Dublin Institute of Technology
ARROW@DIT

Conference papers | School of Computing

2009-01-01

Case Studies in Thin Client Acceptance

Paul Doyle
Dublin Institute of Technology, paul.doyle@dit.ie

Mark Deegan
Dublin Institute of Technology, mark.deegan@dit.ie

Follow this and additional works at: http://arrow.dit.ie/scschcomcon

Recommended Citation
ICIT Journal 2009 ICIT "Case Studies in Thin Client Acceptance"

This Conference Paper is brought to you for free and open access by the School of Computing at ARROW@DIT. It has been accepted for inclusion in Conference papers by an authorized administrator of ARROW@DIT. For more information, please contact yvonne.desmond@dit.ie, arrow.admin@dit.ie, brian.widdis@dit.ie.

This work is licensed under a Creative Commons Attribution-Noncommercial-Share Alike 3.0 License

CASE STUDIES IN THIN CLIENT ACCEPTANCE

Paul Doyle, Mark Deegan, David Markey, Rose Tinabo, Bossi Masamila, David Tracey
School of Computing, Dublin Institute of Technology, Ireland
WiSAR Lab, Letterkenny Institute of Technology
{paul.doyle, mark.deegan, david.markey}@dit.ie,{rose.tinabo, bossi.masamila}@student.dit.ie
david.tracey@lyit.ie

ABSTRACT

Thin Client technology boasts an impressive range of financial, technical and administrative benefits. Combined with virtualisation technology, higher bandwidth availability and cheaper high performance processors, many believe that Thin Clients have come of age. But despite a growing body of literature documenting successful Thin Client deployments there remains an undercurrent of concern regarding user acceptance of this technology and a belief that greater efforts are required to understand how to integrate Thin Clients into existing, predominantly PC-based, deployments. It would be more accurate to state that the challenge facing the acceptance of Thin Clients is a combination of architectural design and integration strategy rather than a purely technical issue. Careful selection of services to be offered over Thin Clients is essential to their acceptance. Through an evolution of three case studies the user acceptance issues were reviewed and resolved resulting in a 92% acceptance rate of the final Thin Client deployment. No significant bias was evident in our comparison of user attitudes towards desktop services delivered over PCs and Thin Clients.

Keywords: Thin Clients, Acceptance, Virtualisation, RDP, Terminal Services.

1 INTRODUCTION

It is generally accepted that in 1993 Tim Negris coined the phrase “Thin Client” in response to Larry Ellison’s request to differentiate the server centric model of Oracle from the desktop centric model prevalent at the time. Since then the technology has evolved from a concept to a reality with the introduction of a variety of hardware devices, network protocols and server centric virtualised environments. The Thin Client model offers users the ability to access centralised resources using full graphical desktops from remotely located, low cost, stateless devices. While there is sufficient literature in support of Thin Clients and their deployment, the strategies employed are not often well documented. To demonstrate the critical importance of how Thin Clients perform in relation to user acceptance we present a series of case studies highlighting key points to be addressed in order to ensure a successful deployment.

1.1 Research Aim

The aim of this research has been to identify a successful strategy for Thin Client acceptance within an educational institute. There is sufficient literature which discusses the benefits of Thin Client adoption, and while this was referenced it was not central to the aims of this research as the barrier to obtaining these benefits was seen to be acceptance of the technology. Over a four year period, three Thin Client case studies were run within the Dublin Institute of Technology with the explicit aim of determining the success factors in obtaining user satisfaction. The following data criteria were used to evaluate each case study in addition to referencing the Universal Theory of User Acceptance Testing (UTUAT) [1].

1) Login events on the Thin Clients.
2) Reservation of the Thin Client facility.
3) The cost of maintaining the service.

1.2 Paper Structure

In section 2 we review the historical background and trends of Thin Client technology to provide an understanding of what the technology entails. Section 3 discusses the case for Thin Clients within existing literature including a review of deployments within industry and other educational institutes. Section 4 provides details of the three case studies discussing their design, evaluating the results, and providing critical analysis. Section 5 takes a critical look at all of the data and sections 6 and 7 provide conclusions and identify future work. This paper is aimed at professionals within educational institutes seeking ways to realize the benefits of Thin Client computing while maintaining the support and acceptance of users. It provides a balance between the hype of Thin Clients and the reality of their deployment.

2 THIN CLIENT EVOLUTION

The history of Thin Clients is marked by a number of overly optimistic predictions that it was about to become the dominant model of desktop computing. In spite of this there have been a number of marked developments in this history along with those of desktop computing in general which are worth reviewing to set the context for examining the user acceptance of this technology. Thin Clients have established a role in desktop computing although not quite the dominant one initially predicted. These developments have usually been driven by increases in processing power (and reductions in the processor costs) in line with Moore's law, but the improvements in bandwidth and storage capacity are having an increasing effect on desktop computing and on Thin Client computing [2] driving the move towards more powerful lower cost desktops but also the possibilities of server virtualisation and Thin Client computing with the ability to run Thin Clients over WANs.

The first wave of computing was one where centralised mainframe computers provided the computing power as a shared resource which users accessed using dumb terminals which provided basic text based input and output and then limited graphics as they became graphics terminals. These mainframes were expensive to purchase and were administered by specialists in managed environments and mostly used for specific tasks such as performing scientific calculations and running highly specialised bespoke payroll systems.

The next wave was that of personal computing, whereby users administered their own systems which provided a platform for their personal applications, such as games, word-processor, mail and personal data. Since then the personal computer has undergone a number of significant changes, but the one of most interest was the nature of the interface provided to the user which has grown into a rich Graphical User Interface where the Personal Computer became a gateway to the Internet with the Web browser evolving into a platform for delivery of rich media content, such as audio and video.

This move from a mainframe centralised computing model to a PC distributed one resulted in a number of cost issues related to administration. This issue was of particular concern for corporate organizations, in relation to licensing, data security, maintenance and system upgrades. For these cost reasons and the potential for greater mobility for users, the use of Thin Clients is often put forward as a way to reduce costs using the centralised model of the Thin Client architecture. This also offers lower purchase costs and reduces the consumption of energy [3].

The challenge faced by Thin Client technology is to deliver on these lower costs and mobility, while continuing to provide a similarly rich GUI user experience to that provided by the desktop machine (a challenge helped by improved bandwidth, but latency is still often a limiting factor [4]) and the flexibility with regard to applications they have on their desktop. Typically, current Thin Client systems have an application on a server (generally Windows or Linux) which encodes the data to be rendered into a remote display protocol. This encoded data is sent over a network to a Thin Client application running on a PC or a dedicated Thin Client device to be decoded and displayed. The Thin Client will send user input such as keystrokes to the application on the server. The key point is that the Thin Client does not run the code for the user's application, but only the code required to support the remote display protocol.

While the term Thin Client was not used for dumb terminals attached to mainframes in the 1970's, the mainframe model shared many of the attributes of Thin Client computing. It was centralised, the mainframe ran the software application and held the data (or was attached to the data storage) and the terminal could be shared by users as it did not retain personal data or applications, but displayed content on the screen as sent to it by the mainframe. From a desktop point of view, the 1980's were dominated by the introduction and adoption of the Personal Computer.

