



# Harvard Computing Group Report

## Security and Desktop Client Architectures The road ahead

*A white paper for technology professionals*

This white paper addresses the topic of security and desktop architectures. It focuses on the challenges and decision points that face IT professionals and executives and answers the following questions:

- Are security risks different in various desktop client environments?
- What are the differences between thin and traditional client architectures?
- How do security systems differ between these platforms?
- How does the choice of desktop architecture affect business continuity planning?
- What are the cost controls and effectiveness of thin client security strategies and systems?
- How can secure thin-client networks be managed remotely?
- What issues affect desktop security in VPNs and partner networks?
- How do you integrate biometric and smart card technologies into desktop architectures?

### Executive Summary

Given the challenges of today's marketplace, IT professionals are looking for ways to improve the cost effectiveness of their technology deployments, while ensuring that their data remains secure. However, they are also looking for solutions that will give them a key advantage over their competitors. Part of this advantage is clearly based on identifying ways to use technology and innovative work practices to improve productivity, reduce costs and improve service to employees, partners and consumers. *Security And Desktop Client Architectures* describes the best ways to deploy desktop client technology to gain a competitive edge and addresses the relevant security issues.

### This report contains:

<b>Differences in Desktop Client Architectures</b>	<b>VPNs and Partner Networks</b>
<ul style="list-style-type: none"> <li>• Traditional desktop architectures and examples</li> <li>• Thin-client architectures and examples</li> </ul>	<ul style="list-style-type: none"> <li>• Desktop clients and VPN security</li> </ul>
<b>Client Security</b>	<b>External Security Systems and Thin Clients</b>
<ul style="list-style-type: none"> <li>• Key security issues</li> <li>• Security differences between thin clients and traditional systems</li> <li>• Business continuity</li> <li>• The problems of residual data</li> <li>• Special considerations for public access terminals and wireless devices</li> </ul>	<ul style="list-style-type: none"> <li>• Integrating biometrics and smart cards for validation and authentication</li> <li>• Encryption for wireless devices</li> <li>• Administrative controls and benefits</li> </ul>
<b>Return on Investment</b>	<b>Summary</b>
<ul style="list-style-type: none"> <li>• Protecting the assets of the organization</li> <li>• Remote management and administration</li> <li>• Application lifecycle management</li> </ul>	<ul style="list-style-type: none"> <li>• Key desktop client security issues</li> <li>• A return to centralized control and security</li> </ul>



## Introduction

IT organizations face innumerable challenges today, including:

- reduced staffing
- reduced budgets
- ever-accelerating introduction of new technologies
- increasing demand from users for new features and applications
- increasing, rather than decreasing, choices for desktop computing environments and connectivity
- increasing variety of handheld, portable, away-from-the-office computing devices
- turbulence and instability in the market that raise questions about vendor longevity
- increased hazards of viruses and other malicious threats
- heightened awareness of, and need for, business continuity and disaster planning services.

The goal of this white paper is to focus attention on many of these issues, specifically as they relate to security and the desktop computing environment. We will examine both the technology and business issues that drive decisions for technology buyers and will explain some of the more significant trends and shifts in buying patterns. Readers should gain fresh insight they can apply to the challenges they face in specifying, selecting, buying and implementing desktop computing products, with a particular focus on security implications.

## Methodology

For this white paper, Harvard Computing Group interviewed numerous users of both thin and fat client desktop platforms. Examples from their use in the real world are shown throughout the paper.

We also made extensive use of general research materials, the published works of respected industry analysts and our own extensive experience in the IT industry. Qualitative comments are supported with quantitative results from our interviews and research.

## Differences in Desktop Client Architectures

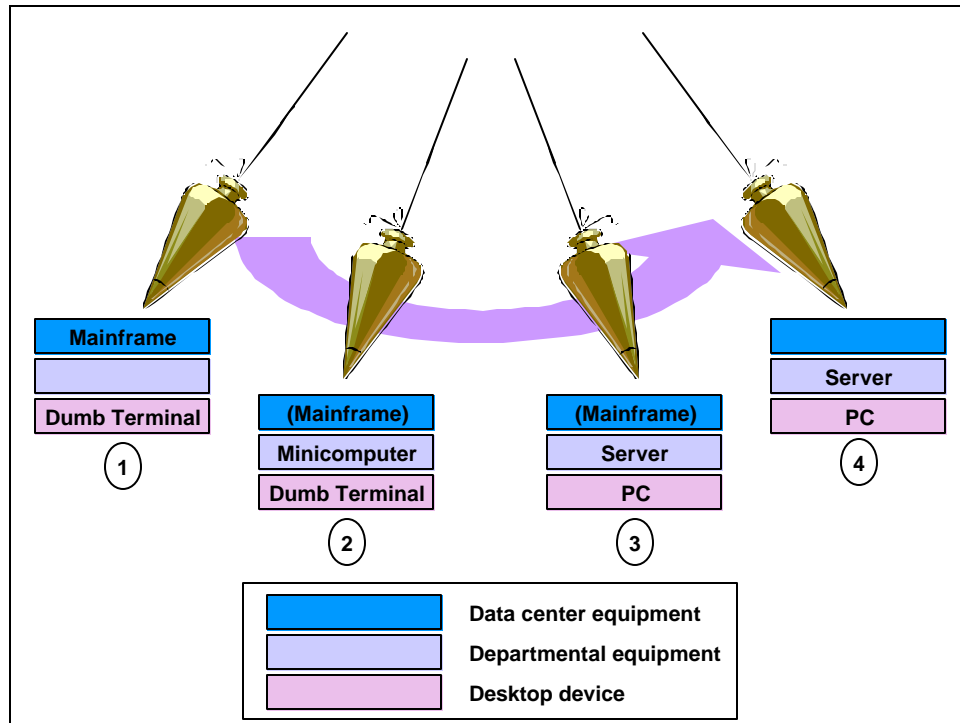
### As the Pendulum Swings

There is a set of continuous, slow pendulum swings that have occurred in the IT industry over the last three decades. Everyone is aware of the general evolution of computing platforms from large, expensive, centrally controlled mainframes to minicomputers and then to PCs and workstations. More recently we have seen a movement back toward the consolidation of computing power into the back office, accompanied by growing use of super-servers and server farms. Though the architectures are radically different, the latter two categories of devices tend to resemble lower cost versions of the mainframes and clustered mainframes of decades past.

