Thin Clients and PCs – A comparative study to find suitability for different computing environments

Tanmay K. Mohapatra

Choosing between thin clients and PCs requires a rational evaluation. Often a correct mix is optimum.

The personal computer (PC) no doubt has created a revolution in the field of computing. In recent years, thin client computing has provided an attractive alternative to the ubiquitous PC. However, both personal computers and thin clients have their own place in business organizations.

A large number of studies have been carried out to examine the technical aspects of various thin client technologies. But very few have proposed a set of holistic guidelines to be followed while making a choice of thin client technologies. In this article I examine the strengths and weaknesses of the PC and various types of thin clients that are available today, the technologies they use and their performance under different environments. Finally, I will propose a few guidelines for determining suitability of each for different computing environments.

A BRIEF HISTORY

With the advent of the PC the centralized computing model of mainframes and time-shared computing gave way to a more distributed model. Personal computers allowed users to use their local workstations to install applications, store data locally and also to do some amount of processing locally. Devices like floppy drives, scanners and printers could be attached directly to the desktop for alternate data input and output. The lower cost (than mainframes), popularity and ease of use of the PC played a major part in its proliferation. The large user base and its open architecture encouraged a large variety of software packages and devices to be developed for it.

The benefit of PCs did not come without some costs. System resources like RAM, CPU and storage had to be adequately supplied to each desktop to cater to the peak requirement. Devices like floppy drives, CD drives, and printers needed to be

connected to individual desktops. Since all desktops were not utilizing all their resources at all times it resulted in redundant and unused resources. Softwares needed to be installed on each individual desktop, resulting in upgrade nightmares. Security control also became an issue, the PC being an open platform and it being under the control of the usually less technically informed end user. Control over software licences became difficult, as users were free to install copies of softwares and also sometimes unauthorized softwares on desktops. At the same time, in larger business and enterprise environments, centralized systems were still used for permanent and shared data storage through databases, and to perform resource intensive operations, as the PC was not capable of such things. Thus apart from the cost of each individual unit being more, larger centralized machines were still required and maintenance costs were higher.

The thin client is an evolution of the PC towards a similar but simplified device. It aims to reduce both the cost of the unit as well as the operating cost. Thin client is a relatively loosely defined term that refers to PC like devices that consist of essential components like the VDU, keyboard and mouse, but they either minimize on or eliminate components like RAM, CPU, disk drives, and CD/floppy drives. Instead they utilize these resources from a central server and multiple thin clients share the same resources. The thin client is therefore cheaper. It is also more reliable because of simpler construction and less number of components (moving mechanical components) that are contained in it. In the thin client model, though the display is on individual desktops, actual applications are executed and data is stored on a central server.

Are we going back to the era of mainframes and time-shared computing where a large server used to serve multiple users operating from relatively dumb terminals? Well, not entirely. Today's business and enterprise computing environment is challenging. The end user needs to use applications with graphical data representation and needs to interact not only using the keyboard but also a plethora of new devices like smart card readers, digital cameras and specialized devices like bank check scanners. Being relatively new technology, thin clients lack the wide variety of software and device support that the PC platform enjoys. The various options available, both between the PC and the thin clients and between the different types of thin clients, often necessitate a detailed study of the requirements vis-à-vis the choices available before making any decision.

THIN CLIENT TECHNOLOGIES

There are quite a few thin client technologies in the market today out of which I have taken up some to be examined.

Windows-Based Terminals

Windows based terminals (WBTs) leverage on the multi-user capability of Windows NT that allows a single Windows server to support multiple users simultaneously. They were introduced first with Windows 3.51 by a product called Winframe (later called Metaframe) from Citrix Systems. Microsoft included this capability in Windows NT 4.0 Terminal Server by incorporating some of the extensions developed by Citrix into the operating system, and it has subsequently become a bundled feature of Windows operating systems (2000/XP/2003).

