



Principled
Technologies®

WHITE PAPER
May, 2006



THIN CLIENTS

Cutting through the hype

Contents

Executive summary	2
Marketing hype	4
Thin client terminology	4
Clear PC advantages.....	5
Thin client territory	9
The many complex other issues.....	10

Executive summary

Intel Corporation commissioned Principled Technologies, Inc. (PT) to analyze and discuss when the competing thin-client and PC platforms are appropriate.

Despite periodic articles and whitepapers claiming that thin clients are the wave of the future, the market share of thin clients remains modest. For reasons intrinsic to the differing natures of PCs and thin clients, thin clients will continue to sell primarily to a few niche markets, while PCs will continue to be appropriate to a far broader range of users.

To understand which niches are appropriate for thin clients and which markets are not, consider three perspectives on thin clients and PCs.

Key areas in which PCs possess advantages over thin clients

PCs have several key inherent advantages over current and proposed thin clients.

Flexibility

The business desktop PC is constantly evolving to include new capabilities. Recent and current examples of such capabilities include:

- Better audio support
- Increasingly better video support
- Support for USB devices (PDAs, cell phones, memory sticks, headphones, etc.)
- The ability to make Internet telephone calls via Voice over Internet Protocol (VoIP)
- Dual monitor support

These capabilities often require not only new software but also new hardware. Though new generations of thin clients can add hardware support for such capabilities, there is typically no way for buyers to add those capabilities to their current thin clients. Instead, they must buy entirely new systems. PCs, by contrast, almost always provide expansion capabilities.

Similarly, just as hardware needs change over time, so, too, do software requirements. These changes can

lead to problems for thin clients, both because many of the latest application developments assume a PC client and because the performance requirements of newer applications can be considerable. Examples of upcoming applications that thin clients may have difficulty supporting include the following:

- VoIP calls (which involve software as well as hardware)
- Microsoft Smart Client support
- Cached Microsoft Exchange support
- Instant messaging

Performance

New business applications bring new and important functions to desktops, but these same applications are often quite demanding. Similarly, new application usage patterns often bring increased performance demands. Both of these types of changes can strain the server processors that thin clients share. Examples include:

- Greatly increased reliance on multitasking applications, as users keep more and more applications active at the same time
- Emerging standards, such as XML, that can consume significant processor resources when you use them for large or complex business documents
- Increased processor demands from applications that involve various types of media, such as graphics, audio encoding, and VoIP

Mobility

More and more business PCs are laptops, untethered mobile devices that let workers function even without a network or when their networks are currently unavailable.

Thin clients, by contrast, are not generally mobile and so do not work well in places, such as airplanes or client offices, that do not have networks or do not open their networks. Those thin clients that are mobile often require significant bandwidth to function well.

Open architecture

One of the biggest attractions of PCs over the years has been the breadth of hardware vendors and choices available to buyers because of the consistently open nature of the PC platform.

The thin client market, on the other hand, is marked by proprietary solutions. The number of



vendors selling substantial quantities of thin clients is fairly small. Having fewer options frequently translates into limited buying and negotiating leverage for customers. Worse, because the different vendors' offerings also come with differing proprietary hardware and software management offerings, buying one vendor's products often amounts to locking into a specific platform. Each such platform typically involves specialized training, consulting, and support costs for IT departments.

Key areas in which thin clients possess advantages over PCs

Thin clients possess some inherent advantages over PCs in a few areas.

Static and undemanding applications

When a desktop user will run only a limited set of well-known programs, and when those programs do not place serious performance demands on the servers supporting the thin clients, thin clients may have enough power to do the job. Users then do not require the added performance or hardware necessary for a more generalized operating environment.

Simple data entry and look-up applications that rely on server databases may be an example of this situation.

Ultra-high-security

Many thin clients offer no way for users to remove data, a limitation that can be an advantage in ultra-high-security workplaces. While PCs can also be configured in this manner, doing so may cost them some of their inherent flexibility advantages.

The many complex other issues

Hype is an apparently unavoidable side-effect of much modern marketing, but in the discussions of thin vs. rich clients it often serves to mask a large array of complex issues that buyers need to consider carefully. In many such areas, one side or the other of this debate may claim clear and inherent superiority, but the reality is unavoidably more complex. Key examples include the following:

Total Cost of Ownership (TCO)

Thin client advocates frequently tout lower TCO as a major inherent advantage of thin clients, but the reality is actually much more complicated.

On the surface, thin clients appear to have the edge over PCs by virtue of being simpler devices that contain fewer and less complicated parts. A common perception is that when a thin client breaks or has problems, you don't need to worry about anything; just replace it. It's also easy to interpret some marketing

hype as meaning thin clients are cheaper initial purchases than PCs.

