

Renewed Computer System

Robert Pogson

2008-11-3

Abstract

The youngest computers our students use are 5 years old. The oldest are 10 years old. Boot times are 45s to 120s and it takes 2 to 3 minutes to obtain a usable desktop. The operating systems used, Windows 2000 and XP, although widely used were obsolete when they were released, having little regard for security. The patches that now render them somewhat secure also slow them down. The components that make up these PCs are several times slower than current technology. These PCs make more noise, consume more power, collect more dust, take up more space but run slower than current technology. There simply is no good reason to keep them. The purpose of information technology is to find, create, change and display information as efficiently as possible and these machines are barely able to do that. It is proposed to replace the whole system with newer technology.

Contents

1 Hardware	1
1.1 Plan A - install new thin clients and LCD monitors	3
1.2 Plan B - use the present machines as thin clients	3
1.3 Preferences	3
2 Software	3
2.1 Applications	4
2.1.1 Office Productivity	4
2.1.2 Multimedia	4
2.1.3 Web applications	5
2.1.4 Computer Science	5
2.1.5 Storage	5
2.1.6 Networking	6
3 Conclusion	6

1 Hardware

The performance of the hardware is poor by any measure:

Old	Current
Memory bandwidth 200-1000 MB/s	10000 MB/s
CPU speed 800-1600 MHz 32 bit	2-3000MHz 64bit
Bus bandwidth 800 MB/s	20000MB/s
Hard Drive transfer 20 MB/s	60 MB/s
Storage capacity 40 GB	500 GB
Boot time to useful desktop 2-3 min	25s

The only way to get decent performance from these machines is to ask them to do less than the present workload. It is possible to obtain one or two powerful new machines to run all the applications in the school and to use the old machines to show the pictures on the screen and to receive the clicks from keyboard and mouse[1]. This arrangement is called thin client/terminal server. The performance per dollar and performance is superior to what we could get by replacing all the PCs with new ones because the servers can afford much more powerful disc storage technology, can use far more and faster memory than a PC, and can use more processors

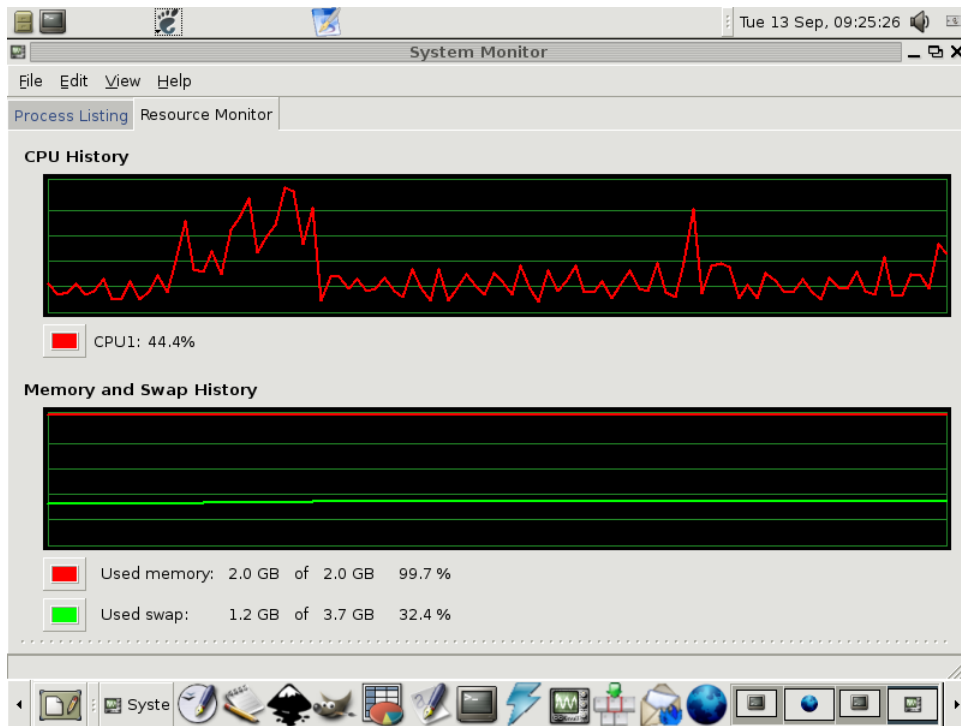


Figure 1: 17 users on a single core 64bit CPU which also ran databases and web applicaitons



Figure 2: e-Box 2300, one of the least expensive thin clients made today. It is a little slow to boot but fits on the back of most LCD monitors and uses 15 watts only.

simultaneously. A server that is affordable can have four or eight processors, each more than capable of running a whole lab.

