Principled Technologies®

WHITE PAPER
December, 2006

THIN CLIENTS
Understanding the costs

Contents

Executive summary

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

As we have discussed in a previous white paper (http://principledtechnologies.com/clients/reports/Intel/ThinClientWP.pdf), thin clients serve certain market niches well. On the surface, cost comparisons of thin and rich clients may seem as simple as checking which device is less expensive to purchase. Thin client advocates also frequently tout lower Total Cost of Ownership (TCO) as a major inherent advantage of thin clients due to their relatively simple architectures.

The real picture, however, is predictably much more complicated. The complication stems from the fact that 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.

These costs come from both the potentially major infrastructure costs necessary to support thin clients initially and the additional support costs organizations will face over the lifetime of the systems. To compare those costs to the costs of using rich clients, organizations must also consider management initiatives in the rich-client space as well as a variety of intangibles and the "future-proofing" that newer generations of rich clients can provide.

In the rest of this white paper we will look at the following key thin client cost areas:

- Thin client hardware
- Application servers
- Network infrastructure
- Application migration
- Training
- Software licenses
- Productivity
- Intangibles
- The future

We will also examine some changes in PCs that have the potential to further alter the PC/thin client cost equation.

We begin, however, by reviewing some terminology that is helpful when discussing thin clients and the infrastructure necessary to support them.

Key thin-client terminology

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

- *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 all users before the era of the PC.
- *Shared server* – A server that allows multiple thin clients to run applications on it simultaneously. A shared server is typically providing server-side computing to the thin clients.
- *Blade PC* – A PC on a card that operates in a chassis in a rack-mounted configuration. To use a blade PC, a user typically runs a thin client device on his or her desktop, and that device in turn uses the blade PC for its computing. Blade PCs were traditionally *dedicated*, meaning one blade PC served one user. Sharing the computational power of a single blade PC among multiple thin clients is certainly possible with more powerful blade PCs, and some vendors are doing so.
- *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 rich clients.

- *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.
- *Citrix Independent Computing Architecture (ICA)* – ICA is a protocol that is closely related to RDP and that also allows a user to connect to a computer running the appropriate underlying software. In this case, that software is the Citrix Presentation Server. As with RDP, the front-end device has traditionally been a thin client, and the other system a server. The front-device, however, could also be a PC running the necessary Citrix software.
- *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.

Costs to consider

Accurate cost assessments of thin clients can be difficult because thin clients require a supporting hardware and software infrastructure that is not always immediately evident. In the following sub-sections, we examine the key types of costs that organizations may face when adopting thin clients.

Thin client hardware

Not many years ago, the purchase price difference between PCs and thin clients was huge. Today's lower PC costs have cut into that difference significantly. It is still the case, however, that many thin client hardware devices are significantly cheaper than PCs--as long as you consider only the cost of the device itself.

The story becomes more complex when you factor in the software and, in some cases, hardware necessary to support the thin clients. 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."

As an example of how complex this picture can be, consider a set of test beds that we assembled during April and May of 2006. Our goal was to set up five-unit test beds that employed different types of clients:

- typical business PCs (Dell OptiPlex 210Ls)
- Sun Ray thin clients (Sun Ray 2s)
- Wyse thin clients (Winterm 5150SEs)

The table below summarizes what we had to spend to set up each test bed. Each test bed used the exact same model and configuration of server, so we do not include that server's cost. We also do not include the cost of monitors, because we used the exact same model monitor on all systems. We also omitted shipping and tax.

Five-client test bed	Price
Sun thin clients	
Sun Fire v240 server	$ 6,924
5 Sun Ray 2 thin clients	$ 1,814
Total	$ 8,738
Wyse thin clients	
5 Wyse Winterm 5150SE	$ 2,055
Citrix 5-user license	$ 1,017
Total	$ 3,072
PCs	
5 Dell OptiPlex 210L	$ 4,215
Total	$ 4,215

As you can see, both types of thin client devices were significantly cheaper than the PCs. The Sun Ray 2 devices, however, required a Sun Fire v240 server to be able to talk to our test Windows server, so the total cost of the hardware necessary to run the Sun Ray 2 thin clients was actually higher than that of the PCs. (Sun has since announced the ability to run the Sun Ray clients directly against the Windows server using RDP, so this cost would go away.) The Wyse clients required a Citrix license, which brought their purchase price up significantly.

Note also that we did not choose anywhere near the cheapest PCs available. Instead, we went for standard OptiPlex systems with Intel processors with HT Technology.

All that said, the per-seat cost of the hardware devices themselves is clearly lower with thin

clients, though by far smaller a margin than in the past.

The total hardware cost, however, is actually greater than this picture shows, because the thin clients by their very mode of operation demand more computing power from servers than PCs and thus lead to higher server hardware costs.

As our tests found (http://www.principledtechnologies.com/clients/reports/Intel/ThinvsPCperf.pdf), in many common business scenarios PC clients place relatively small demands on servers, while thin-client demands can be significant. The following subsection examines the types of application server costs that adopting thin clients can bring.

Application servers

The application servers that support thin clients obviously must run at least two types of software: the programs necessary to communicate with the thin clients and the 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. Because rich clients can run most, if not all, of their applications locally, rich clients typically need significantly less application server resources.

The cost of supplying the application servers is thus likely to be significantly higher with thin clients than with rich clients. Any comprehensive analysis of the cost of deploying thin clients must include those additional server costs.

Because the servers are doing more work, each one will typically be able to support fewer thin clients than rich clients. In the study we mentioned earlier (http://www.principledtechnologies.com/clients/reports/Intel/ThinvsPCperf.pdf), we found that with just five thin clients running the same Microsoft Office-based tasks the response time each user received worsened dramatically. (Rich client users, by contrast, saw either no or much less response time degradation.)