There is a parallel pendulum swing occurring on the desktop as illustrated in the seven stages shown in Figure 1 and Figure 2 on the next two pages.



## The Pendulum Swings Away From Centralized Control



*Figure 1 – The swing of the pendulum away from centralized processing to highly distributed processing.*

In the earliest days of computing, shown as Stage 1 in Figure 1, computers were extremely expensive and required highly specialized knowledge to use and to operate. Security was provided almost exclusively by locking up computers in custom-built data centers. Access to data centers was tightly controlled. As dumb terminals were deployed more widely outside the data center, IT departments increased security by requiring users to enter a user ID and password at the terminal.

In Stage 2, minicomputers allowed distribution of some or all of an organization’s computing functions. Security problems were more difficult (though frequently ignored) because many departmental minicomputers sat in the corner of each department with very little physical security. As minicomputers grew in capability, mainframes became redundant in many organizations.

In Stage 3, PCs replaced dumb terminals and PC-based servers replaced minicomputers. Security became even more of a challenge (though it was often still ignored) because the desktop and server operating systems were typically less robust and offered poorer security controls than minicomputer or mainframe operating systems.

In Stage 4, we entered the “Wild, Wild West” of computing – we built business applications based entirely on desktop PCs and departmental servers with little or no central control. Security was minimal at best. The primary motivation for this style of computing was the widespread availability of inexpensive computers and reasonably widespread knowledge of Windows® by users and administrators. Some desktops were UNIX workstations; many servers were based on UNIX and other non-Windows operating systems.

## The Pendulum Returns to Centralized Control

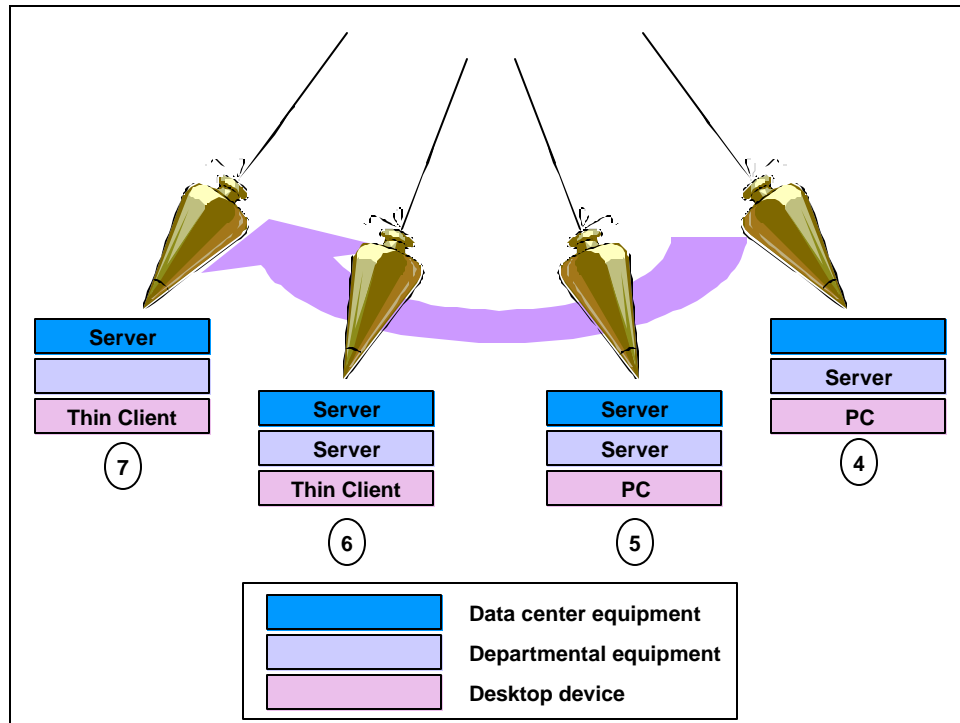


Figure 2 – The return of the pendulum from highly distributed processing back to centralized processing

Stage 5 shows evidence of a return to some centralized control by moving the most critical server functions back to the data center. Taking advantage of centralized control, redundancy and reliability enhances computer and network security throughout the organization.

Stage 6 moves closer toward centralized control by replacing desktop PCs with thin clients. Security is enhanced in many ways that will be described later in this report. For many organizations, however, Stage 6 is just an interim measure. To regain the safety and security of centralized IT that thin clients offer, they must move to Stage 7.

In Stage 7, the pendulum has swung back to a tightly controlled, centralized design in which the majority of computing power resides in secure, reliable, redundant data centers. Desktops consist of thin clients for most applications, with PC clients still used when required for specific applications. IT departments have regained a degree of control and economies of scale that they haven't had since Stage 1 of Figure 1.

Before we discuss thin-client devices and architectures in more detail, we need to review, briefly, the characteristics of traditional desktop architectures.