Other users requiring higher performance and graphics used Unix Workstations from companies like Apollo and Sun Microsystems. The X Window System [5] was used on many Workstations and X terminals were developed as a display and input terminal and provided a lower cost alternative to a Unix Workstation, with the X terminal connecting to a central machine running an X display manager. As such, they shared some of the characteristics of a Thin Client system, although the X terminal ran an X Server making it more complicated than Thin Client devices.

The 1990's saw the introduction of several remote display protocols, such as Citrix's ICA [6] Microsoft's RDP [7] and AT&T's VNC [8] for Unix that took advantage of the increasing bandwidth available on a LAN to provide a remote desktop to users.

Terminal Services was introduced as part of Windows NT4.0 in 1996 and it offered support for the Remote Desktop Protocol (RDP) allowing access to Windows applications running on the Server, giving users access to a desktop on the Server using an RDP client on their PC. RDP is now offered on a range of Windows platforms [9]. Wyse and vendors such as Ncomputing launched terminals, which didn't run the Windows operating system, but accessed Windows applications on a Windows Server using RDP, which is probably still the dominant role of dedicated hardware Thin Clients. Similarly VNC is available on many Linux and Unix distributions and is commonly used to provide remote access to a user's desktop. These remote display protocols face increasing demands for more desktop functionality and richer media content, with ongoing work required in how, where and when display updates are encoded, compressed or cached [10]. Newer remote display protocols such as THINC have been designed with the aim of improving these capabilities [11].

In 1999, Sun Microsystems took the Thin Client model further with the SunRay, which was a simple network appliance, using its own remote display protocol called ALP. Unlike some of the other Thin Clients which ran their own operating system, SunRay emphasized its completely stateless nature [12]. This stateless nature meant that no session information or data was held or even cached (not even fonts) on the appliance itself and enabled its session mobility feature, whereby a smart card was used to identify a user with a session so that with the smartcard the user could login from any SunRay connected to the session's server and receive the desktop as it was previously.

Many of these existing players have since focused on improving their remote desktop protocols and support for multimedia or creating new hardware platforms. There have also been some newer arrivals like Pano Logic and Teradici who have developed specific client hardware to create "zero" clients, with supporting server virtualisation to render the remote display protocols. Also, there are a number of managed virtual desktops hosted in a data centre now being offered.

One of the drivers behind Thin Client Technology, particularly when combined with a dedicated hardware device, is to reduce the cost of the client by reducing the processing requirement to that of simply rendering content, but a second driver (and arguably more important one) is to gain a level of universality by simplifying the variations in the client side environment. This has been met in a number of new ways using Virtual Machine players and USB memory in Microsoft's research project "Desktop on a Keychain" (DOK) [13] and also the Moka5 product [14], allowing the mobility (and security) benefits attributed to Thin Clients. This can be enhanced with the use of network storage to cache session information [15].

It can be seen that Thin Clients have evolved along with other desktop computing approaches, often driven by the same factors of increasing processing power, storage capacity and bandwidth. However, newer trends that are emerging with regard to virtualisation, internet and browser technologies, together with local storage, present new challenges and opportunities for Thin Client technology to win user acceptance. As Weiser said in 1999 in this new era, "*hundreds or thousands of computers do our bidding. The relationship is the inverse of the mainframe era: the people get the air conditioning now, and the nice floors, and the computers live out in cyberspace and sit there waiting eagerly to do something for us*". [16]

3 THE CASE FOR THIN CLIENTS

There are many stated benefits for Thin Clients all of which are well documented [17][18]. While there is no single definitive list, potential system designers may have different aims when considering Thin Clients, these benefits should be clearly understood prior to embarking on any deployment and are discussed below.

3.1 Reduced cost of software maintenance

The administrative cost benefit of the Thin Client model, according to Jern [19] is based on the simple observation that there are fewer desktop images to manage. With the combination of virtualisation environments and Windows Terminal Service (WTS) systems it would not be uncommon for twenty five or more desktop environments to be supported from a single installation and configuration. This reduces the number of upgrades and customizations required for desktop images in computer laboratories where the aim is to provide a consistent service from all systems. Kissler and Hoyt [20] remind us that the "*creative use of Thin Client technology can decrease both management complexity and IT staff time*." In particular they chose Thin Client technology to reduce the complexity of managing a large number of kiosks and quick-access stations in their new thirty three million dollar library. They have also deployed Thin Client devices in a range of other roles throughout Valparaiso University in Indiana. Golick [21] on the other hand suggests that the potential benefits of a Thin Client approach include the lower mean time to repair (MTTR) and lower distribution costs. It is interesting to note that he does suggest that the potential cost savings for hardware are a myth, but that administration savings still make a compelling case for using Thin Client technology.

3.2 Enhanced Security

Speer and Angelucci [22] suggest that security concerns should be a major factor in the decision to adopt Thin Client systems and this becomes more apparent when referencing the Gartner Thin Client classification model. The Thin Client approach ensures that data is stored and controlled at the data-centre hosting the Thin Client devices. It is easy to argue that the user can retain the mobility of laptops but with enhanced security and the data is not mobile, just the access point. The argument is even easier to make when we consider recent high-profile cases of the theft of unencrypted laptops containing sensitive medical or financial records. The freedom conferred on users of corporate desktop and laptop PCs undermines the corporation's obligations in relation to data privacy and security. Steps taken to protect sensitive data on user devices are often too little and too late. Strassmann [23] states that the most frequent use of a personal computer is for accessing web applications and states that the Thin Client model demonstrates significantly lower security risks for the corporation. Five security justifications for adopting the Thin Client model were proposed.

1) Zombie Prevention
2) Theft Dodging
3) File Management
4) Software Control
5) Personal Use Limitations

Strassmann concedes that Thin Clients are not necessarily best for every enterprise and every class of user, but for enterprises with a large number of stationary "non-power" users, "Thin Clients may present the best option in terms of security, cost effectiveness and ease of management."

3.3 User Mobility

User mobility can refer to the ability of a user to use any device, typically within the corporation's intranet, as a desktop where the user will see a consistent view of the system, for example, SunRay hot-desking. While user profiles in Microsoft Windows support this, it is often only partially implemented. Session mobility can be viewed as the facility for users to temporarily suspend or disconnect their desktop session and to have it re-appear, at their request, on a different device at a later time. This facility removes the need for users to log-out or to boot-up a desktop system each time they wish to log-in. Both of these potential features of Thin Client technologies help to break the sense of personal ownership that users often feel for their desktop or laptop computers. It is this sense of personal ownership which makes the maintenance and replacement of corporate PCs a difficult task, and this feeling of ownership and control is often a reason why users resist the adoption of a centrally controlled Thin Client to replace their desktop, whereas this is exactly why IT management may want to adopt it.