If you delve deeper, however, you find the truth is nowhere near that simple. To evaluate TCO accurately, you have to consider both the lifetime cost of the thin clients themselves and the cost of the necessary supporting server and network infrastructure. When you include all these costs, thin clients may actually cost more than PCs, especially the new low-priced PCs.

Improvements (both past and future) to PC security and manageability further reduce the support cost advantages of thin clients. Good procedures also affect these costs. For example, by using techniques such as disk imaging you can minimize the repair and user downtime costs of PCs and bring them much more in line with thin clients in these areas.

Initial purchase costs for PCs are also often the lower than those of thin clients.

Manageability

PCs (and PC operating systems) from five or ten years ago were much more difficult to manage than thin clients. Hardware and software improvements, especially those in current and soon-to-be-released PCs, are changing the picture dramatically. Manageability initiatives from companies such as Intel and Microsoft have put PCs and thin clients on much more of an equal footing in this area.

Security

Most security problems are ones where thin clients either have little or no advantage. Even in those areas where thin clients do have an advantage, improvements in the PC platform have made client security almost as strong as that of servers. Both platforms, for example, must deal with the issues of data migration via USB drives and removable media or not offer those capabilities.

Complexity

Managing PCs certainly involves complexity, but thin clients open their own complex issues. Any thin client buyer, for example, will face issues, such as dealing with back-end server processing requirements and providing sufficient bandwidth, that will require significant IT resources and specialized expertise. These resources and expertise will represent significant costs in large organizations and may well be limited or non-existent in smaller ones.

User perception

IT organizations understandably want to minimize their costs and in some cases view maximizing their control of desktop computing as a way of doing so. Users equally understandably want to maximize their productivity. These desires can come into conflict and leave different groups with very different perceptions of thin clients. Thin clients are frequently more attractive to IT organizations than to the users on whose desks they will end up.

Hybrid solutions

In some situations, users are turning to a hybrid answer that uses thin-client support software on PC hardware. By doing this, users get PCs on which they can perform typical current (and future) office functions with good performance. In addition, when users need to run those few applications with either high server processing or high security demands, they do so via thin-client software, typically from Citrix, on their PCs. This approach can provide a winning combination in cases, such as medical settings, in which users need both general-purpose applications and a few tightly controlled, server-based programs.

In the remainder of this white paper, we discuss each of these perspectives in greater detail.

Marketing hype

One of the biggest challenges in comparing PCs and thin clients is pushing aside the marketing hype to get to the truth. Coverage of thin clients by analysts and the press tends to be sporadic and generally uneven. Most research is either vendor-driven (e.g., news about Sun acquiring Tarentella) or on the sensational side (e.g., “Think Thin”). Multiple outlets also seem to cover the topic at about the same time, as if they were working from similar topic calendars. In particular, we found multiple white papers in summer 2004 and multiple stories in summer 2005, most of which portrayed thin clients as The Next Big Thing.

Worse, much of the data seems to echo in the trade press and in vendor claims long after the data has stopped being current. For example, consider a TCO study by Zona Research, Intelliquest, and META Group Study that claimed the following advantages for thin clients:

- 80 percent less maintenance per year
- 25 percent capital cost savings
- 34 percent lower maintenance costs
- 23 percent lower operating costs
- 25 percent increase in user productivity

In online research we found citations of this study in many places, including HP’s thin-client pitch. None of the citations we found, however, noted that this study appeared on April 27, 1999 and is currently extremely difficult to locate. Even the research groups that did the study are in various stages of disappearance or acquisition. Those factors do not prove that the conclusions are wrong, but the study is certainly old. By contrast, up-to-date, unbiased analyses are fairly rare.

Even the claims about the increasing effects of thin clients could benefit greatly from objective analysis and appropriate context. For example, the IDC market analysis report, *Gathering Steam: IDC Enterprise Thin Client Forecast and Analysis, 2004-2009* (by Bob O’Donnell in October 2005) projects 46 percent growth in thin-client sales in 2005 to nearly 2.4 million thin clients, with 1 million of those in the U.S. and Western Europe. The report forecasts that vendors will sell 5.3 million thin clients 2009. Those numbers appear quite large and compelling—until you place them in context. Another IDC study, which eWEEK quoted in its December 20, 2005 story, *PC Shipment Growth to Slow in '06, Report Shows* (http://www.eweek.com/article2/0,1895,1904019,0_0.asp), projects vendors will end up having sold 200 million PCs (desktops and notebooks combined) in 2005.

Finally, in comparing thin clients and PCs it’s vital that buyers compare apples to apples. Both PCs and thin clients are evolving. PC initiatives from Intel, Microsoft, and others are profoundly improving that platform. Newer thin clients offer newer hardware and features aimed at addressing many of their previous weaknesses.