### Traditional Desktop Architectures

Older desktop architectures fall into two categories: dumb terminals and PCs. In the former case, there is no processing power on the desktop, merely the means to accept input, display (or print) output and to communicate. In the latter case, the PC adds local processing power to those capabilities. Further, the PC contains the double-edged sword of having local software and storage. While software and local disk drives are viewed as a tremendous advantage by users, they are the bane of many IT departments and are

A **dumb terminal** is a computer terminal with three basic capabilities: it can accept input from a user, typically via a keyboard; produce output on a screen or on paper; and communicate to a computer. It has neither local storage nor any significant local processing power.

much of the impetus for interest in thin clients, as described in ensuing sections of this paper.

Despite the argument to be made for thin clients, there are still applications and circumstances when fat clients may be more appropriate. Traveling workers are a prime example. When workers unplug from the local network and need to compute while disconnected, for example in a hotel or on an airplane, they need local storage and software. Without pervasive network connections, a traveling thin client is of little to no value.

A **fat client** is usually a PC or workstation that has local storage for data and software, a local processor and the ability to communicate with both local and remote computers.

Table 1 contains examples of appropriate use of fat-client technology.

### Traditional Client Application Examples

Examples	Why Fat Clients Make Sense	Security issues
Telecommuting	<ul style="list-style-type: none"> <li>• Need access to same files at home and in office</li> <li>• Network access speeds from home are insufficient (except over cable modem and DSL)</li> </ul>	<ul style="list-style-type: none"> <li>• Virus prevention and control</li> <li>• Maintaining integrity of passwords and password access</li> <li>• Maintaining integrity of data</li> </ul>
Traveling workers	<ul style="list-style-type: none"> <li>• Same as for telecommuting, plus</li> <li>• Need to work offline, e.g., in an airplane</li> </ul>	<ul style="list-style-type: none"> <li>• Same as for telecommuting</li> <li>• Notebook theft</li> </ul>
High-performance applications (graphics design, CAD/CAM, financial modeling)	<ul style="list-style-type: none"> <li>• Requires serious computing power, data storage and handling on local machine to perform adequately</li> <li>• Ability to work without network connection</li> </ul>	<ul style="list-style-type: none"> <li>• Same as for telecommuting</li> </ul>
Instrumentation/specialized hardware	<ul style="list-style-type: none"> <li>• PC provides parallel, serial, USB and bus connections for add-on hardware</li> </ul>	<ul style="list-style-type: none"> <li>• External ports could be used by hackers</li> </ul>
Content/knowledge creation	<ul style="list-style-type: none"> <li>• Heavy users of office productivity software or desktop publishing applications may not be satisfied with thin clients</li> </ul>	<ul style="list-style-type: none"> <li>• Same as for telecommuting, plus</li> <li>• Maintaining integrity and correct versions of installed software</li> </ul>

*Table 1 – Fat client application examples and security issues*

### Thin-Client Architectures

In some respects, a thin-client desktop is a return of the pendulum back to the era of dumb terminals, but with several significant differences. First, the underlying technology of a modern thin client bears very little resemblance to the decades-older hardware of a dumb terminal; for example, thin clients support color, graphical user interfaces and sophisticated, high-speed networking capabilities. Second, thin clients, especially when equipped with LCD panel displays are also physically thin – they require far less space and power than older CRT-based dumb terminals. They also generate less heat, which can be very important in large facilities.

A **thin client** is a terminal device that includes a monitor, keyboard and mouse, like a PC or workstation; however, it typically lacks a hard drive, floppy disk, CD-ROM and central processing unit (CPU). Software required to make the thin client function is stored within the device on a read only memory (ROM) chip; however, the majority of user software runs on a server in another location.

Thin terminals employ a graphical user interface but can revert to a character-mode interface for compatibility with legacy applications. Most thin clients do not support attachment of external devices, though higher-end terminals may support add-ons like biometric devices or cash drawers that are connected via parallel ports, serial ports, PCI cards or Universal Serial Bus (USB).

Despite the technological differences, however, a large part of the appeal of thin clients is a return to the data center-based control, administration and software management of the mainframe/dumb terminal era. As much of the remainder of this paper will detail, the cost, security and administrative advantages of thin client computing make a compelling business and technology case. The advantages are sufficiently compelling that market research firm IDC projects thin client terminal sales to reach nearly nine million per year in 2005, up from



approximately 1.3 million in 2001.

Table 2 below includes examples of appropriate use of thin client technology.