3.4 Environmental Costs

In the article "*An Inefficient Truth*" Plan [24] reveals a series of "truths" supported by a number of case studies directed at the growing costs of Information and Communication Technologies. One such case study is of *Reed Managed Services* where 4,500 PCs were replaced with Thin Clients, and a centralised blade server providing server based virtualised desktops. Savings are reported as follows:

1) 5.4 million kWh reduction,
2) 2,800 tonnes of CO_2 saved annually
3) Servers reduced by a factor of 20
4) IT budget cut by a fifth

Indeed there are many deployments focused on obtaining energy savings through the use of Thin Clients. In a case study where SunRay systems were introduced into Sparkasse a public German Bank, Bruno-Britz [25] reports that the savings in electricity costs alone were enormous. The University of Oxford has deployed SunRay Thin Client devices in their libraries citing the cooler and quieter operation as factors in their decision. These devices, having no local hard disk and no fan operate at a lower temperature and more quietly than traditional PCs. This characteristic has environmental implications from noise, cooling and power consumption perspectives.

3.5 Summary of Benefits

In summary, we can extract the benefits observed within literature and case studies as follows:

1) Increased security as data maintained centrally
2) Reduced cost of hardware deployment and management and faster MTTR
3) Reduced administration support costs
4) Environmental costs savings
5) Reduced cost of software maintenance
6) Reduced cost of software distribution
7) Zero cost of local software support
8) The ability to leverage existing desktop hardware and software
9) Interface portability and session mobility
10) Enhanced Capacity planning
11) Centralised Usage Tracking and Capacity Planning

3.6 Thin Clients vs. Fat Clients

Thin Client technology has evolved in sophistication and capability since the middle of the 1990s, however the "thickness" (the amount of software and administration required on the access device) of the client is a source of distinction for many vendors [26][27]. Regardless of "thickness", Thin Clients require less configuration and support when compared to Fat Clients (your typical PC). In the early 1990s Gartner provided a client-server reference design shown in Figure 1. This design provides clarity for the terms "thin" and "fat" clients by viewing applications in terms of the degree of data access, application and presentation logic present on the server and client sides of the network.

The demand for network based services such as email, social networking and the World Wide Web has driven bandwidth and connectivity requirements to higher and higher levels of reliability and performance [28]. As we progress to an "*always on*" network infrastructure the arguments focused against Thin Clients based on requiring an offline mode of usage are less relevant. The move from Fat Client to Thin Client is however often resisted as individuals find themselves uncomfortable with the lack of choice provided when the transition is made, as observed by Wong et al.[29].

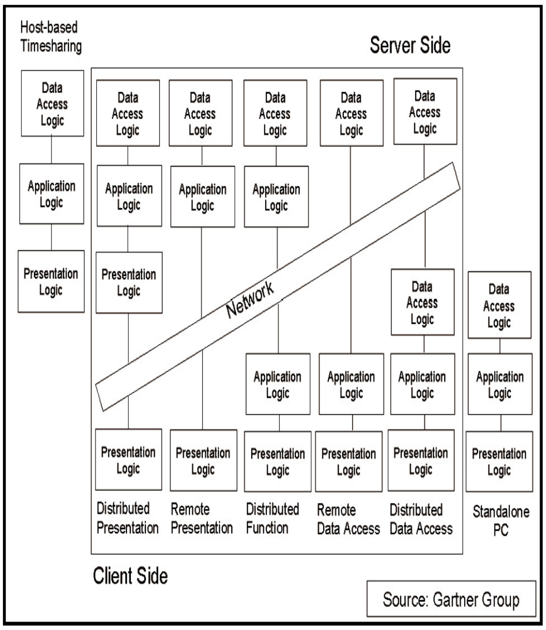

Figure 1: Gartner Group Client/Server Reference Design

4 CASE STUDIES

No matter how well documented the benefits of Thin Clients may be, there is still an issue of acceptance to be addressed. While it may be tempting to assume that the implementation of technology is a technical issue and that simply by building solutions a problem is effectively solved, evidence would point to the contrary. As there can often be a disparity between what is built and what is required or needed. Too often requirements gathering, specification definition and user consultation are forgotten in the rush to provide new services which are believed to be essential. In essence the notion of "***if we build it they will come***" is adopted, inevitably causing confusion and frustration for both service provider and the user. For example, during Sun Microsystems' internal deployment of its own SunRay Thin Client solution many groups and functions sought exemptions from the deployment as they believed that their requirements were sufficiently different to the "generic user" to warrant exclusion from the project. The same arguments still exist today and it is often those with a more technical understanding of the technology who are the agents of that technology's demise. By providing interesting and often creative edge cases which identify the limitations of a technology, they can, by implication, tarnish it as an incomplete and flawed technology. In the case of Thin Clients, it should be accepted that there are tradeoffs to be made. One of the appealing aspects of the Fat client is its ability to be highly flexible which facilitates extensive customization. However not every user will require that flexibility and customization. Thin Clients are not going to be a silver bullet addressing all users needs all of the time.

All three case studies were evaluated under the following headings in order to allow a direct comparison between each. These criteria were selected to ensure that there was a balance between the user acceptance of the technology and the technical success of each deployment.

1) Login events on the Thin Clients
2) Reservation of the Thin Client facility
3) The cost of maintaining the service

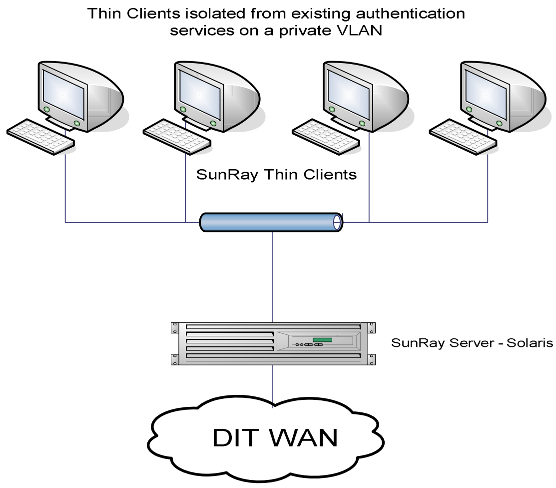

Figure 2: Case Study 1

4.1 DIT Case Study 1

In 2005 the DIT introduced the SunRay Thin Client technology into the School of Computing. In a similar approach to many other technology deployments the strengths of the technology were reviewed and seen as the major selling points of the deployment. In the case of SunRay there was a cheap appliance available which would provide the service of graphical based Unix desktops. Centralised administration ensured that the support costs would be low and the replacement requirements for systems for the next five years would be negligible. In essence the technological and administrative advantages were the focus of this deployment. Few of the services offered within the existing PC infrastructure were included in the deployment. This deployment sought to offer new services to students and introduced Thin Clients for the first time to both students and staff.

4.1.1 Design

A single laboratory was identified for deploying the SunRay systems and all PC in that lab were replaced with SunRay 150 devices. A private network interconnect was built which ensured that all data sent from the clients traversed a private network to the SunRay server. The initial design of this case study is shown in Figure 2 and it allowed students within this new Thin Client lab access to the latest version of Solaris using a full screen graphical environment as opposed to an SSH command-line Unix shell which was the traditional method still used from existing computing laboratories. A new authentication system was introduced based on LDAP which required students to have a new username and password combination which was different to the credentials already in use within the Active Directory domain used for the existing PC network. The reason for this alternative authentication process was due to the difficulty of authenticating on a Unix system using Active Directory. Once the server was running, the Thin Client laboratory was ready to provide graphical based Unix login sessions at a considerable reduced price when compared to an investment of Unix workstations for each desk. In total 25 Thin Client devices were installed which were all connected to a single Solaris server. In summary the key components within the design were as follows:

1) The service was on a private network
2) New authentication process was introduced
3) New storage mechanism was introduced
4) Devices were all in the same location
5) Service provided was a CDE desktop on Solaris
6) Graphical desktops running on Linux servers also accessible

4.1.2 Results

The login events are a measure of the general activity of the devices themselves and were considered to be a reasonable benchmark for comparison with existing laboratories within the institute. One interesting point is that the comparison of facilities is not necessarily relevant when the facilities provide different services. Due to the fact that Unix instead of Windows was provided meant that, with the exception of those taking courses involving Unix, the majority of students were unfamiliar with the technology and did not seek to use the systems.