Thin-client terminology

Thin clients come in several different styles, each of which uses a somewhat different architecture to delivering computing power to users. We will use a 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 central server rather than on the 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. 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. Most blade PCs are *dedicated*, meaning one blade PC per user.
- *Thin client* – A thin client is a computer that relies on a server or other back-end system to supply most or all of its computation needs. While these devices often possess a lot more intelligence than most terminals did in years past, their main task is displaying information. Newer thin clients have begun to include the ability to run browsers and some other applications locally. The more they provide such abilities, however, 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 the necessary software.
- *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 underlying software. In this case, that software is the Citrix Presentation Server. As with RDP, the front-end device is typically a thin client and the other system is a server. The front-device, however, could also be a PC running the necessary 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.

Clear PC advantages

PCs have been the mainstay of corporate computers users for quite some time now. As we noted in the

Executive summary, PCs offer key characteristics—flexibility, performance, mobility, and open architecture—that have made them great for corporate users. In the next few subsections, we look in more depth at each of these four characteristics and how requirements for them strongly favor PCs over thin clients.

Flexibility

Business requirements are constantly evolving, and the technology tools supporting those requirements have to change to keep up with them. The business desktop PC is constantly evolving to provide new capabilities that businesses need. Whether these changes necessitate new hardware, new software, or both, the PC's flexibility usually makes it able to handle them.

Hardware flexibility

Most business PCs have room for new hardware that can provide new capabilities. Though new generations of thin clients can add support for those capabilities, existing thin clients typically are not expandable. Adding expansion capabilities is a step toward making thin clients more PC-like; some of the current generation of thin clients have enough bells and whistles to be PCs in everything but name.

Audio capabilities are often useful, and most PCs provide them. Some thin-client users have encountered problems with the lack of audio support of many thin clients and their underlying server-side computing. A slightly humorous example appeared in the Shark Tank column of ComputerWorld on May 17, 2005 (www.computerworld.com/departments/opinions/sharktank/0,4885,101779,00.html). The writer recounted a story of a user whose company had its voice mail system integrated into its e-mail system—an integration with potentially large business benefits. When users received voice-mail messages, they would also receive email notifications with the voice-mail messages attached. This user regularly had other employees call him back but have no idea what was in the voice-mail messages he had left them.

The problem turned out to be that those users were working on thin-client systems. They could see the email notifications about his voice-mail messages, but they could not play those messages. The article quoted the user as saying, "They don't have sound capabilities on their machines, so they can't play their voice messages that way. This has been going on for two years



and no one ever complained. They would read in their e-mail that they had a message from someone and then just call that person." The advantages of the integrated voice-mail and email system were obviously lost on these thin-client users.

Current-generation thin clients often support audio. The models in the Wyse thin-client product line, for example, all include audio support. Unfortunately, not all of the server-side software does a good job of supporting audio. User complaints on newsgroups indicate that Microsoft RDP does not support audio very well. Citrix's ICA protocol, by contrast, does seem to do a good job.

The closed nature of most thin clients means that to give audio support to users who lack it would involve buying new thin clients, not simply adding sound cards and/or speakers to the existing systems.

Audio is a small and rapidly fading issue compared to many of the other hardware developments of the past decade or so.

For example, more and more users employ a range of portable electronics, such as PDAs, USB drives, digital cameras, cell phones, iPods, and so on. Until recently, thin clients not include USB ports, so their users would have no way to connect portable devices to them. Even when the physical USB hardware was present, users often encountered frustrating problems getting the necessary software to work—even when that "fix" was as simple as enabling USB support on the server.

Many current-generation thin clients do include USB ports and have the server support software necessary to allow USB devices to function. Allowing such devices, of course, can undermine the security argument vendors often use to help promote thin clients, so in this case flexibility in thin clients can come into conflict with security.

Video has emerged as an important new application in many companies. Organizations are using video for everything from training to presentations to conferencing. PCs either support video or are flexible enough to make it fairly easy to add such support. The hardware and software limitations of many thin clients, by contrast, combine to make video playback somewhere between difficult and impossible. The basic problem is that to keep bandwidth usage down, support servers send screen updates to thin clients as infrequently as possible. This approach works fine for typical Windows applications, but it does not work as well for images, such as video or Flash animations, that change rapidly. The options are to live with jumpy

images or, when possible, to greatly boost the bandwidth consumption.

The thin-client solution is the usual one vendors propose when new hardware capabilities are necessary: buy a new thin client. Some newer thin clients (such as the Wyse thin clients running with Windows XP embedded or, optionally, with Windows CE) address this problem by supporting Windows Media Player directly on the client.

