purchasing software licenses.

Software licensing costs summary	
Item	**Cost**
Microsoft Windows Server 2003 R2 Enterprise Edition License	$308,200
Microsoft Windows Server 2003 Client Access Licenses (CALs) per device	$0
Microsoft Windows Server 2003 Terminal Server (TSCALs) per device	$400,000
Citrix Access Suite 4.0 Connection Licenses – Enterprise	$3,833,600
Microsoft Office Standard Edition 2003	$0
Total cost	**$4,541,800**

Figure 9: Thin client software licensing costs summary.

Server software costs beyond the applications themselves are also not always easy to identify, but they typically exist. Such costs may include Windows server software, remote access software such as Citrix, new thin-client management software, and so on.

We assume that our stalking horse plans to use Windows servers along with Citrix access software to support its 60 applications. Among those applications is Microsoft Office. We assume that the stalking horse already owns the necessary Windows Server user licenses (CALs) and the Microsoft Office licenses for use by the PCs that the thin clients are replacing. The company can transfer these licenses to the thin clients.

Even so, the total cost for Windows Server licenses, the Terminal Server CALs, and the Citrix licenses would run about $4,550,000, including discounts and assuming 100 percent user concurrency on the Citrix servers (source: http://pcuniverse.com).

Productivity

Companies adopting thin clients may also face productivity costs if they do not carefully manage their network infrastructure and supporting servers.

We discussed potential network issues earlier. In general, organizations moving to thin clients will need to assess and possibly improve both the reliability and the bandwidth of the networks supporting those clients. Otherwise, users can pay response-time penalties when their bandwidth requirements overcome the network.

Server performance can also affect productivity. As our study (http://www.principledtechnologies.com/clients/reports/Intel/ThinvsPCperf.pdf) showed, even as few as 5 thin clients running the same office applications at once could bog down a server and experience dramatically worse response time. These problems typically occur at key times, such as the beginning of the day, during crunch times, and so on, when many users are trying to do computationally demanding work at the same time.

These problems are likely to get worse, rather than better, over time, because of the increasing trend toward multitasking. Users accustomed to the power of PCs are increasingly performing multiple computing tasks simultaneously. Having multiple applications active at the same time and frequently switching between them is a common way to work in most organizations. Each application adds a demand to the host server. Business users increasingly expect to perform more than one of these tasks at a time and still experience reasonable performance. Consequently, organizations deploying thin clients can meet such expectations only by making sure the supporting servers have enough power to run all those applications for all active users.

Lost productivity is especially costly for large companies such as our stalking horse and its 8,000 thin client users. The cost per hour of downtime would be $160,000 using an average salary and benefits of $200 per person per day. Even a period of network degradation that results in a productivity/time loss of 25 percent would cost upwards of $40,000.

We did not include these costs in our earlier tables, but organizations evaluating thin clients should certainly consider them carefully.

Intangibles

Adopting anything new or different typically means experiencing some adjustment cost. Users accustomed to the way PCs operate will have to learn how to use thin clients effectively.

Many organizations face the potential for user backlash at having to move from PCs that they think of as theirs to a shared server with a terminal on it. Though users' opinions generally do not factor into cost equations, they can result in work attitudes that affect productivity.

The future

One of the hardest expenses to estimate in any analysis is opportunity cost, and that is certainly

the case with thin clients. One of the biggest advantages of PCs over the years has been their versatility and their ability to run new, never-before-thought-of applications and hardware.

Thin clients, by contrast, are limited devices. Doing more on them means either doing more on the server or replacing the clients. Hardware changes in the future will always be a challenge for thin clients. Where PCs are extensible--admittedly at the IT cost of installation--thanks to their open architecture, thin clients are by nature usually not upgradeable. Features they exclude, such as new I/O technology ports, are often the very things that interesting future technologies require.

Similarly, future software of interest may either swamp servers or not work at all. Real-time communications applications, such as VoIP or Microsoft Office Communicator, often require substantial system resources. In a recent white paper (http://www.principledtechnologies.com/clients/reports/Intel/SkDCWP.pdf), we examined the effect of VoIP combined with multitasking applications and found that the processor demands were indeed high. Only the PCs with dual-core processors were able to keep up with that demand. Similar demands from multiple thin-client users would quickly swamp the capabilities of most servers currently providing processing power to those clients.