Login events on the Thin Clients:

The login events were extracted from the Solaris server by parsing the output of the ***last*** command which displays the login and logout information for users which it extracts from the /var/adm/wtrmpx file. The number of login events per day was calculated and plotted in the graph shown in Fig. 3. Immediately obvious was the low use of the system. Given that the nature of the service did not significantly change over the course of the three years that the system was in place with the exception of semester activity in line with student presence in the institute, it is clear that there was low utilization of the service. The graph shows raw data plotted, where login events were less than 10 per day.

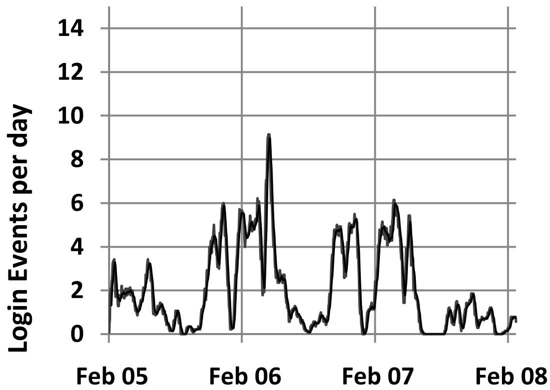

Figure 3: User Login Events

Reservation of the Thin Client Facility:

Each laboratory may be reserved by staff for the delivery of tutorial sessions and exercises. The hourly reservations for this laboratory were reduced as a result of the introduction of Thin Clients with only 1 to 2 hours being reserved per day. One of the primary reasons for the reduction in the use of this facility was the fact that it had now become special purpose and the bookings for the room were limited to the courses which could be taught within it.

The Cost of Maintaining the Service:

A detailed analysis of cost savings associated with the introduction of Thin Clients within our institute and specifically the costs associated with this case study was performed by Reynolds and Gleeson, [30]. In their study they presented evidence of savings in relation to the cost of support, the cost of deployment and a basic analysis of the power consumption costs. They review both the system and the software distribution steps associated with Thin Clients and PC systems and present a point of quantifiable comparison between the two. Key findings of this analysis were as follows:

1) Time spent performing system upgrades and hardware maintenance was reduced to virtually zero as no hardware or software upgrades were required.
2) A single software image was maintained at the central server location and changes were made available instantly to all users.
3) No upgrade costs were incurred on the Thin Clients or server hardware. All systems have

remained in place throughout both case studies. The devices in this lab are now 8 years old and are fulfilling the same role today as they did when first installed.
4) The Thin Client lab is a low power consumption environment due to the inherent energy efficiency of the Thin Client hardware over existing PCs. This can provide up to 95% energy savings when compared to traditional PCs [24].

4.1.3 Analysis

There has been extensive research in the area of user acceptance of technology, but perhaps the most relevant work in this area is the Unified Theory of Acceptance and Use of Technology (UTAUT) [1] which identifies four primary constructs or factors;

a) Performance Expectancy
b) Effort Expectancy
c) Social Influence
d) Facilitating Conditions

While there are additional factors such as *Gender*, *Age* and *Experience*, within the student populations these are for the most part reasonably consistent and will be ignored. It should be stressed that although the UTAUT was developed for an industry based environment it is easily adapted for our purposes. It was felt that this model serves as a relevant reference point when discussing the performance of the case studies.

Clearly Case Study 1 failed to gain acceptance despite belief that it would in fact be highly successful at its inception. We review the case study under the four UTAUT headings to identify the source of the user rejection of the Thin Clients.

a) Performance Expectancy
This factor is concerned with the degree to which the technology will assist in enhancing a users own performance. Clearly however the services provided an advantage to those students who wished to use Unix systems. Since the majority of courses are based on the Windows operating system it would be reasonable to assume that there was no perceived advantage in using a system which was not 100% compatible with the productivity applications used as part of the majority of courses.

b) Effort Expectancy
This factor is concerned with the degree of ease associated with the use of the system. One of the clear outcomes of Case Study 1 was that students rejected the Unix systems as it was seen to be a highly complex system, requiring additional authentication beyond what was currently used in traditional laboratories.

c) Social Influence
This is defined as the degree to which there is a perception of how others will view or judge them based on their use of the system. Clearly by isolating the devices and having it associated with specialized courses, there was no social imperative to use the labs. Unix as a desktop was relatively uncommon in the School at the time of the case study and there would have been a moderate to strong elitist view of those who were technical enough to use the systems.

d) Facilitating Conditions
This is defined as the degree to which an individual believes in the support for a system. At first glance this does not appear to be a significant factor considering that the services were created by the support team and there was considerable vested interest in seeing it succeed. However additional questions asked by the UTAUT include the issue of compatibility with systems primarily used by the individual.

Each of the UTAUT factors can be considered significant for Case Study 1. Many of the issues raised hang on the fundamental issue that the new services offered on the Thin Client were different to existing services and for all practical purposes seen as incompatible with the majority of systems available to students elsewhere. The fact that the technology itself may have worked flawlessly, and may have delivered reduced costs was irrelevant as the service remained under utilized. Given that the reason for this lack of acceptance was potentially inherent in the implementation of services and not due to failings in the technology itself it was clear that a second case study was required which would address the issue of service.

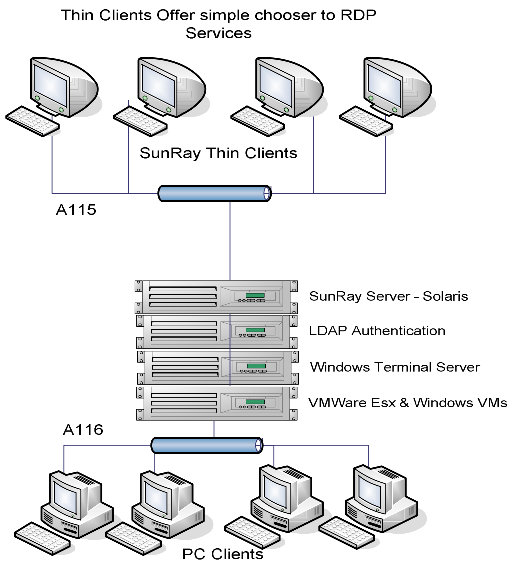

Figure 4: Case Study 2

4.2 Case Study 2

The second case study is a modification of the basic implementation of the first case study with changes focused on increasing student acceptance of the Thin Client facility. Removing the Unix centric nature of the existing service was central to the system redesign. It was decided that additional services could be easily and cheaply offered to the Thin Client environment providing users with the ability to access more compatible services from within the Thin Client environment. Figure 4 identifies the key components within the design.