For many companies, the next big extension of the desktop platform may well be the addition of support for telephone calls via Voice over Internet Protocol (VoIP). The potential savings of VoIP, as well as the potential for better audio quality and the additional functionality possible once calls are all digital, are becoming increasingly difficult for many companies to ignore. PCs may need to add a microphone to make such calls, and newer processor generations are certainly much better at those calls than older ones, but the PC platform can handle them. Thin clients, by contrast, are often out of luck for VoIP. The ability to accept audio input, which is obviously vital if you want to make a call, is one that many thin clients do not have. (Microsoft RDP also does not currently support it.) Thin client vendors certainly can and probably will address these limitations, but the lack of flexibility of the underlying systems means the vendors will probably do so by introducing new models.

Future hardware innovations, such as support for multiple monitors and new types of input devices, will almost certainly continue to arrive first and work best on the hardware of the flexible PC platform. Thin client support will have to wait for new models.

Software flexibility

Just as hardware needs change over time, so, too, do software requirements. While PCs routinely work with newer versions of existing applications and entirely new programs, these changes can lead to problems for thin clients. Crucially, many of the latest application developments assume a PC client. Consider these current and upcoming applications that expect PCs and may pose problems for thin clients:

- VoIP (software as well as hardware needs)
- Microsoft Smart Client
- Cached Microsoft Exchange
- Instant Messaging

We have already discussed VoIP, so we will concentrate here on the others.

Microsoft Smart Client uses XML Web services to take advantage of local processing power on the client PC. By doing so, it can deliver improved performance and functionality over browser-based applications. The reason is that Smart Client applications are not as subject to bottlenecks on the network or Internet servers as browser-based software. Smart Client applications can run application logic on the client so they can provide significant functionality even when they are offline. Microsoft's .NET helps ease the development of such applications. They are likely to become increasingly common.

Microsoft Outlook 2003 clients using Cached Exchange Mode can perform most email-related tasks from the local client. This operating mode improves email performance for items in the local copy of the mailbox by reducing the number of requests to the server for data. This reduction allows the Exchange server to support more users and can significantly reduce the network bandwidth consumption of email flowing between clients and servers. It also lets users keep working when network interruptions occur. These differences can be particularly important for the Outlook experience of two very important groups of users: those working remotely and those located in branch offices.

Instant Messaging (IM) software has changed from something kids run to keep in touch with their friends to an important business communication tool for many organizations of all sizes. IM supplements traditional communications avenues, such as email and telephone calls. IM is a real-time tool with minimal cost and management requirements. The problem thin clients face when trying to support IM is how to run multiple different user IM sessions on the same server; most IM software expects a local PC. As with many of the software issues we discuss, vendors will find workarounds for this one or create new versions to address it. With thin clients, however, these versions will turn IM from a virtually no-management-cost option that users handle on their own into another IT responsibility.

Performance

The software demands of new applications, such as VoIP, lead to the next major advantage for PCs. Applications such as VoIP can require substantial system resources. Though those requirements can pose a problem for older desktop PCs, they can easily become severe issues when a single server must support the computing performance needs of multiple thin clients simultaneously. More generally, new

business applications can bring new and important functions to desktops at the cost of significant processing requirements. Similarly, new application usage patterns, such as the increasingly common work habit of running many applications at the same time, can increase the computing power that users require. These increases will often greatly strain, if not exceed, the computing capabilities of the shared server processors on which thin clients rely.

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—necessary for recording audio or transmitting video, for example—can be quite substantial.

All of these changes are occurring in the context of a greater change in work habits: users are more commonly performing multiple computation tasks simultaneously. These tasks may involve combinations of such things as security-related functions (e.g., virus scanning), communications applications (e.g., email), and general productivity application tasks (e.g., document preparation and formatting). Regardless of the types of tasks involved, business users increasingly have come to expect that they can perform more than one of these tasks at a time and still experience reasonable performance.

In a recent Principled Technologies' 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 latest generation 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.

Mobility

More and more business PCs are laptops, untethered mobile devices that let workers

accomplish significant computing tasks even when they are not attached to a network. Notebooks now often outsell desktop PCs. IDC predicts (ref. *eWEEK*, December 20, 2005 *PC Shipment Growth to Slow in '06, Report Shows*, <http://www.eweek.com/article2/0,1895,1904019,00.asp>) that will be the norm in the U.S. by 2008 or 2009. Mobile systems are also getting increasingly more powerful, with capabilities that rival those of their desktop counterparts.

Thin clients, by contrast, are not generally mobile and so do not work well in places, such as airplanes or client offices, that do not have networks or do not open their networks. Those thin clients that are mobile often require significant bandwidth to function well.

This area is one that many note as one of the biggest problems for thin clients. Consider the following examples.