Few people would argue against the observation that over time applications have tended to require more performance. That trend is only continuing as more applications move to support data formats, such as XML, that can place significant performance demands on the underlying processor. These trends only accelerate when you consider new applications and application features that involve media content such as pictures, sound, and video.

Advances in desktop operating systems and office applications also often lead to increased processor demands. The heavier processor requirements of Microsoft Office 2007 and the stronger but potentially more graphically intense interface of Microsoft Windows Vista are likely to pose challenges for thin clients as those packages become commonplace.

Thin Client Hardware

We do not include the costs of the thin clients themselves because we assume that the stalking horse company was already planning on replacing the PCs with either PCs or thin clients. The purchase price of thin clients can be appealing, although price differences between PCs and thin clients are diminishing. Bob O'Donnell's guidance from the 2005 IDC thin-client market analysis included this observation: "One critical point that thin client vendors need to plan for is the rapidly approaching price points of low-cost PCs….the purchase price-only benefits of thin clients will eventually go away or get so small as to be inconsequential."

The HP Compaq t5135 Thin Client that we selected for our stalking horse company is a basic thin client workstation that supports both the RDP and ICA protocols. It costs around $199 without a monitor.

A range of "fatter" thin clients have come onto the market over the last several years. Some of these fatter versions offer more advanced features, such as a JAVA Virtual Machine (JVM) for running JAVA applications locally. Local web browsers and local terminal emulators are also available. The more thin clients provide such abilities, of course, the closer they come to being PC-like rich clients. These devices are beyond the scope of this study.

PC changes

No discussion of the issues to consider about the costs of adopting thin clients would be complete without at least a brief mention of the changes in ownership costs of PCs. Multiple PC industry efforts are aimed at improving PC manageability and reducing the management cost differences between PCs and thin clients. Initiatives from such companies as Intel, AMD, and Microsoft have put PCs and thin clients on much more of an equal footing in this area.

Thin clients traditionally have possessed a management cost advantage in deployment, moving, and repair. These advantages stem from a single feature: the ability to set up a thin client anywhere there is a network connection and access the supporting server. Whether the problem is to set up a new user (deployment), transfer a user to a new location (moving), or deal with a broken system (repair), the solution is the same: drop in a new thin client. The thin clients within a given type are largely interchangeable because most, if not all, of each user's online state resides on the supporting server.

PCs, however, are greatly reducing these advantages by moving toward similar ways of management, but without the need for so much work on the server side.

One key management policy change is to keep copies of most, if not all, critical data on servers.

This has long been possible but has often required appropriate policies and user compliance. Remote PC management, server-enforced policies and profiles, and automated data backups are helping in this area.

Another key change is the ability to easily maintain a minimal number of consistent PC disk images, so setting up users on different PCs becomes a much simpler endeavor. Initiatives such as Intel's Stable Image Platform Program (SIPP) and AMD's Commercial Stable Image Platform (CSIP) will help address this issue. These initiatives enable systems to share images and thus be much easier for IT departments to maintain.

Operating system vendors such as Microsoft and Red Hat are also helping by providing tools that aid in desktop installation and monitoring, as well as software and patch deployment.

Finally, other management solutions such as Intel's Active Management Technology (AMT) can help simplify the IT management challenge by supporting remote maintenance by IT staff nowhere physically near the PCs.

Reducing manageability costs is clearly a major imperative for the PC industry as a whole, so we can reasonably expect continuing and significant improvements in this area.

Summary

The cost of adopting thin clients may seem as simple as the cost of buying the systems, but the real story is far more complex and expensive. Thin client rollout costs actually include a broad range of hardware, software, and staff expenses. Organizations adopting or considering moving to thin clients should consider this broad array of potential cost sources if they want to accurately assess the true cost of deploying thin clients.