### Thin Client Application Examples

Examples	Why Thin Clients Make Sense	Security issues
Office workers	<ul style="list-style-type: none"> <li>• Dramatically enhanced reliability</li> <li>• Central control of desktop environment</li> <li>• Ease of installing frequent updates to web browser and office productivity applications</li> </ul>	<ul style="list-style-type: none"> <li>• Users cannot load software onto device</li> <li>• Central control of user environment</li> <li>• No document storage on local device, ensuring proper access control and backup</li> </ul>
School classrooms and computer labs	<ul style="list-style-type: none"> <li>• Dramatically enhanced reliability</li> <li>• Significantly reduced opportunity for students to misuse or abuse school-provided software</li> </ul>	<ul style="list-style-type: none"> <li>• Better control over student use</li> <li>• Users cannot load software onto device</li> </ul>
Data entry	<ul style="list-style-type: none"> <li>• Higher reliability</li> <li>• Lower cost</li> </ul>	<ul style="list-style-type: none"> <li>• Central control of user environment</li> </ul>
Factory floor applications	<ul style="list-style-type: none"> <li>• No moving parts or fans that might deteriorate in dirty environment</li> <li>• Increased reliability and ease of use</li> </ul>	<ul style="list-style-type: none"> <li>• Central control of user environment</li> </ul>
Exam room, nursing station and public access terminals in healthcare facilities	<ul style="list-style-type: none"> <li>• Quiet operation (no moving parts or fans)</li> <li>• Slim form factor</li> </ul>	<ul style="list-style-type: none"> <li>• Central control of user environment</li> <li>• No data storage on terminal, reducing risk of inappropriate access to confidential data</li> </ul>
Kiosks	<ul style="list-style-type: none"> <li>• Near-zero maintenance</li> <li>• Total control and administration from central location</li> <li>• Light weight and small size</li> </ul>	<ul style="list-style-type: none"> <li>• Reduced risk of hacking by casual users</li> <li>• Reduced risk of damage to more expensive hardware</li> </ul>
Telecommuting	<ul style="list-style-type: none"> <li>• Increased reliability of hardware</li> <li>• No software installation or upgrades</li> <li>• More effective security</li> <li>• More effective use of low-speed, dialup lines when that is all that is available</li> </ul>	<ul style="list-style-type: none"> <li>• Central control of user environment</li> <li>• No data storage on terminal, reducing risk of inappropriate access to confidential data</li> </ul>

Table 2 – Thin client application examples and security issues

## Client Security

### Key Security Issues

The security issues for all types of client devices are widely varied and can include any or all of the items shown in Table 3 below.

Security Issue	Examples
Control of physical access to the device	<ul style="list-style-type: none"> <li>• Room security via conventional door locks, card key access, biometric devices</li> <li>• Keyboard locks</li> </ul>
Control of logical access to the device	<ul style="list-style-type: none"> <li>• User ID/password</li> <li>• Card key</li> <li>• Biometric devices</li> <li>• Device timeouts</li> </ul>
Control of the software environment on the desktop	<ul style="list-style-type: none"> <li>• Policy software</li> <li>• Operating system "lock down" features</li> <li>• Use of devices without local software</li> </ul>



Security Issue (cont'd.)	Examples
Application access restrictions	<ul style="list-style-type: none"> <li>• User ID/password</li> <li>• Card key</li> <li>• Biometric devices</li> <li>• Application timeouts</li> </ul>
Data security on the client	<ul style="list-style-type: none"> <li>• Encrypted files on hard drive</li> <li>• Use of devices without local storage</li> </ul>
Data security on the network	<ul style="list-style-type: none"> <li>• Encryption</li> </ul>
Viruses	<ul style="list-style-type: none"> <li>• Anti-virus software</li> <li>• Use of devices without diskette or CD-ROM drives</li> </ul>

Table 3 – Client security issues

Table 4 highlights the security differences between fat and thin clients for each of the security issues listed in Table 3.

### Security Differences Between Thin-Client and Traditional Systems

Security Issue	Fat Client	Thin Client
Control of physical access to the device	<ul style="list-style-type: none"> <li>• No significant difference</li> </ul>	<ul style="list-style-type: none"> <li>• No significant difference</li> </ul>
Control of logical access to the device	<ul style="list-style-type: none"> <li>• User ID/password very common</li> <li>• Inactivity timeouts not very common</li> <li>• Biometric devices and card keys can be attached but may be difficult to configure depending on operating system and firmware compatibility</li> </ul>	<ul style="list-style-type: none"> <li>• User ID/password very common</li> <li>• Inactivity timeouts very common</li> <li>• Low-end thin clients may not support hardware or software required for attachment and use of external devices; high-end thin clients certainly do provide this support</li> <li>• For suitably equipped thin clients, attaching biometric devices can be as simple as plugging them in</li> </ul>
Control of the software environment on the desktop	<ul style="list-style-type: none"> <li>• Many tools from many vendors, but assembling a complete solution is daunting</li> <li>• Vast majority of installed PCs are completely unmanaged</li> </ul>	<ul style="list-style-type: none"> <li>• Huge advantage of thin clients results from lack of hard drive; the software environment is totally controlled from a central site</li> </ul>
Application access restrictions	<ul style="list-style-type: none"> <li>• See notes above regarding user ID/password, biometric devices and inactivity timeouts</li> </ul>	<ul style="list-style-type: none"> <li>• See notes above regarding user ID/password, biometric devices and inactivity timeouts</li> </ul>
Data security on the client	<ul style="list-style-type: none"> <li>• Local hard disk and memory cache present enormous potential for data compromise</li> <li>• Diskette and CD-ROM provide opportunity for introduction or removal of data or software</li> </ul>	<ul style="list-style-type: none"> <li>• Lack of hard drive removes the threat of unintended data access</li> <li>• In some situations, memory cache must be flushed periodically to ensure against unintended access to residual data</li> <li>• Lack of diskette and CD-ROM eliminates the threat</li> </ul>
Data security on the network	<ul style="list-style-type: none"> <li>• PC traffic via TCP/IP is not encrypted without the use of additional software</li> </ul>	<ul style="list-style-type: none"> <li>• Network traffic is typically encrypted within the ICA / RDP<sup>1</sup> protocols</li> <li>• The ICA protocol supports up to 128-bit encryption</li> </ul>

<sup>1</sup> Citrix' Independent Computing Architecture (ICA) protocol and Microsoft's Remote Desktop Protocol (RDP) manage communications between thin clients and terminal servers.