David Berlind in his ZDNet *Between the Lines* blog dated June 13, 2005 used the title *Notebook - the thorn in the thin client's side* (<http://blogs.zdnet.com/BTL/?p=1495>). In that entry he wrote, "But one cat that the thin client approach has never quite skinned is the mobility cat." Dan Tynan wrote a generally pro-thin-client set of stories in the July 14, 2005 issue of InfoWorld under the general title of *Think thin* (http://www.infoworld.com/infoworld/article/05/07/14/29FEthin_3.html). In one of those pieces, he said, "The Gartner Group's Martin Reynolds sees thin clients as a natural choice for niche markets and companies moving operations offshore, where data security can be an especially thorny issue, but he doesn't see them replacing mainstream business desktop computers. In part because they're less powerful and flexible than PCs, and partly because many businesses are moving away from desktops and toward notebooks, which can pose problems for traveling execs who need to log in to the corporate LAN to access their apps."

One of Tynan's side stories in the *Think thin* collection was *The skinny on thin client mobility* (http://www.infoworld.com/article/05/07/14/29FEthinmobile_1.html?s=feature). In this piece, he went into more depth about the issues thin clients face with mobility. He noted, "True thin clients, by definition, require a constant network connection. To untether a thin client and use it for anything more than an expensive fashion accessory, it must travel within a pervasive wireless cloud." He added, "For traveling executives, mobility and thin client computing make for an uneasy compromise between the power of working offline and the insecurities it brings. The result is frequently a hybrid solution using conventional

hardware." By "conventional hardware" he means, of course, a PC.

Tynan does quote a Wyse executive, who currently carries a notebook PC, as saying that in two years connectivity will be universal and thin clients will be able to function anywhere. We are nearly halfway into that two-year period, and network access, though improving, is certainly far from universal even in the U.S.

One Gartner study even has the rather direct title, *Use Citrix Now for Remote Access, but Not Yet for Mobility*, which pretty much sums up the state of mobility for thin clients.

It is certainly reasonable for thin-client advocates to note that eventually more comprehensive high-speed Internet access, whether from pervasive Wi-Fi or Wi-Max networks or other technologies, will make mobile thin-client devices practical. Business purchasers trying to address the immediate and near-term computing needs of their staff, however, must deal with the world as it is today. Any business notebook user who has tried and failed to get wireless and VPN access to work from an unfamiliar location will attest that being able to work offline is often an essential capability.

Proprietary

One of the biggest attractions of the PC platform over the years has been the breadth of hardware choices and vendors available for it. The open architecture of the PC platform makes it possible for many different vendors to offer a broad variety of products to run on and enrich the platform.

The thin-client market, by contrast, consists of a set of proprietary hardware architectures. Add-on options are thus scarce. Perhaps more importantly, each different vendor's offerings typically come with differing proprietary hardware and software management tools. To deal with these different products, IT departments frequently must incur significant costs for special training and vendor consulting and support.

In addition, because the number of vendors selling substantial numbers of thin clients is fairly small, buyers often have limited negotiating options.

IDC's 2005 thin-client market analysis shows three vendors with over 80 percent of the market:

- Wyse: 43.1 percent
- Neoware: 29.8 percent
- HP: 10.0 percent

The next largest vendor has only 3.3 percent of the market.

The small number of vendors might not be an issue if all were selling products that shared a common basic architecture and could work together easily, but that is not the case.

For one thing, the thin clients from these vendors operate using a wide variety of operating systems, including Windows CE, Windows XP embedded, various flavors of Linux (including proprietary ones), and proprietary ones.

Those vendors also each have their own management software and solutions. Wyse, for example, offers the Wyse Infrastructure Management Software, while HP sells its own Rapport Administrative Software.

Of course, PC management software can also be proprietary, but many general-purpose management products will work with many brands of PCs.

Finally, IT departments that are contemplating supporting a blend of PCs and thin clients face the prospect—and the costs—of having to deal with the multiple management programs necessary to deal with this sort of mixed environment.

Thin-client territory

As we discussed in the Executive summary, while PCs remain the platform most users prefer, thin clients also possess some inherent advantages over PCs in certain environments. We review the two key such environments in more depth in the next two subsections

Static and undemanding applications

Some work environments involve only well-known applications that do not place significant processing demands on the support server and that seldom change. When desktop users run only a limited set of such programs, thin clients offer an advantage over PCs, because users do not have to pay for the added performance and hardware necessary for a more generalized operating environment. Support costs may also be lower for thin clients in such settings, because user work environments will change rarely.

A May 31, 2005 *InfoWorld* article (*HP pushes old thin-client strategy with new devices*: weblog.infoworld.com/article/05/05/31/HNhpthinclient1.html) makes this point while extolling the virtues of thin clients. It quotes Greg Schmidt, a product marketing manager for HP as saying, “[A thin-client

architecture] is recommended for groups of workers who use a limited and predefined set of applications.” The article notes that such groups as call-center staffers, financial analysts, and bank tellers may be good candidates for thin clients.