Using the present machines as thin clients will work but the drawbacks other than performance still remain: heat, noise, dust collection and clutter. The performance of the system would be determined by the performance of the server which can run many thousands of applications at once at full speed (in human terms). Humans click on something, get a picture and stare at it for a few seconds before clicking again. The server can supply such service to many hundreds of users at once. Where full-screen video is needed by many users, regular thick clients should be used but for browsing, word-processing and running many applications, thin clients do very well because we take time to absorb information from the screens.

In 2006, a system of terminal servers was built to run the K-12 school at Easterville, MB. Four servers costing \$1200 each could run the whole school of 153 PCs, doing up to 60000 context changes per second while users were getting full service. CPU load was rarely over 20%. That was with dual-core CPUs on the server and DDR memory. These days we have quad-core CPUs and DDR2 memory at lower prices. A single server could do the job for about \$3000 today. We paid \$139 for small thin clients then. Similar units cost about \$100 today. Smaller units cost as little as \$80.

It would be feasible to take either of two strategies for the system. Either install new thin clients, or use the old machines as thin clients. The feasibility of the former approach has been verified in the lab in Room of the high school by using an 8 year old machine as a thin client of a 4 year old server. The performance is outstanding. New thin clients would have the benefit of smaller footprint, noiselessness, dustlessness, and

lower power consumption. Either approach should be supplemented by portable/notebook PCs with projectors for video. The approach of replacing old machines with similar new machines costs twice as much for lower performance and guaranteed obsolescence in a few years. Thin clients can show pictures for a decade or longer.

1.1 Plan A - install new thin clients and LCD monitors

	Quantity	Cost
New thin clients	100 @ \$100	\$10000
Terminal Server with quad CPU, 16GB RAM, 3 TB storage	2@\$2000	\$4000
New 17 inch LCD monitors	100@\$150	\$15000
8 port network switches	18@\$40	\$720
24 port network switches with 2 gigabit/s ports	2@\$100	\$200
	Total	\$30000

1.2 Plan B - use the present machines as thin clients

	Quantity	Cost
Terminal Server with quad CPU, 16GB RAM, 3 TB storage	2@\$2000	\$4000
24 port network switches with 2 gigabit/s ports	2@\$100	\$200
	Total	\$4200

1.3 Preferences

Plan A is preferable as it is lower maintenance, less noise, less space, power and dust. The performance would not be better than Plan B. It depends what value is placed on the benefits of a machine designed to be a thin client rather than using a thick client as a thin client. One could use Plan B as a trial or phased in Plan A. One could also scrap the CRTs and use LCD monitors with Plan B. That would save a lot of space and power.

2 Software

The software we use is descended from NT from 1993 onward[2]. It was developed by people who had worked on the earlier VMS and OS/2 systems. In the beginning, it was a decent operating systems based on sound principles. However, the marketing droids in Microsoft completely corrupted the system by insisting so many interlocking features be added that security was seriously compromised. Waves of malware swept the IT industry as a result[3]. Our system has only remained workable by constant patching and occasionally, disconnection from the Internet. There is no certain means of protecting our system.

At the same time as the security has been weak, features have been added with each “Service Pack” rendering the system ever slower. Windows constantly fragments its disc storage and memory so that the more software packages and the more processes that run, the slower the system operates.

Better software is based on UNIX operating system principles: shared memory for libraries and applications saves a lot of memory and precise separation of roles provides security. When Windows was designed, it was a single-user with no security. NT was better, but integrating the web browser into the operating system threw away all that security because the browser can easily “execute” an image or media file and give it total control of the machine. Also, by the time of XP, sharing was virtually eliminated from Windows[4]. Microsoft could not implement any kind of security with shared memory. UNIX could. No amount of patching can fix a browser built like that. That feature was built into Windows to provide features Netscape could not in order to stifle competition.

There are several UNIX operating systems which are feasible:

- OpenSolaris - SUN Microsystems - very robust and efficient on servers but lacking many drivers and applications needed by desktop operating systems
- FreeBSD - derived from Berkley Software Distribution - renowned for stability on the server but, again, limited by available drivers and applications
- GNU/Linux - written from scratch by the Linux Kernel group of Linus Torvalds and many others and the GNU organization started by the organizer of the Free Software Foundation, Richard Stallman

Of these, only GNU/Linux has the broad base of contributors and users to make it reliable in desktop systems. It is used on about 100 million desktops around the world[5]. Many thousands of developers[6] from every corner of the world contribute software under a free licence, mostly the GNU Public Licence which permits use, modification, examination and distribution under the same terms. There is no charge for use of this software and there is no charge for obtaining updates over the Internet from distributions such as Debian. The collection of software for GNU/Linux is amazing. There are over 100 thousand free software projects and the Debian distribution has over 20 thousand packages built and tested for mutual compatibility[7]. It may take only seconds to install an application using the Debian packaging system. In the lab, I have students install on a few old machines a minimal operating system from Debian. It takes 30 minutes and the machine can be used as a thin client. If we were to install in bulk we could use a networked image transfer system such as Clonezilla to do many at once.