4.2.1 Design

The most important addition to the second case study was the provision of additional services which were similar to those available in PC labs. This was to ensure that students could use this facility and have an experience on a par with the PC labs. A new domain was created where Unix and Windows shared a common authentication process. Due to difficulties integrating Unix and the existing Windows authentication process, the new Domain was built on the LDAP system with SAMBA providing the link between the new Windows Terminal Servers and the LDAP system. While students could now use the same username and password combination for Windows and Unix systems this was not integrated into the existing Windows authentication process. Students were still required to have two sets of credentials, the first for the existing PC labs, and the second for access to a new domain containing a number of Windows Terminal Servers and the original graphical Unix desktop. While the Thin Clients now provided Windows and Unix graphical desktops, the new Windows Domain was also accessible from existing PC labs via RDP connections to the Terminal Servers. This allowed classes to be scheduled either inside or outside of the Thin Client laboratory. In addition to providing Windows Terminal Services (WTS), student owned virtual machines were now also available. Due to the fact that most services were now available from all locations, the ease of access to the services from within the Thin Client lab was improved by providing users with a menu of destinations upon login. This new login script effectively provided a configurable redirection service to the WTS and Virtualisation destinations using the ***rdesktop*** utility [31] which performed a full screen RDP connection to specified destinations. An interesting outcome of this destination chooser was that any RDP based destination could be included regardless of the authentication process used. This would however require a second authentication process with the connecting service. The new services provided were as follows:

a) A general purpose Windows Terminal Server with mounted storage for all students and staff.
b) Course specific Windows Terminal Servers for courses where there were specific software requirements not common to all students.
c) Individual Virtualised desktops for students in specific modules where administration rights were required.
d) All services were made available from both the Thin Client and PC labs as they were available over the Remote Desktop Protocol RDP.
e) Provisioning of an easy access point to all services from within the Thin Client environment which was not available from PC systems.

4.2.2 Results

The data gathered for Case Study 2 was evaluated under same three headings as per case study 1.

1) Login events on the Thin Clients
2) Reservation of the Thin Client facility.
3) The cost of maintaining the service.

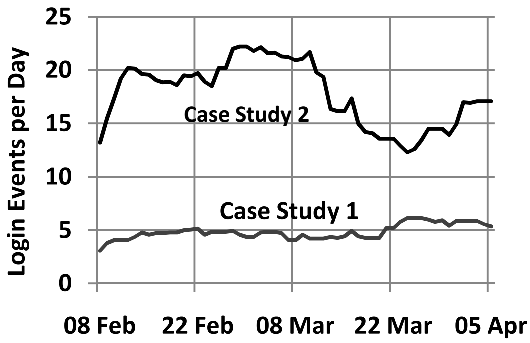

Figure 5: User Login Event Comparison

Login events on the Thin Clients:

Figure 5 shows a comparison of activity during the same time period for the two case studies. To identify trends in the data a displacement forward moving average was performed on the data as shown in Eq. (1).

$$SMA = \frac{L_t + \ \dots \ L_{t+13}}{14} \qquad (1)$$

It is clear that for the same time period there was a significant increase in the use of the system as the number of login events increased by a factor of 4. Once again the login events were extracted from the Solaris server by parsing the output of the ***last*** command.

Reservation of the Thin Client Facility:

The changes to the Thin Client facility were announced at the start of the second academic semester as a *PC upgrade* and the number of room bookings increased as shown in Figure 6 from 6 hours a week to 20 hours a week. This was due to the use of the room as a Windows based laboratory using the new WTS and virtualisation services.

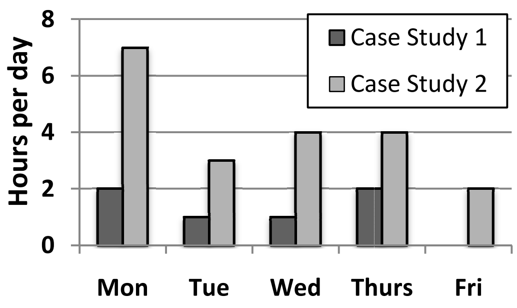

Figure 6: Thin Client Room Reservations

The Cost of Maintaining the Service:

All of the benefits observed from the first case study were retained within this case study. The addition of terminal services reduced the reliance of students on Fat Client installations. Students are now using virtual machines and terminal servers on a regular basis from all labs.

4.2.3 Analysis

This second case study certainly saw an improvement over its earlier counterpart and students and staff could now access more familiar services from the Thin Client lab. Given the dramatic increase relative to the earlier results it could be stated that the introduction of the more familiar services increased the acceptance of the facility. Both case studies demonstrated equally well that it is possible to obtain the total cost of ownership benefits using a Thin Client model, but the services offered has a dramatic affect on user acceptance. It is useful to review the outcome in relation to the UTUAT.

a) *Performance Expectancy*
Given that new services such as personalised virtual machines were now available, staff and students could identify a clear advantage to the system where administration rights could be provided in a safe manner, allowing more complex and previously unsupported activities to take place. For example, the Advanced Internet module for the MSc. students could now build and administer full web servers which could remain private to the student ensuring that no other student could access or modify a project which was a work in progress.
b) *Effort Expectancy*
Considerable improvements were made in this case study to allow users to access well known environments from both the Thin Clients and PC systems. Students who were taught modules using the new WTS or virtual environments were trained on how to access the systems, and once they used them they continued to do so throughout the year. Those who did not have modules being taught using these new services were still required go through a new login/access process which was not well documented. For example within the Thin Client labs the new username/password combination was required to access the choice of destinations from the devices. This acted as a barrier to use even though emails were sent to students and information on how to access these accounts were posted in the labs. Usernames were based on existing student ID numbers.
c) *Social Influence*
Little changed in this case study for those who did not have a teaching requirement based on the new services.
d) *Facilitating Conditions*
With the provision of WTS services and virtual machines which provided Windows environments the issue of compatibility was reduced. However two issues remained which were not addressed. Firstly while users could now share a common data store between systems on this new domain there was no pre-packaged access to the data store on the existing PC domain. While it was technically possible to combine both under a single view, this required user intervention and additional training which was not provided. Secondly the sequence of steps required to access choices from the Thin Clients was a non-standard login process which now required a second login, the first of which was at a Unix graphical login screen. For many this initial login step remained as a barrier to using the system.

The most striking result from this case study is that while the second case study demonstrated significant increase in acceptance and use, the PC environments remained the system of choice for students, as shown in Figure 7. In this graph we show the typical use PC laboratory within the same faculty. Thin Client use remained less than one third of the use of the busiest computer laboratory. Thin Clients are shown to be capable of providing services equally well to both Windows and Unix users by introducing the ability of students to access their own private desktop from many locations, however this feature alone was not enough to entice users from the existing PC infrastructure. Clearly the introduction of virtualisation to the infrastructure allowed new services to be developed and used from Thin and Fat clients which could be seen as a potential for migrating users to a Thin Client/Virtualisation model, which indeed is a future planned initiative. The results show a definite increase in the use of the Thin Client facilities with data being gathered from the same period over both case studies to eliminate any bias which might occur due to module schedule differences at different time periods during the year.