The Citrix Metaframe is based on the Independent Computing Architecture (ICA) protocol whereas Windows terminal server uses the Remote Desktop Protocol (RDP). Though both perform similarly they do differ in the underlying protocol and in the additional features provided. ICA greatly enhances Windows terminal server's capability through advanced features like stereo Windows audio, direct dialup. The ICA protocol is supported over UNIX platforms through Citrix Metaframe. RDP has an open-source client called "rdesktop" (http://www.rdesktop.org/) that runs on UNIX platforms.

Network Computer

The Network Computer or NC was developed by IBM, Oracle and Sun Microsystems. It is a device attached to an application server from where it downloads Java applications to be run locally. This fact, that applications are actually executed on the NC and not on the server is the most significant difference between NC and other display only terminals. Because of this, an NC does require adequate CPU and RAM on the client as per application needs. Application installation and data storage is done at the server. Some NCs also support protocols, such as X Window and ICA, which allow them to be used as display only terminals for programs running on a server. Being a Java based solution, this is ideal for Java applications. It can also be used for browser based Internet applications as NCs usually come bundled with a browser.

Tarantella Thin Clients

Tarantella leverages on Java and browser technologies to provide thin client capability over a diverse range of client environments – Microsoft Windows PCs, UNIX or Linux Workstations, Internet devices and Network Computers.

In the Tarantella architecture ("Tarantella ® Enterprise 3 ™ Software – A Technical Overview", Tarantella White Papers, July 2002, http://www.tarantella.com/whitepapers/), each user after logging on to the system has a browser based menu of applications called 'WebTop' that lists multiple applications available under the logged in account. On invoking an application, a Java based screen is brought up that connects to one of multiple Tarantella servers to invoke a new

application instance for the account. Tarantella also provides a native client option, which can be installed on the client machine and used instead of the webtop. The Tarantella architecture consists of multiple Tarantella servers to which desktop clients connect. The Tarantella server in turn connects to the application server and logs in as the appropriate user before executing the application. The application server and tarantella server communicate through a 'Protocol Engine' that is different for each different protocol used – e.g. there will be a protocol engine for RDP that can talk to Windows based applications and another for X11 that can talk to UNIX based applications.

The communication between the Tarantella server and the desktop happens through a protocol called Adaptive Internet Protocol (AIP). The AIP protocol measures various parameters on the fly and changes itself appropriately to adapt to the current conditions so that it makes the best use of current conditions. Some of the parameters that it can vary during execution are – color depth, color quality, interlaced images, graphics acceleration, frequency of updates and command compression.

AT&T Virtual Network Computing

The Virtual Network Computing (VNC) is a remote display system for desktop environments. An open source implementation is available. It is based on the Remote Frame Buffer (RFB) protocol. In the RFB protocol, a frame-buffer is maintained by the client, which it periodically updates by fetching data from the server.

Since the refresh is based on client pull rather than server push, it is self-adjusting to network conditions and machine loads. For example, if the network is slow, the update request from the client will reach the server after more time and the update data from the server will contain more changes. This allows for more and more changes to be sent with less and less update requests. The suppression of intermediate changes and transmission of only the final state, results in bandwidth savings.

The RFB protocol can use different mechanisms for encoding the frame-buffer data sent. Raw encoding is the simplest form of encoding where the pixel data in left to right order is sent. This requires minimum processing on the server, but is heavier on the network. The Copy rectangle encoding is used in operations like scrolling or moving of a window where the client already has the pixel data and it just needs to move it to a new location. In addition more complex encoding like Rise and Run length Encoding (RRE), CoRRE encoding and Hextile encoding (Tristan Richardson, Kenneth R. Wood, ORL Cambridge, "The RFB Protocol – Version 3.3", http://www.uk.research.att.com/vnc/rfbproto.pdf) can be used where suitable. RRE is essentially a two-dimensional analogue of run-length encoding that is used to compress and encode rectangles of pixel data.