Similarly, thin clients may be appropriate for certain applications, such as data entry, that have minimal processing requirements. When the full power of a standalone PC is not necessary and the demands on the support servers are modest, thin clients may provide acceptable performance. Many of the thin client case studies in media reports involve applications of this sort. Such pieces often discuss large departments of similar works converting to thin clients. (Few talk about knowledge workers, engineers, or executives switching to thin clients. When those types of users need to run the same server-based applications, they typically access those applications from their PCs.)

An opinion piece by Douglas Schweitzer in the January 9, 2006 issue of *ComputerWorld* (*Thin Is in Again for the New Year*: www.computerworld.com/hardwaretopics/hardware/story/0,10801,107508,00.html) looks in part at using thin clients for simple applications. The article states, “In this scenario, employees using thin-client-connected displays are not truly using “computational” applications—they’re primarily using just a word processing program, although the keyboard and mouse make the experience feel no different than the experience of using a PC.” For such applications, the lower performance and other limitations of thin clients pose no problems.

Ultra-high-security

Many thin clients offer no way for users to remove data, a limitation that can actually be a significant advantage in ultra-high-security workplaces. Such thin clients typically offer no USB support and no removable media of any kind. With such clients, no data ever resides on the client, so even if someone steals the entire client, the organization loses no information.

While it’s certainly possible to configure PCs in the same way, doing so removes some of the inherent flexibility advantages of PCs. It’s also typically easier to obtain thin clients configured this way than PCs.

A different way to address the problem on PCs is via strong security policies: require users to store all data on the server, forbid the removal of data via USB drives, and so on. Enforcing such

policies, however, is obviously much more difficult than with thin clients, where there are no physical options for removing the data. Encryption can also help PCs, and it's essential on laptops and other mobile devices, but thin clients still have a definite advantage for such work environments.

The many complex other issues

Many aspects of the thin client vs. PC debate involve issues on which one side claims a clear superiority but that are actually much more complex. Vendor portrayals and even press coverage of these issues have tended to over-simplify them, because touting large (and frequently largely unsubstantiated) advantages is much easier than exploring issues with multiple complex nuances.

A key reason that many of the apparently simple issues prove to be much more complicated is that technology, both PC and thin-client, does not stand still. In particular, recent and near-term future developments on the PC platform are changing the way buyers need to compare the two approaches.

Total cost of ownership

Total cost of ownership (TCO) is an area that thin-client advocates frequently tout as a key advantage of their platform. The real picture, however, is much more complex.

On the surface, thin clients appear to have a strong advantage over PCs simply by virtue of being simpler devices: simpler means cheaper to buy, fewer parts to break, easier to support, and so on. By centralizing processing on servers, the story goes, thin clients reduce their support costs even further, because IT departments have fewer complex boxes to manage and support.

Dig into this issue at all, however, and many of the claimed TCO advantages of thin clients vanish.

Many thin clients, for example, have a higher initial capital cost than many PCs. This fact is due in part to the decreasing cost of PCs. 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."

When you consider both the cost of the thin clients themselves and the cost of the necessary supporting server and network infrastructure, thin clients may actually cost a great deal more than PCs, especially

the new low-priced PCs. Servers that were adequate to store data that PCs processed often prove inadequate to support the processing needs of thin clients. Similarly, bandwidth that easily handled the file transfers of PC users may not be enough for the constant demands of screen images flowing to thin clients.

The answer to these problems is to boost the server power and bandwidth as necessary to support the thin clients. The cost of doing that, however, is a potentially quite large hidden cost of adopting thin clients.

Necessary improvements often include several obvious requirements: more servers (to handle the processing load), more back-end disk storage (to replace the local storage missing from thin clients), and more bandwidth (to handle the network load). The hidden costs don't stop there, however.

A key less obvious cost is often the price of setting up the necessary server equipment storage areas, which we'll refer to as "equipment closets." Many departments and work groups use networking gear and even a file server that reside in an equipment closet (or other small room) near them. These areas are rarely designed to be miniature data centers. Instead, they're typically just converted small offices or closets (hence the term). The extra servers (or blade PCs and their racks) necessary to support thin clients may often force IT staff to augment the power available in the equipment closet, to improve or replace its HVAC, and to add monitoring for both unauthorized access and potential environmental issues. In many cases, the IT department will find they need to expand the room or relocate the equipment entirely, both of which are expensive options.

For those user groups near an organization's data center, a reasonable option is to move the thin-client support servers and network gear to that data center, but in today's increasingly distributed large organizations, as well as in smaller organizations that lack data centers, many users are simply nowhere near a major data center.

Yet another hidden cost of thin clients stems from their inability to work without the underlying network. If users are working on PCs, they can probably tolerate occasional network issues; thin client users, by contrast, lose the ability to work when the network is not available. The network thus constitutes a single point of failure for all the thin client users on it. Even if the network does not stop entirely and instead simply starts delivering

degraded performance, thin client users may experience poor response times.