The timing and method used to announce the changes was critical to the increase in acceptance. The announcement of the systems as a PC upgrade removed some of the barriers which existed for users who did not feel comfortable with a Unix environment but failed to attract a majority of the students.

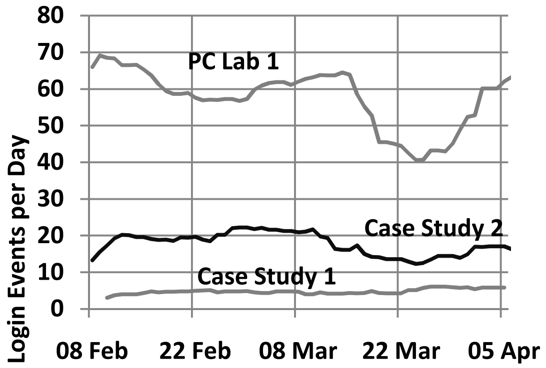

Figure 7: Comparison with PC Computer Labs

4.3 Case Study 3

The third case study was designed using the experiences of the first two case studies and was extended beyond the School of Computing. It was aimed at demonstrating the capability of the Thin Client technology in two different demographic environments, the first was one of the Institute Libraries where PCs were used by students from many different faculties and the second was within the Business faculty where computer system use was provided in support of modules taught within that faculty. This case study expressed the following aims at the outset

1) To demonstrate the use of Thin Client technology within the student population and determine the level of student acceptance of that technology.
2) To implement a number of alternative technologies in order to provide a point of comparison with respect to their overall performance and acceptance.
3) To determine the capability of the existing network infrastructure to support Thin Clients.

4.3.1 Design

Unlike the previous case studies the aim was to insert Thin Clients into the existing environment as invisibly as possible. This meant that existing authentication processes were to be maintained. There were two different authentication processes in place which needed to be support, Novell Client for the Business faculty and Active Directory for the Library. In both cases a WTS system was built which joined to the respective domains. Applications were installed on the Thin Client in order to mirror those that were present on existing PCs in the chosen locations. It was essential that the Thin Clients were not to be identifiable by students if at all possible, and to co-locate them with existing PC systems. To ensure that all devices behaved in a consistent manner to PCs they must boot and present the same login screen as would be expected on a PC in the same location. To achieve this all Thin Client devices with the exception of the SunRay systems used a Preboot Execution Environment (PXE) [32] boot process to connect to a Linux Terminal Server Project server (LTSP). The server redirected the user session to the correct WTS using **rdesktop** where the user was presented with a Windows login screen identical to those on adjacent PC systems.

The SunRay systems were run in Kiosk mode which allowed the boot sequence to redirect the session to a WTS also via the **rdesktop** utility. The WTS were installed on a VMWare ESX Server to allow rollback and recovery of the servers. This however was not central to the design of the case study and only served as a convenience in sharing hardware resources between multiple servers. The only concern was the potential performance of the WTS under a virtualised model. Given that the applications were primarily productivity applications such as word processing and browsing, and that the maximum number of users allowable on any WTS was 25 (based on the number of devices which were directly connected to the WTS) this was considered to be within the acceptable performance range of the architecture. This assumption was tested prior to the case study being made accessible to students with no specific issues raised as to warrant further restructuring of the architecture

Seventy five Thin Clients were deployed in six locations. The following Thin Client devices were used as shown in Figure 8 and Table 1.

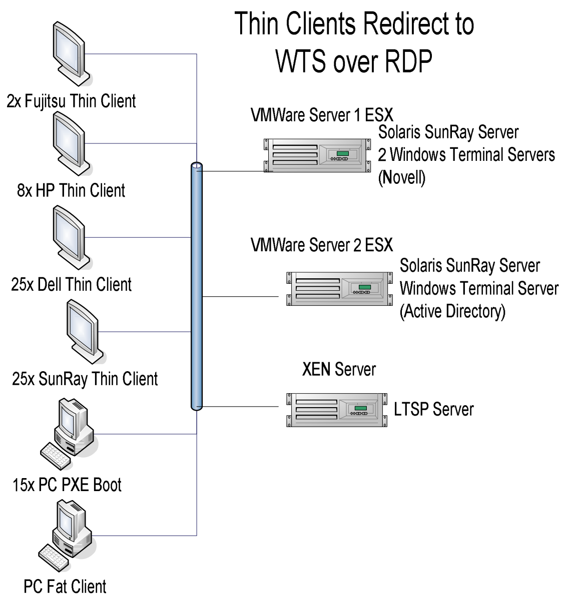

Figure 8: Case Study 3

Table 1: Thin Clients deployed

Device	Boot Mode	Quantity
Dell GX260	PXE Boot PC	15
Dell FX 160	PXE Boot TC	25
HP T5730	PXE Boot TC	8
Fujitsu FUTRO S	PXE Boot TC	2
SunRay 270	SunRay	25

4.3.2 Linux Terminal Server Project

LTSP works by configuring PCs or suitable Thin Clients to use PXE-Boot to obtain the necessary kernel and RDP client used as part of this project. These are obtained from a TFTP server whose IP address is provided as a DHCP parameter when the client PXE-Boots. As part of the DHCP dialogue, devices configured to PXE-Boot are given settings by the DHCP server. These include; *TFTP Boot Server Host Name* and *Bootfile Name*.

The necessary settings were configured on each of the DHCP servers serving the relevant locations within the DIT so as to point any PXE-Boot devices to the relevant LTSP boot server and to specify the kernel to be loaded by the PXE-Boot client. Using these settings the PXE-Boot clients load a Linux kernel and then an RDP client which connects to one of the three WTS used as part of this case study.

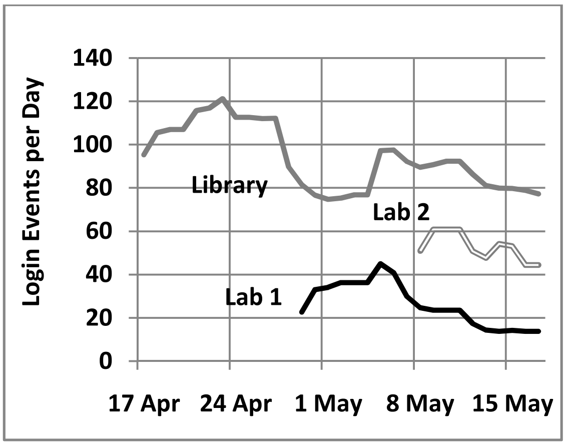

Figure 9: User Login Event Comparison

4.3.3 Results

Use of the Thin Clients was recorded using login scripts on the Windows Terminal Servers which recorded login and logout events. As expected the use of the Library systems exceed the use of the laboratories but both were in line with typical use patterns expected for each location. What was immediately obvious was that each location had a higher utilization than the previous two case studies but comparable with the PC labs shown in Figure 9. One of the difficulties with the comparison however is that the final case study was performed at a different point in the teaching semester and use of the systems declined as students prepared for examinations. Lab 1 was a "quiet lab" located remotely from the primary labs within the Business faculty and traditionally did not have high use. Lab 2 was a more central location and again as expected this exhibited greater user activity. The systems remained in operation continually for the period of the case study which was over one month during which data was collected from the three WTS systems.