Linux (X Windows) Clients

The X Windows protocol, which is predominantly used in the UNIX world for graphical displays, was developed by the X consortium led by M.I.T. (The X11 protocol, http://www.x.org/X11_protocol.html). An X Windows setup has two major components – an X client that sends the UI update messages and an X server that displays on the screen and also sends user inputs back to the client. The X protocol uses reliable TCP/IP connections using port numbers beyond 6000. It has the reputation of being a fat protocol, consuming a lot of bandwidth. The X Windows protocol has been profiled (John Danskin, Pat Hanrahan, "Profiling the X Protocol", May 16 1994, http://graphics.stanford.edu/papers/profiling/) to analyze the network resource requirements for different message types and many attempts have been made to minimize the network usage of X Windows protocol. The introduction of lightweight X (LWX) in X11R6 is designed to reduce the bandwidth requirements to a point where it can be used even over a dialup modem line. XFCE (http://www.xfce.org/) is a desktop environment developed using this lightweight protocol.

The protocol is based on high-level graphics commands, unlike other protocols. It uses an eager server-push model for sending the display commands. Communication in X is asynchronous; that is, the client does not block to receive a response from the server once it has sent some command. It is possible to send several commands in a sort of batch mode and receive responses for them at a later point of time. Not all commands require a response from the server. The X Windows protocol can support up to 24-bit color.

Sun Ray

Sun Ray is a thin client product from Sun Microsystems that provides access to applications running on multiple operating systems, including Solaris, Microsoft Windows NT and other UNIX platforms. The Sun Ray client uses a proprietary protocol to communicate with the server. Some attempts have been made to reverse engineer the Sun Ray protocol (Paul Evans, "SunRay Protocol Documentation", 2003, http://www.srcf.ucam.org/~pe208/uni/work/sunray/prot.ps). When started up, the client obtains the address and port number (typically 7009) of an authentication server through DHCP. Then it connects to the authentication server over TCP and exchanges certain sequences of authentication data. After authentication, the communication of display data happens through UDP packets. UDP, being connection less, is lighter than TCP but does not offer the reliability of TCP. The Sun Ray protocol uses 3D drawing primitives, similar to the VNC RFB protocol with screen refresh based on an eager server push model. It can support up to 24-bit color.

COMPARISON

The display mechanism is key in thin client technologies; the protocol used determines the kind of display it is suitable for and the amount of resources required for it to perform. Thin client protocols to some extent need to provide support for external devices and peripherals. Usability with external devices and peripherals is another key parameter that decides the activities it is suitable for; thin clients do not yet have very good support for devices and peripherals when compared to the PC. The cost, of course, is a factor to be considered while taking decisions.

Display Quality and Resource Requirements

Display protocols play an important part in two critical areas. The first affected parameter is the display capability; certain protocols are tuned towards certain types of displays. Video display involves rapidly changing screens and colors whereas an office productivity application would mostly involve static screens. The second affected parameter is resource requirements like CPU and bandwidth. CPU resources on both server and client machine can be critical. Table 1 lists the display protocols that are used by the thin clients discussed in the previous section.

Thin Client	Display Protocol
WBT	RDP, ICA
Tarantella	AIP
AT&T VNC	RFB
Linux (X Windows)	X
Sun Ray	Proprietary

Table 1. Display protocols used by different thin client systems

Yang, Nieh, Selsky and Tiwari ("The Performance of Remote Display Mechanisms for Thin-Client Computing", *Proceedings of the 2002 USENIX Annual Technical Conference*, http://www.usenix.org/events/usenix02/full_papers/yang/yang.pdf) and Nieh, Yang and Novik ("A Comparison of Thin-Client Computing Architectures", *Technical Report CUCS-022-00*, Network Computing Laboratory, Columbia University, November 2000, http://www1.cs.columbia.edu/~library/TR-repository/reports/reports-2000/cucs-022-00.pdf) have done extensive measurements of performance of thin client display mechanisms. According to their findings in the Web and Video performance test, X and AIP protocols are the fastest given adequate network bandwidth (100 Mbps). Reducing bandwidth has the biggest negative impact over X and Sun Ray. The ICA, RDP, AIP and VNC protocols were able to deliver sub second response over bandwidths as low as 768 Kbps and most protocols failed to deliver sub second performance over bandwidth lower than 128 Kbps. However, overall