Consequently, moving to thin clients will typically mean buying either more or more powerful servers than would be necessary to support the same number of rich 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.

One option is to put the servers in the organization's existing data centers. Doing so is likely to mean increasing the available bandwidth out of the data center; we discuss this topic further in the next section.

This approach can also be a problem for highly distributed organizations, because in such groups many user communities may not be near a corporate data center. For those groups, the local equipment room, often known as the "wiring closet," may be the only viable solution.

Many organizations call these rooms wiring closets because that's what the rooms used to be. 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 of servers.

A related problem is the fact 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 the cramped and often over-heated closets.

Addressing these issues raises costs, either expenses to find additional real estate or the costs of more power and cooling for the servers--or both. IT organizations may also then want to consider adding environmental monitoring gear to the closets to make sure the servers stay within their operational guidelines.

A final potential cost issue involves securing the wiring closet. Though it's certainly the case that any room with networking infrastructure components should be secure, organizations may be lax in this area when the only contents are cables and switches. When the rooms contain servers that hold potentially quite sensitive corporate and/or client data, however, security must be a paramount concern. Increasing the security can add the cost of key-card access or something similar.

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, but the bandwidth requirements can nonetheless be significant.

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 vs. the amount necessary to support simpler applications exchanging simple screen image data between the clients and the server.

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 are often connected 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.

Any organizations not yet running 100-Mbps (or faster) Ethernet will also very much want to consider upgrading to that level before adopting thin clients.

The nature of thin clients means, of course, that they 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.

The network thus constitutes a single point of failure for all the thin client users on it. The cost of avoiding this problem should factor into the analysis of any organization moving to thin clients.

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're running. That is not the case in a multi-user environment, where multiple thin clients are sharing the server. The problem may be as simple as a temp file with a hard coded path name, or the application may not work at all and may require architectural and programming changes to enable it to run correctly on shared servers supporting thin clients. 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.

To avoid encountering these problems during normal business operations, companies considering thin clients must test all such applications in a server-based computing environment before they deploy the thin clients.

The cost of these tests and the costs of any application modifications they prove necessary should factor into the overall thin client cost analysis.

Training

Because thin clients are ideally drop-in replacements for PCs, it's tempting to imagine that no training costs are necessary. Most organizations should plan, however, on some training costs for three different groups:

- users
- support staff
- IT staff

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 pluses 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 who are likely to have used PCs are unlikely to have used thin clients and so may require additional orientation.

User training is unlikely to be a significant cost, but it will be a cost.

Support staff

Training the people who support those users is likely to be a more substantial cost. 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 old first line of defense, "Power down your system," does not work as well 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 administration setup, so support staff members are likely to need specialized training in this area as well. Finding staff with training in a particular type of thin-client administration will generally be harder than finding staff with PC administration experience.

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.

IT staff

The central IT teams that manage the servers and network infrastructure are the groups that will require the most thin-client training. IT network infrastructure analysis and support skills will play key roles, because as we noted earlier, thin client users are only as productive as the network and servers that support them.

The people 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 or 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 current thin clients may run Windows CE, Windows XP embedded, Linux, or a proprietary operating system. The 2005 IDC thin client market analysis predicts in 2009 the market will be split fairly evenly with 25 to 30 percent using each of Windows CE, Windows XP embedded, and Linux; 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.

Software licensing

One area of potential complexity that's easy to overlook is software licensing. On the good side, it's possible that by carefully managing usage an organization can end up needing fewer total licenses for user software than the number of users.

The bad news is that the cost of setting up for thin clients may include 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 particular versions currently in use may not be licensed for servers.

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.

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

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 five 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 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 switching actively among them is a common way to work in most organizations. Each application adds a demand to the host server. Business users have increasingly come to expect that they can perform more than one of these tasks at a time and still experience reasonable performance; organizations deploying thin clients can meet that expectation only by making sure the supporting servers have enough power to run all those applications for all active users.

Intangibles

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

Organizations may well like some of the changes. For example, many people move files between work and home on USB flash drives or CDs. Stopping that may appeal to an organization's security chief, but it will also potentially lower the work people do away from their workstations.

More important in many organizations is 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 equation, they can result in real attitude and thus productivity effects.

The future

One of the hardest costs to estimate in any analysis is opportunity cost. So it is with thin clients. One of the biggest advantages of PCs over the years has been their versatility, their ability to run new types of applications and hardware.

Thin clients, by contrast, are typically extremely limited devices. Doing more on them means either doing more on the server or replacing them.

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 their nature usually not upgradeable. Features they exclude, such as USB ports, which are missing from many thin clients, are often the very things that interesting future technologies, such as USB microphones or cameras for voice or video over IP, 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 high indeed. 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 of the 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. Though decoding media files places a moderate demand on systems, the processing requirements of encoding—which is necessary, for example, for recording audio or transmitting video—can be quite substantial.

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.

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 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 three areas:

- deployment
- moving
- repair

All of these advantages stem from a single feature: the ability to set up a thin client anywhere there's 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 change is to keep at least copies of most, if not all, key data on the server. This has long been possible but has often required appropriate policies and user compliance. Remote PC management, server-enforced policies, 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 significant improvements in this area to continue to appear.

Summary

The cost of adopting thin clients may seem simple to assess, but as this paper has shown, that cost is actually spread across a broader range of hardware, software, and staff areas than may be initially apparent. Organizations adopting or considering moving to thin clients should consider a broad array of potential factors if they want to accurately assess the true cost of thin clients.

Principled Technologies provides 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.

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