4.3.4 User Survey

Once the case study was running a desktop satisfaction survey which employed the Likert scale [33] was conducted to obtain feedback from students using the Thin Client systems. The design of the questionnaire was such that students were asked to identify their desktop using a colour coded system which was known only to the authors. Each of the Thin Clients and a selection of PC systems (which were not PXE booted) where targeted for the survey to allow a comparative analysis between all Thin Clients and existing PC systems to be performed. The survey did not reference Thin Clients in any of the questions but rather sought feedback on application use and overall satisfaction with the performance of the system through a series of questions. There were 234 responses recorded for the survey. The key questions in the survey were as follows.

1) *Please rate the overall performance of the machine you are currently using*
2) *Please identify the primary reason you used this computer*
3) *How would you rate your overall satisfaction with this desktop?*
4) *Would you use this desktop computer again?*

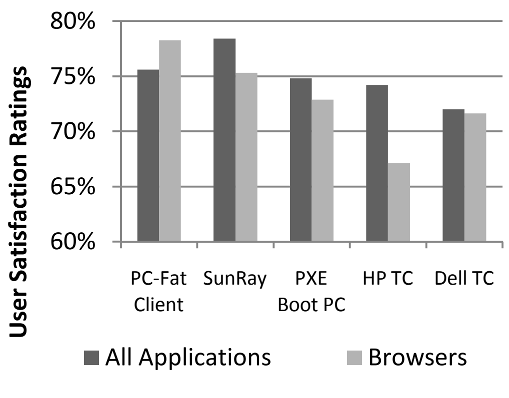

Figure 10: User satisfaction rating of desktop performance

The issue of overall performance was broken down by the device used which was identified using the colour coded scheme described earlier. Figure 10 below represents the average rating of satisfaction reported by users broken down by device and primary application in use. Since over 50% of responses identified "*Browsing*" as the primary reason for using the machine there are two satisfaction ratings provided as a point of comparison. Figure 11 shows the combined rating of users responses to overall satisfaction with desktop, desktop performance and application performance.

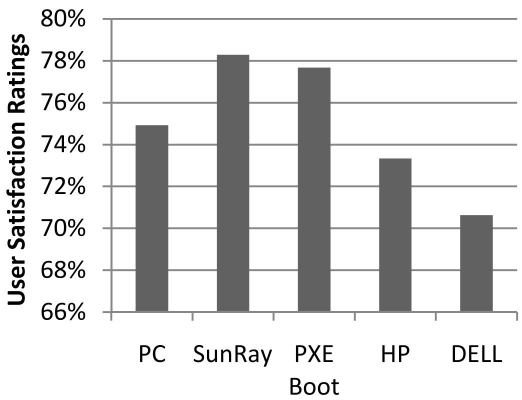

Figure 11: Combined rating of desktop performance

4.3.5 Analysis

This final case study while shorter in length than the other case studies demonstrated significant progress in user acceptance. As part of the survey users were asked if they would consider reusing the system and as can be seen in Figure 12 there was significant support for the systems.

The small number of responses representing those who did not wish to reuse the system cited USB performance as the primary cause of their dissatisfaction. This was identified early in the testing of the Thin Clients that all systems performed noticeably slower than the PC systems in this respect. Questions regarding the primary storage method used by students were added to the survey as was a satisfaction rating. From the results in Figure 13 it is clear that while the PC systems did perform better when users primarily used USB storage, the satisfaction in storage performance for all other options were comparable. The HP satisfaction rate had a low survey response rate and hence was not considered significant in our analysis given the small number of data points.

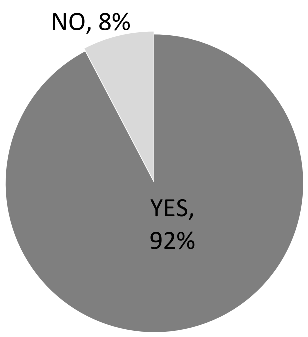

Figure 12: User Response "Would you use this system again"

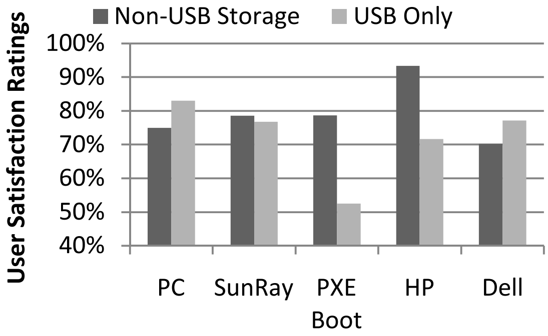

Figure 13: Storage Satisfaction Rating

By making the Thin Clients as invisible as possible and comparing satisfaction and user access to the existing PC systems it was clear that for the majority of users there was no apparent change to the services provided. Integrating into the existing authentication process was an essential feature of this case study as was the presenting of a single authentication process at the WTS login screen. Efforts were also made to ensure that the applications installed on the WTS were configured to look and feel the same as those on the standard PC. As with the previous case studies it is useful to review the case study in relation to the UTUAT.

a) *Performance Expectancy*
   With the exception of increasing the number of desktops in the Library, the primary deployment mainly replaced existing systems, so users were not provided with any reminders that they were using a different system. In effect there was no new decision or evaluation by the user to address the questions which were relevant in the previous case studies.
b) *Effort Expectancy*
   The reuse of the existing login/access procedure which was well known and part of the normal process for students using existing PC systems again allowed for this factor to become mainly irrelevant. Usernames, passwords, applications and system behaviour were identical to those on the PCs.
c) *Social Influence*
   Without perceiving a difference in service, social influence as a factor was also eliminated. Only the SunRay systems had different keyboards and screens, and as these screens were of higher resolution than existing PCs they were if anything seen as a more popular system.
d) *Facilitating Conditions*
   Unlike the previous case studies support for the facility was more complex. Different levels of expertise and engagement were required. Thin

Clients were now part of a larger support structure where many individuals were not core members of the technical team who built the systems. However given that only three support calls were raised during the case study there was little pressure on this factor either. The calls raised were not in fact directly related to the Thin Client devices, but rather the network and the virtual environments used to host the centralised servers.

5 CRITICAL ANALYSIS

The UTUAT provides a useful reference point in understanding some of the factors affecting acceptance of the Thin Clients. In the first case study the primary barrier to acceptance was the incompatibility of the new system with the existing system. Students were not motivated to use the new system as there were few advantages to doing so and considerable effort in learning how to use it. The second case study while more successful still failed to gain acceptance despite the expansion of services offered being comparable with existing Windows services. The session mobility and access from anywhere feature, while useful did not overcome the resistance of users to migrate to the Thin Clients. Thin Clients still required separate credentials and the login process was still different to the PC systems. The third and final case study was designed to provide the same existing services as the PC only using a centralised server and Thin Client model. No new services for the user were provided. The primary aim was to have the systems indistinguishable from the existing installation of PCs, effectively running a blind test for user acceptance. Once the users accepted the new systems, further machines could be deployed quickly and cheaply. The total cost of ownership and centralised support savings demonstrated in the first two case studies were just as relevant in the third case study.