performance at lower bandwidths is good with ICA and RDP protocols. Video performance is poor at lower bandwidths in all display protocols. AIP, X and Sun Ray are more able to support a broader range of applications, particularly multimedia applications. ICA, RDP and VNC are quite bandwidth efficient for web applications, but not quite as good in video and multimedia performance. Caching and compression are useful in low bandwidth situations where the extra CPU power consumed can be justified by savings in scarce network bandwidth.

In Table 2 below I attempt to rate the PC and thin clients based on display quality and performance attributes. The NC has not been rated in this table as the performance and capabilities of NC depend to a large extent on the protocol implemented by the NC and the kind of application that is being used. Similarly, Table 3 summarizes the computing and bandwidth resource requirements of PCs and various thin client platforms.

	Video Quality	**Web Page Performance (Minimum Latency)**
PC	Excellent	Excellent
WBT	Fair	Fair
Tarantella AIP	Good	Good
AT&T VNC	Poor	Fair
X Windows	Excellent	Good
Sun Ray	Good	Fair

Table 2. Display quality and performance comparison

	Bandwidth Efficiency	**CPU Efficiency at Client**	**CPU Efficiency at Server**
WBT	Good	Excellent	Fair
Tarantella AIP	Fair	Fair	Good
AT&T VNC	Excellent	Poor	Fair
X Windows	Fair	Excellent	Excellent
Sun Ray	Poor	-	Fair

Table 3. Resource efficiency comparison

Protocol Capabilities and Device Support

With the PC as a desktop attaching a local device is much easier. It being the most widely used desktop has a wide range of device drivers available. The applications accessing the devices execute on the same hardware to which the device is attached, which makes things easier.

Using local devices with thin clients requires support from the display protocol, the application software and device drivers. Though most thin clients support local drives and local printers, the range of devices supported is not extensive and overall support for locally connected devices is poor. Multimedia support is another achilles heel of thin client systems.

The Windows XP based thin client is pretty much full featured. Citrix Metaframe is also a matured product in thin client arena and it provides many advanced featues like COM port mapping, stereo sound and clipboard redirection. However, Windows based thin clients require higher initial investment. Linux is cheaper and less resouce hungry but has poor support for multimedia and peripherals ("Analysis of Thin Client Operating Systems Windows CE, Embedded XP, Windows XP, and Linux", NCD white paper, 2003, http://www.ncd.com/news/2003/ThinClientOScomparisons.pdf).

Installation and administration costs

Thin clients are simpler to install and setup than a fat client. In most cases it will just require one to plug in the power and the network cable. With simple hardware and minimal moving parts they are very less likely to breakdown compared to PCs. Even when a thin client does break down, it is very easy to replace the hardware with another piece of hardware. Down time is minimal and productivity is higher; as data is stored on the server the user can resume working immediately.

With thin clients, operating system and software upgrade are one-time activities to be done on the server compared to multiple installations on each individual PC. Thin clients systems have better longevity than fat clients, and do not require the desktop hardware to be upgraded frequently. Though thin clients provide lot of savings in administrative costs, the initial cost of many thin clients may be quite high and may be as high as that of standalone PCs. Many existing but outdated PC hardware can be converted to thin clients thus increasing their lives by many more years.