Combine all these factors, and it's often the case that the true and total initial cost of deploying thin clients can be significantly higher than that of deploying PCs.

TCO does not, of course, stop with the initial purchase and setup. The cost of supporting and maintaining computing devices over time is often substantially more than the cost of their initial installation. Thin client advocates have traditionally pointed to these costs as being areas in which thin clients have the edge over PCs.

The truth is again more complex. PCs typically have higher costs for user administration, software deployment, and security and virus protection, but good use of techniques such as disk imaging can minimize the repair and user-downtime advantages of thin clients. Thin clients may also have higher server administration and management costs, particularly because operations staff often incur training costs to enable them to support the new platform.

Finally, recent and upcoming changes in the PC platform are eroding the manageability advantages of thin clients, as we'll discuss in the next section.

Manageability

PCs (and PC operating systems) from five or ten years ago were much more difficult to manage than thin clients. Hardware and software improvements, especially those in current and soon-to-be-released PCs, have changed this picture dramatically. Manageability initiatives from companies such 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 edge in three areas:

- deployment
- moving
- repair

All of these advantages are basically aspects of a single feature: the ability to set up a thin client anywhere there's a network connection and access the supporting server or blade PC. 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 are largely interchangeable, because most, if not all, of each user's online state resides on the supporting server or blade PC.

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 and automated data backups are helping in this area.

Another key change is to 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.

In addition, such other emerging management solutions as Intel's Active Management Technology (AMT) and Microsoft's Systems Management Server (SMS) can help simplify the IT management challenge by supporting remote, (and optionally after-hours) maintenance by IT staff nowhere physically near the PCs.

Finally, 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.

Security

Thin clients, by virtue of their closed and limited nature, appear to possess large, inherent security advantages over PCs.

Closer analysis, however, reveals that most security problems that companies face are ones in which thin clients have little or no advantage.

Any such analysis must begin by noting that there is no one security problem. Instead, there are many types of security threats, including:

- physical theft
- software "invaders," from viruses, worms, and Trojan horses to other types of malware and even spam
- unauthorized access
- "phishing" (gaining access via "social engineering" or tricking people)
- theft of services

The CSI/FBI 2005 Computer Crime and Security Survey discusses the leading types of computer security losses. This survey includes the estimated dollar amounts of loss by 639 responding companies. It shows the leading causes of loss to be the following:

- viruses (33 percent)
- unauthorized access (24 percent)
- theft of proprietary information (24 percent)
- denial of service (6 percent)
- insider net abuse (5 percent)
- laptop theft (3 percent)
- financial fraud (2 percent)
- misuse of public Web applications (2 percent)

Thin clients do little to address any of these types of losses. For example, though it's certainly the case that thin clients themselves are not commonly subject to viruses, the computers (servers and/or blade PCs) that support them are. Several of the remaining items on the above list are ones that can affect users regardless of the types of computing systems they use.

The most useful aspect of thin clients in these areas is the ability of IT departments to control them centrally, but such controls are increasingly possible with PCs as well.

In addition, some of the biggest developments in PCs in the last couple of years have related to security, and more important developments are on the way. Both AMD and Intel offer technologies, which Windows XP supports, to prevent a common exploit in which malicious code overruns a buffer and executes the resulting code. AMD's Enhanced Virus Protection and Intel's Execute Disable Bit are included in their latest processors and may significantly help in the battle against viruses. These improvements will also, of course, help servers and blade PCs (and thus aid thin clients indirectly).

Emerging virtualization technologies such as Intel's Virtualization Technology (VT) and AMD's Pacifica offer even more protection. These technologies let systems virtualize hardware, which in turn allows them to contain attacks from viruses and other malware to a single virtual machine on a PC while the others on the PC keep working.

Thin clients appear to have an obvious advantage over PCs in theft of proprietary information and laptop theft, but, again, a closer look reveals a different picture. The theft of information tends not to be from the physical theft of a hard drive. Instead, it is more likely to come via devices such as USB memory sticks or CDs. Both these types of devices are possible with

thin clients, and it's also possible to disable either of them on PCs.

Laptop theft really is unrelated to thin clients; if a company wants to avoid this problem, it has only to refuse to allow users to have laptops. In today's world, however, such a prohibition is almost unthinkable for many users—which, of course, means those users value the portability and disconnected work abilities of laptops and so could not be working on thin clients.

Complexity

Managing PCs certainly involves complexity, as thin-client advocates are quick to point out. Thin clients also open their own sets of complex issues, such as dealing with back-end server processing and managing to provide sufficient bandwidth, that require significant IT resources and expertise.

It's easy to portray thin clients as simple. After all, what are they but simple devices communicating over a network to a supporting server or blade PC? That portrayal, however, masks several types of complexity.