2.1 Applications

2.1.1 Office Productivity

This covers most of what I see done in the labs: generating, finding, organizing and publishing information. We have some very popular packages here:

- OpenOffice.org - a full-power office suite originally intended to replace MS Office and it still does that well but also many find it superior. SUN Microsystems created OpenOffice when they decided to buy a company producing StarOffice instead of buying 20000 licences for Office. It was cheaper. OpenOffice provides superior productivity. PDF generation and automatic indexes are just a couple of the wonderful features in the software. More than 100 million installations exist.
- FireFox - a powerful web browser with features superior to earlier versions of IE. 30% of browsers seen by servers on the Internet use FireFox.
- LyX - a smooth type-setting application that makes writing a pleasure. This document was produced with it. One concentrates on writing and leaves the type-setting to LyX.
- Gimp - a powerful image processing application does everything except colour matching.
- xpdf - a good PDF reader.

2.1.2 Multimedia

Although thin clients are not the best for video, they will do for part-screen stuff as seen while browsing, and they can do audio very well.

- XMMS - a player for MP3s and such
- MPlayer - a player for audio and video like Flash and MP3 and MPEG
- Audacity - a sound editor. Cutting and pasting sounds... Making tracks. Mixing them.
- Hydrogen - a sound synthesizer for percussion, beats and such (I am not musical...)
- Zynaddsubfx - another music synthesizer
- Ardour - a wonderful sound mixer
- Cinepaint - video editing
- Blender is an integrated 3d suite for modeling, animation, rendering, post-production, interactive creation and playback (games).
- Xawtv - is a video viewer that works with webcams

2.1.3 Web applications

Windows does not have these, usually, but we can install on our terminal servers web applications like an interactive website from the Internet. The users access the web application via the browser and there are no network delays. The slowest access is perhaps 2s.

- phpMyAdmin - a general database access tool for MySQL
- phpBB - a bulletin board with moderators and all that
- MediaWiki - the same software that powers Wikipedia.org can be used in a school for collaboration/accumulation of a knowledge base and it is searchable
- SWISH-e - a very fast search engine can be accessed from ordinary web pages by CGI or PHP scripts to find documents by keyword and phrase, just like GOOGLE. It can index and search word-processing and PDF documents as well as text and HTML.
- Gallery - a database for images. Annotations are searchable. A million images is no problem.
- java applications and javascripts and PHP scripts made in-house or downloaded

2.1.4 Computer Science

GNU/Linux is much better than Windows for CTS - Information Processing, particularly in language processors and library for software development. Students can easily work on desktop machines, servers or their own machines all with the same software. We do not have to count licences so they can install as many copies as they wish.

- Programming languages - java, perl, pascal, php, python, C, C++, FORTRAN, Modula-2, etc.
- Databases - PostgresQL, MySQL, DB3
- Services - Apache web server, ftp, tightVNC, OpenSSH (allows secure file transfer and remote control of clusters of machines), e-mail, DNS, OpenLDAP, NTP, NFS, Samba, web proxy, web caching and filtering
- Virtual machines - permit students to install on the working server without affecting anyone else on the system
- Clustering
- Networking
- Searching
- Sorting

All these are standard applications on GNU/Linux machines but they permit CTS students to explore IT/Computer Science with no extra costs.

2.1.5 Storage

GNU/Linux can use complicated storage arrays for redundancy and speed that are not available on a desktop system usually. Instead of one hard drive in the box, we can have six or more in each of two servers giving us much more speed of access and storage capacity. Presently, our systems have to access the shared files over a 100 mbits/s network cable (12 megabytes/s). On the terminal servers, our applications and processes could access local hard disc storage and cached files with no network delay or collision which is a constant annoyance now. We would keep all that traffic off the network. By having many more hard drives in the system, the probability of any file being near a read/write head is much higher and seek times will be lower and transfer rates higher. We can read from multiple hard drives at once for greater throughput. Two hard drives may be used as one drive with twice the normal throughput and three pairs may store identical copies of files for redundancy[8]. More than two hard drives would have to fail to interrupt services. Using 500 gB to 1000 gB hard drives would give us about ten times our current storage capacity with much better performance. A pair of big drives can transfer 120 megabytes per second, ten times our present speed (over the network).