Energy and space saving

Thin clients are much smaller in size than PCs. In places where desk space is at a premium, thin clients can come as a savior. Since they do not have many heat generating parts, many of them implement convection cooling mechanisms and do no

require fans. This makes them energy efficient and practically noiseless. Measurements done (Wyse Technology Inc., "Desktop Energy Consumption – A comparison of Thin Clients and PCs", Whitepaper, September 2001) indicate that thin clients can consume as low as 25 watts of energy compared to 170 watts consumed by a PC. In a big organization with a lot of terminals, power savings by using thin clients can be significant.

Security

By their very nature, thin clients are more secure than fat clients. Though the user can access the data from any place, the data is stored on the central server only. A thin client device stripped of external storage devices (CD ROM, Floppy drives) does not allow copies of data to be taken without authorization. It also prevents viruses being inadvertently introduced into the system and incompatible or unauthorized softwares being installed by users.

CHOOSING THE RIGHT PLATFORM

The decision is two pronged. First you will have to evaluate whether to use PCs or to go for thin clients. If thin clients satisfy your requirements better, a second evaluation needs to be done to determine the type of thin client to be used. The key to the decision is likely to be arrived at from the answers to a set of questions.

PC or Thin Client?

Whether to go for PCs or for thin clients can be decided from the following questions.

What kind of usage will the desktop be put to?
Thin clients do not yet enjoy the kind of widespread support from different hardware devices and software packages that the PC platform does. Therefore, if the requirement is extensive use of external devices and usage of a variety of software packages, then the PC platform will be a better choice. Similarly the PC platform will be more suitable for high graphics and video based applications, where the performance of thin clients has been found to be bad because of higher resource usage, lower quality of video and slower operation. However for usage as an Internet browser and word processor or simple applications or browser-based applications the thin clients may be an ideal solution.

What kind of network connectivity is available between the terminal server and the desktops?
Thin clients require reliable networks. Since data is stored and applications are

executed in the central server, a disruption of network will result in loss of productivity. In contrast to the thin clients, a PC allows local applications to be used. The local applications can serve as a backup in case of network failure and allow the user to continue, albeit with limited functionality.

How many terminals are to be used?
With a large user base, thin clients provide substantial savings in running costs including maintenance costs and energy costs.

Which Thin Client?

If a thin client platform is to be used, the next set of questions will most likely be asked to decide on the type of thin client that can be used -

Is it going to be used in a LAN or WAN environment?
A thin client based on X Windows and Sun Ray has higher bandwidth requirements than others. Thin clients based on the ICA and RDP protocol are the least bandwidth consuming. Another option could be to have a de-centralized thin client based system where thin client servers are deployed in each region. Applications running on regional servers then refer to central servers to access common data stores and services.

What is the software licensing cost?
Commercial thin client software is usually charged with a per-user licensing model. An X Windows based thin client can be setup using open source software like Linux. An open source client for RDP protocol is also available for use as a Windows thin client, though that will still attract per user licensing fees for the Windows server.

What is the hardware cost?
Is it possible to re-use existing hardware? Some thin client systems like Windows based terminals, Tarantella AIP and Linux based X Windows allow existing but possibly outdated PCs to be re-used.

What kinds of applications are to be used by the users?
If the applications are predominantly UNIX based, then Sun Ray, Linux X Windows, or NCs can be used. Whereas if they are predominantly Windows based (e.g. MS Office, VB applications), then Windows based terminals or Tarantella AIP will be more suitable.

CONCLUSIONS

Thin clients are an attractive alternative to PC based systems and quite a few different thin client architectures are available today to choose from. However the thin client cannot entirely replace the PC platform. It is important to take a closer look at the requirements and based on the features decide on the feasibility and the closest fit.

With the kind of myriad requirements in large organizations it may not be possible or appropriate to satisfy all requirements completely by any single platform. A judicious mix of both PCs and thin clients can often turn out to be the optimum computing environment for the organization.

Tanmay K. Mohapatra *is a Technical Architect associated with the Banking Products Division of Infosys Technologies Ltd., Bangalore, India.*