6 CONCLUSION

There is considerable literature in support of Thin Client technology, and while there may be debate regarding the finer points of its advantages the issue has been and continues to be one of acceptance. Acceptance for Thin Clients as a technology is often confused with the non technical issues arising from the deployment. The UTUAT helps distinguish between technical and non-technical issues and as shown within our case studies, the way in which the technology was presented to the user had a higher impact on acceptance than had the technology itself. This point is highlighted by the fact that the Thin Client devices which were not widely used in first case study were integrated seamlessly into the third case study. These three case studies provide data centric analysis of user acceptance and identify the evolving designs of our deployments. To gain acceptance of Thin Clients within an educational institute our case studies identified these key factors.

1) Locate the Thin Clients among the existing PC systems, do not separate them or isolate them.
2) Ensure that the login process and credentials users are identical to the existing PC systems.
3) Ensure that the storage options are identical to the existing PC systems
4) Focus on providing exactly the same services that already exist as opposed to focusing on out new services.

By ensuring we ran a blind test on the user population where Thin Clients co-existed with PC systems, and where the services offered were indistinguishable by the user, we were able to show a user satisfaction rating of 92%. No significant bias was evident in our comparison of user attitudes of desktop services delivered over PCs and Thin Clients.

7 FUTURE WORK

Additional case studies are planned which will focus on acceptance of Thin Clients within the academic staff population and will evaluate the relevance of some of the proposed core technological advantages within that environment such as session mobility, Desktop as a Service, and dynamic lab reconfiguration and remote access using WAN and not just LAN environments.

8 REFERENCES

[1] V. Venkatesh, M.G. Morris, G.B. Davis, and F.D. Davis, “User acceptance of information technology: Toward a unified view,” *Mis Quarterly*, 2003, pp. 425-478.

[2] J.D. Northcutt, “CYB Newslog - Toward Virtual Computing Environments.”

[3] D. Tynan, “Think thin,” *InfoWorld*, Jul. 2005.

[4] S.J. Yang, J. Nieh, M. Selsky, and N. Tiwari, “The Performance of Remote Display Mechanisms for Thin-Client Computing,” *IN PROCEEDINGS OF THE 2002 USENIX ANNUAL TECHNICAL CONFERENCE*, 2002.

[5] T. Richardson, F. Bennett, G. Mapp, and A. Hopper, “Teleporting in an X window system environment,” *IEEE Personal Communications Magazine*, vol. 1, 1994, pp. 6-13.

[6] Citrix Systems, “Citrix MetaFrame 1.8 Backgrounder,” Jun. 1998.

[7] Microsoft Corporation, "Remote Desktop Protocol: Basic Connectivity and Graphics Remoting Specification," *Technical White Paper, Redmond*, 2000.
[8] T. Richardson, Q. Stafford-Fraser, K. Wood, and A. Hopper, "Virtual network computing," *Internet Computing, IEEE*, vol. 2, 1998, pp. 33-38.
[9] Microsoft Corporation, "Microsoft Windows NT Server 4.0, Terminal Server Edition: An Architectural Overview," Jun. 1998.
[10] J. Nieh, S.J. Yang, and N. Novik, "A comparison of thin-client computing architectures," *Network Computing Laboratory, Columbia University, Technical Report CUCS-022-00*, 2000.
[11] R.A. Baratto, L.N. Kim, and J. Nieh, "Thinc: A virtual display architecture for thin-client computing," *Proceedings of the twentieth ACM symposium on Operating systems principles*, ACM New York, NY, USA, 2005, pp. 277-290.
[12] B.K. Schmidt, M.S. Lam, and J.D. Northcutt, "The interactive performance of SLIM: a stateless, thin-client architecture," *Proceedings of the seventeenth ACM symposium on Operating systems principles*, Charleston, South Carolina, United States: ACM, 1999, pp. 32-47.
[13] M. Annamalai, A. Birrell, D. Fetterly, and T. Wobber, *Implementing Portable Desktops: A New Option and Comparisons*, Microsoft Corporation, 2006.
[14] "MokaFive, Virtual Desktops," *http://www.mokafive.com/.*
[15] R. Chandra, N. Zeldovich, C. Sapuntzakis, and M.S. Lam, "The Collective: A cache-based system management architecture," *Proceedings of the 2nd USENIX Symposium on Networked Systems Design and Implementation (NSDI'05).*
[16] M. Weiser, "How computers will be used differently in the next twenty years," *Security and Privacy, 1999. Proceedings of the 1999 IEEE Symposium on*, 1999, pp. 234-235.
[17] M. Jern, ""Thin" vs. "fat" visualization clients," *Proceedings of the working conference on Advanced visual interfaces*, L'Aquila, Italy: ACM, 1998, pp. 270-273.
[18] S. Kissler and O. Hoyt, "Using thin client technology to reduce complexity and cost," *Proceedings of the 33rd annual ACM SIGUCCS conference on User services*, ACM New York, NY, USA, 2005, pp. 138-140.
[19] M. Jern, ""Thin" vs. "fat" visualization clients," L'Aquila, Italy: ACM, 1998, pp. 270-273.
[20] S. Kissler and O. Hoyt, "Using thin client technology to reduce complexity and cost," New York, NY, USA: ACM, 2005, pp. 138–140.
[21] J. Golick, "Network computing in the new thin-client age," *netWorker*, vol. 3, 1999, pp. 30-40.
[22] S.C. Speer and D. Angelucci, "Extending the Reach of the Thin Client.," *Computers in Libraries*, vol. 21, 2001, pp. 46 - .
[23] P.A. Strassmann, "5 SECURE REASONS FOR THIN CLIENTS.," *Baseline*, 2008, p. 27.
[24] G.A. Plan, "An inefficient truth," *PC World*, 2007.
[25] M. Bruno-Britz, "Bank Sheds Pounds.," *Bank Systems & Technology*, vol. 42, 2005, p. 39.
[26] "Sun Ray White Papers," *http://www.sun.com/sunray/whitepapers.xml.*
[27] B.K. Schmidt, M.S. Lam, and J.D. Northcutt, "The interactive performance of SLIM: a stateless, thin-client architecture," Charleston, South Carolina, United States: ACM, 1999, pp. 32-47.
[28] S. Potter and J. Nieh, "Reducing downtime due to system maintenance and upgrades," San Diego, CA: USENIX Association, 2005, pp. 6-6.
[29] I. Wong-Bushby, R. Egan, and C. Isaacson, "A Case Study in SOA and Re-architecture at Company ABC," 2006, p. 179b.
[30] G. Reynolds and M. Gleeson, "Towards the Deployment of Flexible and Efficient Learning Tools: The Thin Client," *The Proceedings of the 4th China-Europe International Symposium on Software. China (Guanzhou). Sun Yat-Sen University. (2008).*
[31] "rdesktop: A Remote Desktop Protocol client."
[32] B. Childers, "PXE: not just for server networks anymore!," *Linux J.*, vol. 2009, 2009, p. 1.
[33] R. Likert, "A Technique for the Measurement of Attitudes," *Archives of Psychology*, vol. 140, 1932, pp. 1–55.