Appendix A: Key thin client terminology

Thin clients come in several different styles, each of which uses a somewhat different architecture for delivering computing power to users. Different vendors and analysts often use different terms for the same concepts and approaches. To avoid confusion, we use the following consistent set of terms, which we list in alphabetical order, to help make clear which thin client approaches and technologies we're discussing.

- *Citrix Independent Computing Architecture (ICA)* – ICA is a protocol that is closely related to RDP and that allows a user to connect to a computer running the appropriate underlying software. In this case, that software is the Citrix Presentation Server. The front-end device may be a thin client or a rich client (PC) running the necessary Citrix software.
- *Microsoft Remote Desktop Protocol (RDP)* – RDP is a protocol that allows a user to connect to a computer running Microsoft Terminal Services. Typically, a thin client front-end device will use RDP to connect to a server. The front-end device, however, could also be a PC running software that uses this protocol.
- *Server-side computing* – A form of computing in which almost all of the processing (with the usual exception of the display of information) happens on a shared server rather than on a client system. A single server almost always serves the computational needs of multiple clients. (If each server supported only one client, the servers would effectively be acting as PCs that happened to sit far away from the users.) The server may in turn use other devices, such as file servers, to meet some of its requirements, such as disk storage. Server-side computing is basically what mainframes and minicomputers provided to users before the era of the PC.
- *Shared server* – A server that allows multiple thin clients to run applications simultaneously. A shared server typically provides server-side computing to the thin clients. The collection of shared servers is the *Access infrastructure.* When a company also uses a shared server to run end user applications, that server is an *Application Server.*
- *Terminal Services* – Terminal Services is a set of software capabilities that is part of newer versions of Microsoft Windows. It lets a user access applications on a server over a network connection using RDP.
- *Thin client* – A thin client is a computer that relies on a server or other back-end system, such as a blade PC, to supply most or all of its computation needs. While such a device often possesses a lot more intelligence than most terminals did in years past, its primary task is displaying information. Newer thin clients have begun to include the ability to run browsers and some other applications locally. The more thin clients provide such abilities, of course, the closer they come to being PC-like clients.

Appendix B: Detailed company profile

This appendix provides a detail profile of the relevant information we used for our stalking horse company.

Detailed stalking horse company information*	
Cost factor	**Value**
Number of desktop computing systems	32,000
Number of sites	40
Number of desktop hardware vendors the company supports	3
Desktop computing refresh cycle	4 year
Percentage of desktops that automation management software supports	75%
Number of desktop PC applications the company deploys per year	102
Number of PC models the company deploys per year	8
Number of desktop PC inventories per year	10
Patches, audits, and security incidents per year	516
Average annual burdened rate for staff - level 1	$80,611
Average annual burdened rate for staff - level 2	$110,484
Average annual burdened rate for staff - level 3	$144,988

Figure 10: Detailed relevant information for our stalking horse company.

*Source: *Measuring the Value of Intel® vPro™ Technology in the Enterprise,* Wipro Ltd. (2006).

About Principled Technologies

We provide industry-leading technology assessment services. We bring to every assignment extensive experience with and expertise in all aspects of technology testing and analysis, from research into new technologies, to the development of new methodologies, to testing with existing and new tools.

When the assessment is complete, we know how to present the results to a broad range of target audiences. We provide our clients with the materials they need, from market-focused data to use in their own collateral to custom sales aids, such as test reports, performance assessments, and white papers. Every document reflects the results of our trusted independent analysis.

We provide customized services that focus on our clients' individual needs. Whether the technology involves hardware, software, Web sites, or services, we offer the experience, expertise, and tools to help you assess how it will fare against its competition, its performance, whether it's ready to go to market, and its quality and reliability.

Our founders, Mark Van Name and Bill Catchings, have worked together in technology assessment for over 20 years. As journalists they published over a thousand articles on a wide array of technology subjects. They created and led the Ziff-Davis Benchmark Operation, which developed such industry-standard benchmarks as Ziff Davis Media's Winstone and WebBench. They founded and led eTesting Labs, and after the acquisition of that company by Lionbridge Technologies were the head and CTO, respectively, of VeriTest.

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