Security Issue (cont'd.)	Fat Client	Thin Client
Viruses	<ul style="list-style-type: none"> <li>• Strong threat of virus introduction via diskette or CD-ROM</li> <li>• Anti-virus software and virus definitions must be continually updated on every client</li> <li>• Virus propagation via local hard drive</li> </ul>	<ul style="list-style-type: none"> <li>• No threat of virus introduction via diskette or CD-ROM</li> <li>• Anti-virus software and virus definitions maintained on server</li> <li>• Virus propagation minimized without local hard drive</li> </ul>

Table 4 – Security differences between fat and thin clients

## Business Continuity

### Survivability

Backup, redundancy and survivability have always been key issues for certain industry and government sectors. However, the terrorist events of September 2001 have brought these concerns to the forefront for a vast number of organizations. How does an organization deliver uninterrupted services in the face of minor disruptions like power outages or local floods? How can those same organizations react to catastrophes like the destruction of a data center or major communication network outages?

The challenges are multiplied for organizations that rely on PCs to host and execute applications and to store data. It is not hard to imagine that 50% or more of an organization’s critical data might be lost if all of the PCs in an entire facility were destroyed.

Thin clients provide several advantages for business continuity, stemming in large part from the server-based storage of data and software. It is literally possible to interrupt an application session on a thin client – whether intentionally or not – and resume that session from an entirely different physical location on another thin client. Furthermore, the application session will resume in exactly the same place and with the same data on the screen. Because a thin client screen merely displays the results of application processing that happens at a central site, the physical location of the thin client is largely irrelevant.

Dr. Scott Barrett, Director of Technology for the Conroe (TX) Independent School District, summed up one of the most compelling arguments for use of thin clients. After a very positive experience deploying more than 1200 Wyse thin clients in classrooms and computer labs, he began rolling out thin clients to administrators and teachers. Many long-time PC users complained at first, until Barrett said to them, “Would you rather have something that works all the time or something that works part of the time?”

The flip side of relying on data centers to provide business continuity is that those data centers must be fully equipped to deal with both remote and central site problems. The good news is that all of the major server vendors provide high reliability and fail-over options in their current product lines.

### Managing Upgrades

Today’s browser-based applications are vulnerable to bugs in the underlying browser and web server software on which they rely. The history of Internet software in the last decade is rife with examples of bugs, exploits and attacks against web browsers and servers. The result has been a seemingly endless stream of patches and upgrades designed to protect users from newly discovered vulnerabilities.

Stephen Hart of Dreyer Medical Clinic was in the middle of rolling out an upgraded version of Microsoft Internet Explorer to a nearly equal number of PCs and Wyse thin clients when he was interviewed for this report. He had just spent three days writing the script to perform the installation on each of the 500 PC clients and, despite his preparations, fully expects to encounter a variety of problems while performing the upgrades over the next several weeks. By contrast he expects it will take him about 10 minutes to apply the upgrade to the servers that drive 500 thin clients.

Traditional client installations are faced with a never-ending challenge of installing upgrades and patches on every desktop. Thin client installations have an enormous advantage because they only need to install patches at the server(s) and all



clients are immediately upgraded.

The software upgrade problem is not limited to browser-based applications. All programs require fixes and updates as described in the section titled, Application Lifecycle Management on page 12.

## The Problems of Residual Data

One desktop client computer problem that is often overlooked is that of residual data – data that remains in memory or on disk after users complete their tasks. For generic applications, residual data may not present any problems but for applications involving confidential data or situations in which unintended personnel may have access to client devices, the problem can be significant.

Fat clients present a potentially larger problem because data files are often stored on the local hard drive. Even when all data for PC-based applications is stored on a central server, there are often transaction files or other temporary files that are written to the local hard drive to enhance user performance. Well-written applications will delete the temporary files, but a knowledgeable hacker can access even the deleted files relatively easily. Thin clients avoid these issues entirely by not having local hard drives.

Web browser-based applications, which represent a growing number of all remotely deployed applications, present a similar problem because web browsers typically keep a cache of recently accessed pages. Once again, a thin client eliminates part of the problem by not having an internal disk drive. However, local web browsers (embedded in the thin client) also keep a memory-resident cache of the most recent pages. Unless explicit action is taken, data in the memory cache remains accessible until the web browser is closed, and may still be accessible after the browser is closed. Thin clients with local browsers are not exempt from this problem because they do contain local memory. Thin clients that use a server-based browser do not have this problem, and some thin clients may have local browser security settings that eliminate the local cache and thereby eliminate the risk.

Another way to solve this problem requires software that explicitly flushes data from the memory-resident cache. Atlanta-based StayOnline ([www.stayonline.net](http://www.stayonline.net)) sells in-room Internet access solutions for the hospitality industry and will have more than 5000 Wyse® ([www.wyse.com](http://www.wyse.com)) Winterm™ thin clients in approximately 200 hotels throughout the United States by the end of 2002. Because of their concerns that one hotel guest might inadvertently (or intentionally) view web pages from a previous guest, the company uses Wyse's Rapport™ administrative software to purge the memory cache in all of the thin clients every night at 3:00 AM. They have even developed a method to purge instant messenger-style buddy lists to enhance the security of their clients' guests.

**StayOnline**  
*Stay Connected...StayOnline*

*Figure 2 – StayOnline uses thin clients to provide reliable Internet access to hotel guests*

## Special Considerations for Public Access Terminals

Kiosks and other public access terminals are generally very good candidates for thin clients because they are seldom used for highly demanding, processor-intensive applications. In addition, they are designed for casual users, not power users. Public access applications generally require only minimal data input and present graphical output to the user. In only a few cases public access terminal may require

### KEY QUESTIONS – PUBLIC ACCESS

- Can I afford to send someone to repair a failed public access terminal or would I prefer to control it from a central site?
- Do I need local storage?
- Do I need special hardware that only a PC can provide?
- Are viruses or the threat of hacking an expected problem?

special hardware attachments or software that would require use of a PC or workstation.