2.1.6 Networking

The Internet runs on GNU/Linux mostly because of its prevalence in routers and servers. We could add amazing capabilities to our terminal servers like:

- a website for every user account
- interactive websites limited only by imagination
- teachers interacting with students remotely
- chatting by text
- locally controlled e-mail...
- kid filtering with DansGuardian[9]
- CloneZilla[10] for doing Install Fests (similar to Ghost but free software)
- complete local repository so no need to access the Internet to download and install software

3 Conclusion

This article outlines what is possible with a modest expenditure using modern technology. It could be implemented as PLAN A, PLAN B, a combination or phased in. It makes no sense to continue operating at this low level of performance when an order of magnitude improvement can be had for so little expense. A new system would be much lower in maintenance and easier to update.

If, for whatever reason, it is not possible to request funding for such a project, there should be at least a serious effort at improving our school in this area by forming an IT planning/advisory committee in the school so that when the time is right, a plan will be ready for implementation.

References

- [1] Linux Terminal Server Project at <http://ltsp.org/> "The Linux Terminal Server Project adds thin-client support to Linux servers. LTSP is a flexible, cost effective solution that is empowering schools, businesses, and organizations all over the world to easily install and deploy desktop workstations. A growing number of Linux distributions include LTSP out-of-the-box. – Shiny new thin-clients and legacy PCs alike can be used to browse the Web, send e-mail, create documents, and run other desktop applications. LTSP not only improves Total Cost of Ownership (TCO), but more importantly, provides increased value over traditional computing solutions. LTSP workstations can run applications from Linux and Windows servers. – Linux thin-clients have proven to be extremely reliable because tampering and viruses are virtually non-existent. It's distributed under the GNU General Public License, meaning it's free and always will be. "
- [2] Article, "Windows NT" in Wikipedia at http://en.wikipedia.org/wiki/Windows_NT . Note the original requirements for 2000 and XP, which we use, were 32 and 64 MB of RAM. We now use 1024 MB because so much has been added over the years.
- [3] http://eval.symantec.com/mktginfo/enterprise/white_papers/b-whitepaper_exec_summary_internet_security_threat_2008.en-us.pdf *"the release rate of malicious code and other unwanted programs may be exceeding that of legitimate software applications."*
- [4] http://www.winsupersite.com/showcase/whistler_fusion.asp *"the OS will actually be copying the shared DLLs and other files that the application installs into a private location, so that the application will "think" that it is running normally, as it would on a legacy Windows system"* There is very little security in this behaviour and it throws out all the advantages of sharing, mainly much reduction in memory requirements. Vista is the extreme example, requiring 8 times as much memory to run more slowly than XP and 32 times as much memory as Windows '95. Tests on Windows terminal servers show that twice the memory per process is required. This permits UNIX-like operating systems to run twice as many processes on a server as do Windows operating systems.

- [5] Prevalence of GNU/Linux on the desktop: Web statistics from the English-speaking world range from 1 to 4% at <http://marketshare.hitslink.com/report.aspx?qprid=11>, http://www.onestat.com/html/aboutus_pressbox54-windows-vista-global-usage-share.html (2007 data), http://w3schools.com/browsers/browsers_os.asp, <http://sec.gov/Archives/edgar/data/320193/000119312508156421/d10q.htm> (Apple's unit sales per quarter are 2.5 million machines, about 3% of PCs sold in a quarter), <http://www.idc.com/getdoc.jsp?containerId=prBR21102808> "*As a result of 10.7 million PCs shipped in 2007, Brazil has elevated its global position in the PC market to the 5th largest market in the world. Behind the United States (64 million), China (36 million), Japan (13 million) and the United Kingdom (11.2 million), and far ahead of India (6.4 million, 9th place),*" The government of Brazil actively promotes GNU/Linux as do the governments of Russia, and China. China has as many users of PCs as the USA but they share PCs. 24 million "netbook" PCs are expected to ship in 2008, of which half use GNU/Linux. On Google Trends, Linux gets one-third the hits of Windows, Ubuntu gets $\frac{1}{8}$ the hits of Windows, and $\frac{1}{4}$ the hits of Vista at <http://google.com/trends?q=Linux%2Cubuntu%2Cwindows%2Cxp%2Cvista&ctab=0&geo=all&date=all&sort=0>
- [6] Sourceforge hosts more than 100000 free software projects at <http://sourceforge.net>
- [7] <http://debian.org> the largest and one of the oldest distributions of GNU/Linux software.
- [8] RAID, Redundant Array of Independent Discs, allows increased reliability and speed from hard drives. see <http://en.wikipedia.org/wiki/RAID>
- [9] Dansguardian a kid-filter that can blacklist/whitelist/scan for malware/scan for undesirable content and cache files from the Internet. <http://dansguardian.org/>
- [10] DRBL - Diskless Remote Boot Linux - provides the possibility of booting thick clients from the server and cloning PCs with Clonezilla. see <http://drbl.sourceforge.net/>