The first lies in the *communication over a network*. The network is everything to thin clients, so it must be completely reliable and adequately fast or thin client users will not be able to do their work. (We looked at the hidden infrastructure costs of the supporting network in the section on TCO.) Network management, like PC management, is always improving, but it's still the case that maintaining network quality is often a costly job that requires significant expertise.

An even bigger area of hidden complexity is in the software arena. Most applications today target the PC platform. Running those applications on thin clients often means running many instances of them on shared servers. Even though both the PC and the server may be running similar versions of Windows XP, applications are often not designed for multi-user environments.

In particular, legacy custom line-of-business applications written for PCs frequently assume that the application has the ability to commandeer the full resources of the PC. This assumption will not be accurate in a multi-user environment.

Consequently, a company will often need to at least port some applications to server-based computing. Other applications may not work at all and may require architectural and programming changes to

enable them to work correctly on shared servers supporting thin clients.

The final type of complexity lies in the thin clients themselves. Despite their apparent simplicity, they are not all the same. Instead, buyers considering thin clients face a wide variety of decisions, many of which will have significant ramifications down the line. 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 far from a given. Various current thin clients run Windows CE, Windows XP embedded, Linux, or a proprietary operating system. Choosing the right operating system is not an easy matter and will require additional analysis. 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, with 12 to 15 percent employing other operating systems, including proprietary ones.

User perception

PC advocates sometimes point to this area as a key advantage for PCs, but the truth is, as with the other topics in this section, more complex. The correct answer often depends on whom you ask.

IT organizations understandably want to minimize their costs and in some cases view maximizing their control as a way of doing so. Users equally understandably want to maximize their productivity. These desires can come into conflict and leave different groups with different perceptions of thin clients, with thin clients often far more attractive to IT organizations than to the users who would have to work on them.

In many, if not most, organizations, it was the users, not the IT departments, who brought in PCs. Many users initially did so in part as a reaction against terminals and central control, though ultimately the productivity of the PC platform carried the day. Users like the flexibility, power, and independence of their desktop and notebook PCs. Users will also not react well to network outages that leave their thin clients (and, consequently, them) dead in the water.

ComputerWorld's Frank Hayes sums up this area well in his January 2, 2006 counterpoint (*Sound-off on Thin Clients: Dead in the Water*:

www.computerworld.com/hardwaretopics/hardware/story/0,10801,107292,00.html) to Mark Hall's pro-thin-client opinion piece. Hayes says, "Remember, bringing in PCs to replace terminals was never IT's idea. Users forced desktop computers on us, starting more than 25

years ago when they smuggled in Apple IIs running something called a spreadsheet. Users have been forcing innovation on us ever since. IT has been fighting it all the way. For users, thin clients are the ultimate IT 'no.' And if we try to force thin clients on users -- sneaking in at midnight to steal their flexible, innovation-oriented PCs and replace them with glorified terminals -- we'll have an all-out war on our hands. That's a war we'll lose. Users make money for the company. We don't. So we have a choice. We can try to sell users on the idea of voluntarily swapping their PCs for thin clients, and good luck to anyone who wants to try. Or we can forget about thin clients, recognize that, after a quarter of a century, we've lost the fight against desktop computers, and focus on a battle we can win: the battle against 'no.'"

Hybrid solutions

Some users are turning to a hybrid answer that involves thin-client support software running on PC hardware.

Consider, for example, a company that has a legacy, high-security application designed to run on a central system. The company needs its users to run that application, and they need the application data to stay secure, so they do not want the program or the data to touch a user's local computing device. (Healthcare organizations, for example, may have such applications.) Those requirements argue for thin clients.

Those users, however, need to be able to do other work, some of which involves applications such as Internet Explorer and Microsoft Office. The users also want to be able to work remotely and offline, often on laptops. Those requirements argue for PCs.

Choosing thin clients would satisfy the company's requirements, but it would cause performance issues for users running office applications on shared servers as both processing load on the servers and data load on the network will degrade performance.

Going with PCs would cause the company problems, in large part because a key security requirement is that no application data reside on a user's system. Even if they allowed the legacy application to run locally, it might well not work well on branch-office systems due to network latency issues; it was designed to run on a central server.



In such cases, a good compromise may be to let the users have their PCs but run thin-client software, such as the Citrix client software, on those PCs when they need to access the legacy application.

The users then get the full power of PCs for their typical office functions. The company gets the security of the legacy application running only on the server. No one pays for the latency issues with the distributed version of the legacy application, because the application runs on the server.

Recognizing the complexity of these many issues and cutting through the hype are key steps in making an informed choice between PCs and thin clients. Though thin clients are certainly appropriate for some niches, PCs will continue to be the dominant computing platform and the right choice for a far broader range of users.

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.
4813 Emperor Blvd., Suite 100
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.