The case for thin clients as public access terminals is even stronger because of their other attributes. First, because they are in public locations, they may be more susceptible to physical attacks and abuse, so a device that can be maintained and repaired easily and cheaply will provide considerable advantage. Second, kiosks, especially in the retail world, may be relocated frequently, so a device that is lightweight, small and easily portable is desirable. Third, eliminating local software and being able to manage public access terminals from a central location provides significant labor savings. Finally, thin clients eliminate concerns about PC-knowledgeable users hacking into the device for nefarious purposes.

## Special Considerations for Wireless Devices

Although the majority of thin clients today are installed on wired LANs, the number of wireless devices is growing. The Health Services Department of a large Midwestern county, for example, is building a new hospital that will open in the second half of 2002. While the majority of their thin clients will be wired, they have installed 150 wireless access points and will make extensive use of wireless thin clients on movable carts. They will also support wireless tablets for use by health care providers.

Despite the obvious labor savings for installing and relocating wireless devices, there are inherent security risks as well, primarily because all data to and from the terminal travels through the air. It is only a minor challenge for someone with a bit of knowledge and suitable hardware (e.g., a wireless protocol analyzer or a laptop equipped with a wireless access card) to intercept vital transmissions.

The organizations surveyed for this report are at both ends of the spectrum on their concern about the security of wireless transmission. The Conroe (Texas) Independent School District, for example, uses wireless laptop carts that are rolled from classroom to classroom. They are not using encryption but are not worried about snooping because they've checked that the radius of wireless transmissions is within school boundaries.

Ensuring that wireless transmissions remain within facility boundaries is one aspect of the wireless interception problem. Organizations must also consider that an intruder could be an employee or member of the community who has permission to enter the premises, but does so with a wireless-equipped laptop or thin client. While such a person would still require a user ID and password to access network resources, the potential does exist for snooping or surreptitious capture of network data.

The Midwestern county hospital referred to above, on the other hand, is already deploying encryption and VPN technology for the roving thin clients in their new building. They decided early on that the Wired Equivalent Privacy protocol (part of the IEEE 802.11b wireless networking standard; see *Encryption for Wireless Devices* on page 14) was not sufficient. They have implemented more rigorous encryption technology to ensure patient confidentiality.

### KEY QUESTIONS – WIRELESS

- Is the broadcast radius of my wireless LAN fully contained within my facility? What if some of my wireless devices are located on the outside edge of a building near the boundary of my facility?
- What are the legal ramifications of *not* using encryption?
- Is the Wired Equivalent Protocol (WEP) used in today's 802.11b wireless LANs sufficient protection for my data?
- Do today's widely available 802.11b wireless devices provide adequate speed or do I need to wait for 802.11a or 802.11g?



## Return on Investment

Reducing Total Cost of Ownership (TCO) alone is a compelling reason to move to thin clients for nearly all organizations. For a study group of 2500 clients and 35 servers, Gartner Group reports a thin client TCO of at least 25% less than unmanaged or poorly managed PCs<sup>2</sup>. The payback period for rolling out the sample set of 2500 thin clients, even taking into account the cost for new servers and server-side software licenses, is approximately three months.

Gartner also reports, however, that the TCO differential between thin clients and extremely well-managed PCs may be negligible. Well-managed PCs, according to Gartner, are ones that are locked down to prevent software installation or changes by users and are managed with a set of tools that allow central administration. In a real sense, “well-managed PCs” by their definition, are little more than elaborate thin clients. It is important to realize, however, that Gartner’s well-managed PCs still contain moving parts and, despite best efforts, may not be locked-down as securely as network administrators would like them to be.

## Protecting the Assets of the Organization

To a growing extent, all organizations are driven by the information contained within their computer systems. Consequently, they are placing increasing emphasis on securing it, while at the same time, making it accessible whenever and wherever it is needed.

The Conroe Independent School District sees one small but important advantage of thin clients in the elimination of floppy disk drives. Staff now use thin clients from home instead of carrying disks back and forth. Not only does this eliminate the risk of staff members carrying virus infections back and forth, but it also ensures that data on the servers is always current. Standard data center backup procedures also assure complete backups.

## Remote Management and Administration

One of the key benefits of thin clients is that all maintenance and control functions are centralized. While this requires a greater investment at the central site for redundancy, fail-over and staffing, the near-zero costs of remote maintenance afforded by thin clients makes the case quite compelling, especially as the number of remote locations containing client devices increases. Figure 3 below illustrates the relative cost advantage of thin clients as the number of locations increases (from left to right in the illustration).

**“Spending \$1000 on a desktop improves the situation for a single user, but spending \$1000 on a server helps everyone.”**

Matthew Harris, programmer for a California university research laboratory

In a dramatic example of tangible annual savings from deploying thin clients, a California university research lab has reduced their annual maintenance budget since converting to thin clients. In the past, they allocated an amount equal to 10% of their annual PC purchase budget for maintenance. Now, their annual thin client maintenance budget is 2% of their expenditures on new systems.

<sup>2</sup> Lowber, Peter, *Thin-Client vs. Fat-Client TCO*, 28 September 2001, The Gartner Group, Note Number DF-14-2800.



