newburn

White Paper on

Thin Client Benefits

Introduction

Businesses are growing, streamlining, decentralising, centralising, out-sourcing, in-sourcing, merging and de-merging more rapidly than at any time in the past. Central to many of their concerns is, "how do we make our IT work together, cost-effectively? Do we run disparate operating systems with associated costs, or bite the bullet and integrate/replace them for a significant up-front cost? Do we have to go through this again with the next change?". Desktop builds, hardware and software, will almost certainly be different, as will WAN service providers.

One way of reducing the costs associated with these inevitable changes is through thin-client computing. Thin-client computing is not a new concept. However, it has not been widely promoted, since major PC vendors have understandably been reluctant to promote a concept that will impact their highly profitable PC-upgrade spiral. Vendors are now realising that they cannot stop the word spreading and are embracing the technology. Thin-client computing is not the run-from-server installations sometimes used by major Windows applications. That method still installs several files on the local PC and requires signficant bandwidth to operate at a useable speed.

Application service provision (ASP) is a high-profile flavour of thin-client computing. The technical solution is very similar, it is the licensing scheme and the out-sourced nature of the provider that it different. ASP has yet to come of age in the UK, primarily owing to software rental licensing agreements. However, this technology looks set to grow. By adopting a thin-client computing solution now, you benefit from the cost savings now, but are in a position to transition to ASP very quickly if desired.

Glossary

The industry uses a number of terms to describe thin-client computing, with some of them being apparently, but incorrectly, interchangeable. For the purposes of this paper, I will use the following:

- thin-client computing - running applications from the server and distributing only keyboard, video and mouse (KVM) updates and possibly audio;
- terminal servers - the application servers at the heart of thin-client computing. The software is normally from Microsoft, although other suppliers are making inroads, they are often deployed in conjunction with Citrix or other functionality enhancing products;
- fat-clients - PCs used in a traditional manner, with applications stored on the PC and data stored on servers;
- thin-clients - dedicated, solid-state devices, providing connections into the thin-client computing environment. The only processing done locally is KVM and sound. These are the modern-day 'intelligent' version of dumb terminals;
- tubby-clients - PCs that have an operating system, and possibly some applications, installed on them, but use a locally installed client to connect to the thin-client computing environment for all, or some, of their applications.

What is thin-client computing?

The fundamental approach behind thin-client computing is very simple. Instead of running applications locally on PCs with all of their associated challenges and costs, applications run centrally with only keyboard, video and mouse (KVM) updates transmitted across the network. Bandwidth usage is minimal compared to traditional PC/server environments, with wireless LAN being ideal for the clients. The server backbone linking the terminal servers, data servers, mail servers, and so on is the only LAN connection that needs high capacity.

In a traditional fat-client environment, applications are stored locally, and data is stored centrally. When a file is opened, the entire file is transferred to the local PC, with the results being saved back across the LAN/WAN to the central storage area. Server/client architecture (such as SQL and Oracle), handle this process slightly differently, but processing still takes place at the local PC. This requires high bandwidth to each PC. If you have downloaded an e-mail attachment over a dial-up connection, you'll be familiar with the issues!

Why thin-client computing?

Thin-client computing lowers costs and improves the service offering in several key areas.

Hardware

With more complex software being distributed, desktop hardware upgrades have been accepted as a necessary evil. The currently accepted useful life of a PC is 2 years, although often depreciated over 3 years. As a PC becomes under-powered for a particular need, it is reallocated, often requiring a different software build.

Thin-client computing turns this acceptance on its head. PCs can be used as tubby-clients until they die, although consider the costs associated with their power usage and maintenance of the base operating system before adopting this approach.

Centralised support

Significant benefits are obtained by centralising the support function, not only in savings, but also in the quality and consistency of the support function. The more diverse the geographical base, the more advantages can be gained.

Many software products provide various forms of remote take-over ability, or shadowing. This ability permits support staff to interact with the users' desktop as they're speaking to them. Performance of these products though is very slow when not on the same physical LAN. With thin-client computing there is no performance drop-off, since all users are running on the same LAN.

Windows 2000 terminal services do not have this capability built-in and add-on products specifically for terminal services will be required. This is expected to change with the .NET range of server products, with Windows XP already having support for remote help built in.

Centralised servers means not having server support staff responding to, or based at, remote sites for server support. Zona Research shows that support costs for 15 PCs in a Windows NT server environment were approximately 500 percent more than in a thin-client computing environment using thin-clients.

Bandwidth

A Microsoft study conducted by NEC and Groupe Bull shows that the highest bandwidth user is a structured task worker, typically performing the same tasks repetitively, eg claims processing, accounts, customer service. These workers would typically use 20Kbits of bandwidth, making LAN performance over dialup connections a reality.

One real-world installation of a tubby-client environment for a 130-user network showed the following:

	Average bandwidth utilisation	Peak bandwidth utilisation
Fat Client	40%	80%
Tubby Client	0.5%	4%

Significant bandwidth issues were being encountered during peak periods prior to the thin-client computing deployment. These figures are consistent with the Microsoft study.

The net effect is that far less bandwidth is required for remote and local sites, further reducing costs in multi-site installations. If your business is in the process of converting older 10Mbs LAN to 100Mbs or even 1Gbs, stop and reconsider.

Have you rejected wireless LAN as being too low bandwidth (currently 11Mbs) and hence too slow for your needs? Reconsider, since this is all you will need on user segments.

Do not upgrade your WAN bandwidth to accommodate increased file transfer needs, particularly e-mail. On the contrary, by switching to a thin-client computing model, you may well be able to downgrade your WAN size, further saving costs.

Power

Power consumption of a thin-client device is 14% of a PC. To place this in perspective, this is 5%, per year, of the thin-client device purchase price. Since a thin-client device will be expected to have a useful life beyond 5 years, the power savings alone will offset 25% of the cost of those devices.

Reduced cooling requirements also lower the costs. Steve Greenberg of thinclient.net, a leading design company for Fortune 500 companies in the US, shows in his power-for-the-people whitepaper, the following savings for a *thin-client* computing environment:

Annual £/kWh	Annual Savings (30 Users)	Annual Savings (100 Users)	Annual Savings (1000 Users)	Annual Savings (2500 Users)
£0.67	£457	£2,107	£21,071	£52,680

These savings are based on a real-world study and have been extrapolated for the UK with Greenberg's assistance. They include the cost of the extra servers required and cooling reductions.