## Relative Maintenance Costs of Fat vs. Thin Clients

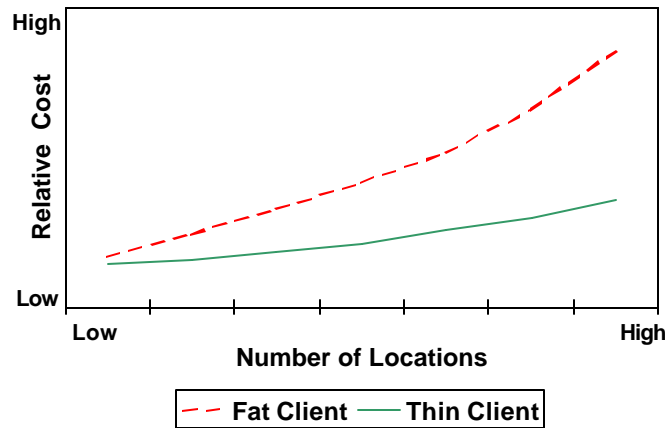


Figure 3 – Relative costs for maintaining fat vs. thin clients as the number of locations increases

While the mean time between failures (MTBF) for current generation PCs may be quite good, the MTBF for thin clients is five to 10 times higher. Most of the added reliability stems from two primary hardware and software factors:

- Elimination of moving parts (which are the most likely to fail)
- Elimination of complex operating system and application software.

Most organizations interviewed for this paper keep a small stock of replacement thin clients on hand for those very rare cases when one fails. Dreyer Medical, in Aurora, Illinois, maintains an inventory equal to about 4% of their thin client population for this purpose.

The Mean Time to Repair (MTTR) for thin clients is also superior to that for PCs for most types of failures. In the worst-case scenario – when the entire device must be replaced – a new thin client can be plugged in and running in less than five minutes. Installing the hardware and configuring the software for a replacement PC could take hours.

### Application Lifecycle Management

The process by which organizations develop, test and deploy new applications can be extremely time consuming and labor intensive. In a traditional fat client environment, the problem is exacerbated by the sheer variety of PC configurations and operating systems that must be tested to ensure that applications function properly on every desktop configuration.

In addition, every new or revised software version – and there may be many in the development cycle for a major application – may require an IT staff person to visit each desktop. Even if an organization uses software management tools that allow them to push changes out to desktop computers, variations in desktop PCs may force IT staff to visit a considerable number of machines for every software update.

A final challenge in the traditional client environment is that software upgrades in large, distributed environments often need to be scheduled over an extended period to ensure sufficient IT staff time to accomplish the upgrades. As a result, users throughout the organization may end up running different software versions for some period of time. During this time, there may be a

loss in productivity because of the questions or problems that arise from the software differences. In addition, this can put an added burden on the support department, further slowing the deployment throughout the organization.

While thin clients don't eliminate all of these problems, they certainly contribute to a vastly smoother and easier application deployment. First, new versions of applications only need to be installed once on the server. Second, assuming an application does not need to take advantage of specialized hardware features on specific thin client models, if the application works on one thin client, it will work on all of them. Finally, because software installed on the server is automatically running on every thin client, the organization does not have to deal with multiple software versions in use at one time.

## VPNs and Partner Networks

The successful use of thin-client devices is dependent upon having sufficient bandwidth between the client and the server. This is very seldom an issue today for LANs that typically run at 10 or 100 million bits per second (Mbps). On wide area network links, or on virtual private networks (VPNs) running over the Internet, available bandwidth is a consideration for deploying thin clients.

The good news, however, is that thin clients often require less bandwidth than fat clients because they are not passing large quantities of data back and forth. Only keystrokes and mouse clicks travel to the server; just screen updates pass in the reverse direction. Consequently, thin clients can be used very effectively for remote access VPN connections, e.g., telecommuters, employees working from customer premises and traveling employees. In addition, thin clients can support partner networks by allowing business partners to access server-based applications through an Extranet.

When used in remote or partner applications, one attribute of thin clients discussed several times in this paper really stands out – the ability to deploy an application once on the server for rapid use on every thin client. It is not practical, or in many cases even possible, to install software on a remote user's or partner's PC. With thin clients, there is no need to do so. Merely connecting to the server through the VPN or partner network accomplishes the same thing.

## External Security Systems and Thin Clients

Password authentication mechanisms are fine in an ideal world, one in which people create easily remembered but highly unique passwords of sufficient length and complexity to prevent others from successfully guessing them or using software to crack them. Unfortunately for organizations that rely on password security, we live in the real world in which people routinely use birth dates or spouse's names or children's names or even repeat their user name as their password. It's not uncommon to find passwords written on PostIt<sup>tm</sup> notes attached to computers.

Doctors and nurses at one healthcare provider interviewed for this report are required to logon to a terminal in each exam room they enter – something they do dozens of times per day. Consequently, many have fallen into the bad habit of using their initials as their password, a particularly egregious violation of password security policy. The organization, on the other hand, feels their hands are tied – if they enforce more complex passwords then the employees may become too frustrated and stop using the system altogether.

Heightened security awareness has caused many organizations to look beyond simple passwords to unique devices or physical attributes of users. The card keys that many organizations use for access to secure rooms or buildings can be applied to computer terminals and applications.



Biometric devices that scan a body part or use voice recognition are increasingly being adapted for application- and device-level security.

Fortunately for organizations considering these technologies now, heightened interest in the use of biometric security devices coincides with recent reductions in prices for the required hardware and software. Recent prices show costs of tens of dollars, rather than hundreds or thousands of dollars per user for smart card and biometric security systems. Given that providing assistance to users who have lost or forgotten their passwords can consume an inordinate amount of time and expense for corporate help desks, it doesn't take very many calls per user per year to justify the added expense of biometric security. And on top of the potential for overall cost savings, an organization gains significantly enhanced security through use of smart cards or biometric devices.

### Integrating Biometrics and Smart Cards



Smart cards are credit card-sized devices that contain a microprocessor and memory, and can store personal or other information. The user inserts the smart card into a reader that uses the information contained on the card to identify the holder. To prevent use of stolen smart cards, users are often required to identify themselves in another way – perhaps by entering a password or using a form of biometric identification as described below. Figure 4 shows a sample smart card.

Figure 4 – Smart card from a division of Schlumberger ([www.schlumbergersema.com](http://www.schlumbergersema.com))

Biometric devices consist of specialized hardware and software that scan fingerprints, hands, faces, or the iris or retina of the eye to compare the proffered body part against a mathematical representation in a database. Each type of biometric device includes dozens or hundreds of separate measurement points against which it makes comparisons. The device in the foreground of Figure 5 is a fingerprint reader from Identix ([www.identix.com](http://www.identix.com)).



Figure 5 – Fingerprint reader

Today, only a limited number of thin-client devices support connection of smart card and biometric devices, in large part because many thin clients lack serial or parallel ports by design. However, thin client manufacturers are adding USB ports and providing supporting software to their devices to allow connection of a broader array of external devices. The need for more rigorous, positive identification of users will drive vendors to meet this need.

### Encryption for Wireless Devices

The classic dilemma for the use of encryption technology is that it is compute-intensive, i.e., it consumes a large number of processor cycles to accomplish the encryption of data. Consequently, many organizations decide not to encrypt data to avoid the risk of degraded application or network performance. Or they employ minimal encryption techniques that cause only minor degradation, if any. Unfortunately, the more modest the encryption, the higher the risk of a security breach. Fortunately, the ICA protocol embedded on most thin-client devices can encrypt thin-client data streams without any noticeable impact on performance because the

underlying protocol places minimal resource requirements on the device.

Wireless LAN products that conform to the 802.11b standard, which encompasses most of the products currently on the market, offer an optional encryption technology known as WEP (Wired Equivalent Privacy). Organizations using thin clients with wireless LANs must explicitly turn on WEP to gain the added benefit of encryption. However, WEP is not foolproof and was publicly cracked in 2001. Hacker tools that can crack any WEP-protected system are readily available on the Internet. This doesn't mean that an organization shouldn't use WEP – doing so is more secure than not using it – but WEP should not be the exclusive means of protection when data confidentiality is a primary concern.

The 802.11 standards committee, along with several vendors and industry consortia, are working diligently to augment or replace WEP.

## Summary

### Key Desktop Security Issues

Organizations deploying new client hardware must consider at least the following security issues:

- How can we protect the confidentiality of our data most effectively?
- How can we deploy local and remote client devices at the lowest acquisition and maintenance costs while still making them secure?
- How can we ensure the protection of data as it travels through our networks?

While, technology selections are almost never simple choices, the answer to a growing number of client deployment questions is to install thin client terminals.

- The economics of thin clients are very favorable.
- They run all of the applications required by task and office productivity workers.
- Ease of administration and maintenance is excellent.
- With appropriate data center procedures, data is highly secure, always backed up and always available.
- LAN-connected terminals inherently provide transmission security; suitable encryption technology exists to protect wireless transmissions.

### A Return to Centralized Control and Security

The pendulum at enlightened organizations has indeed swung away from the “Wild, Wild West” of haphazard PC deployments toward a time when the critical value of an organization's information assets are paramount. Protecting and securing those assets can be done extremely effectively with thin clients. As a final thought, the following quotation from the website Thin Planet summarizes the position of thin clients in today's world:

“In the end, thin-client computing combines the security and control of a mainframe with the interface and function of a PC. As IT organizations seek to stay lean and mean, it seems inevitable that many will cut out the fat in favor of a thin-client solution.”<sup>3</sup>

<sup>3</sup> Thin Planet, Common Sense About Thin Clients, [www.thinplanet.com/tech/generic.asp?f=TDnumber&k=s&v=TD34202](http://www.thinplanet.com/tech/generic.asp?f=TDnumber&k=s&v=TD34202)



The Harvard Computing Group (HCG) provides white papers on technology subjects to the public through its store at [www.harvardcomputing.com/Store](http://www.harvardcomputing.com/Store).

HCG also delivers consulting services, corporate training programs, and seminars in the United States, Europe, and Asia. Contact information is provided below or email us at: [info@harvardcomputing.com](mailto:info@harvardcomputing.com)

**Harvard Computing Group, Inc.**

Regency Park Suite 103  
238 Littleton Road  
Westford, MA 01886 USA  
Phone: +1 978-692-6766  
Fax: +1 978-692-1864  
Email: [info@harvardcomputing.com](mailto:info@harvardcomputing.com)

**Harvard Computing Inc. Europe**

Golden Cross House  
8 Duncannon Street  
London, WC2N 4JF, UK  
Phone: +44 (0)20 7484 5084  
Fax: +44 (0)20 7484 5100  
Email: [inquiries@harvardcomputing.com](mailto:inquiries@harvardcomputing.com)

**Harvard Computing Group Asia**

103-104 E. Rodriguez Avenue  
Ugong, Pasig 1100  
Philippines  
Phone: +63 2 671 2785  
Fax: +63 2 633 1146  
Email: [mailto:philippines@harvardcomputing.com](mailto:mailto:philippines@harvardcomputing.com)

**Harvard Computing Group Thailand**

23<sup>rd</sup> Floor CP Tower  
313 Silom Road  
Bangkok 10500, Thailand  
Phone: +66 2 231 8045  
Fax: +66 2 231 8121  
Email: [tantraporn@harvardcomputing.com](mailto:tantraporn@harvardcomputing.com)

