

Arizona State University West Library Linux Deployment Executive & Technical Summary

By Perry Horner

Coordinator of Library Technology Support & Development (p_horner@asu.edu)

Introduction

The ASU West Library Technology Support and Development (TS&D) department has designed and implemented a successful and innovative computing architecture model to fulfill the Library's need for flexible application offerings, guaranteed reliability, and cost containment. Historically the Library has replaced its public computing hardware every three years, including semester and annual updates and changes to supporting software. These three-year cycles provided an opportunity to assess future needs for the next cycle replacement. In 1999 the Library was introduced with a proposal by its technology department to make a drastic change in its core computing architecture- migrate everything from Microsoft Windows to Linux. This bold proposal did not stop with just an operating system change; it introduced a hybrid concept of merging the best of thin-client technology (dumb terminals) with fat-client reliability (local hard drive OS and application storage) at an affordable cost. The project proposal was named "E3" identifying that this action would be the third evolution of the Library's computing architecture. The foundation E3 established using a Linux platform enables the Library to tap into the open source collective to complement the Library's existing expertise, address users needs with fewer barriers towards solutions, and allow for creative endeavors that were not available while using Microsoft Windows. The following details the proposal's charter.

Business Problem

Long-term Library technology support funding is not stable. As the Library continues to fill existing vacant positions to meet the growing needs of the campus community, flexible spending from salary savings to support Library technology will diminish to a level making current technology funding requirements unsustainable. Software licensing costs and new hardware costs continue to rise due to pricing or quantity. The Library cannot continue to sustain its technology infrastructure using its current architecture. Nor can the Library take full advantage of its technology resources using the current architecture.

Statement of Work

- Develop and implement a new computing architectural model that will be sustainable over the long-term and facilitate creative developments to meet business goals.

Project Objectives

- Migrate both client workstations and servers from Microsoft Windows to Linux
- Develop a user workstation experience that contributes to individual goals
- Reduce Total Cost of Ownership (TCO) by 50%
- Develop a foundation for long-term technology infrastructure planning
- Provide a long-term technology funding strategy

Success Factors

- End-user workstations provide core Library resource services
- TCO reduced compared to previous technology fiscal cycles
- Diskless workstation environment improves reliability and performance

Project Dependencies/Constraints

- Requires dedicated network segment
- Library is the sole end-to-end support provider for equipment and software

Architecture Explained

The new library workstations offer ASU and non-ASU users access to Internet resources using the Firefox browser, access to floppy, Zip 250, and AFS network storage, Microsoft file format viewer applications through the use of WINE, a feedback agent, station usage and tracking agent, timeout screensaver, PDF viewer, and a disk management tool. The workstations take advantage of network PXE booting and receive

a clean disk image into a compressed 1GB RAMdisk. The workstations utilize Kerberos authentication and LDAP directory services for accounts. Since the disk image is in memory, once the power is dropped to a workstation, or the reset button is pressed, all footprints are erased. Thus, any drive corruption, cookies, history, etc. is removed with a single press of the reset button. When the workstation starts up, a clean image is served to the workstation with a boot time of less than one minute. If a user logs out, all history and files created or changed at login are removed and restored to their original state. This process creates a fat-client workstation through a thin-client discovery process. Since the client machines do not have a hard drive, access to the RAMdisk is using memory access pipes and not standard I/O pipes. This makes disk access faster than any hard drive available. By eliminating the hard drive, a major point of hardware failure is removed as well as a component that generates large amounts of heat that could harm other system components. At the time of procurement, traditional vendor offerings that met the necessary specifications were cost prohibitive with current funding- workstations @ \$1,100 per unit, servers @ \$18,000 per unit. After a build vs. buy analysis, it was determined that the Library could build its own workstations at less than \$500 per unit and servers at less than \$5,000 per unit using readily available, industry standard, off-the-shelf components. After training, student employees were able to assemble 71 workstations at a rate of 30 per day for an hourly wage cost of only \$312 dollars total for all labor. For troubleshooting purposes, each client station reports its system log to a system log-caching server. The Library has a Public Interface Team that designs the “look and feel” user interface. TS&D then developed and implemented the interface based on guidelines and criteria set forth by the Library’s Public Interface Team. Each client workstation allows the user to read and save to both Iomega ZIP and floppy disk removable media. These media devices have permissions set to deny Execute requests. A patch to the Linux kernel called *AutoMounter* allows a user to insert and remove these media at will. Any writing to the devices takes place immediately and does not reside in a RAM cache. The user also can press the RESET button on the workstation faceplate. This will start a warm reboot. The user cannot use the POWER button. This has been disabled through the use of a key-lock switch on the 5 ¼ bay faceplate.

Boot Servers

Two boot servers share the boot request services through distributed load balancing and fail-over. Each server has a list of primary/preferred clients. If the services are not available to a client from its primary server, the other server will receive the request and fulfill it. Security is managed using a combination of allowable MAC addresses and intrusion detection software. Although a MAC address can be spoofed, the client boot image is matched to the unique hardware of the Library Public Access workstations. This means the spoofing workstation must have the same network card, video card, RAM size, and storage devices as the Library Public Access workstations in order to function. Security is further strengthened by only allowing specific IP addresses access to the servers via firewall rules for the server network cards. Logging of the various Boot Server services and activities are placed within a log volume. These logs are then replicated to a remote log server. Some rules are created around each type of log identifying the format of logging, subsequent triggers, and lifespan. The Boot Servers have cron scheduling for initiating a remote Wake-on-LAN broadcast to client workstations. Every morning at a scheduled time, the servers broadcast a Wake-on-LAN to specific MAC addresses. The valid receiving workstations then power up and initiate the PXE boot process. In the evening, the servers initiate a remote shutdown/poweroff command to each valid workstation.

Challenges

Pharos UniPrint

The increasing cost of supporting printing in the Library has initiated the investigation into a managed printing solution. Currently ASU employs a print management solution provided by Pharos called UniPrint. The trick for TS&D was to find a way to seamlessly implement the ability to take advantage of this system. Through much trial and error, a solution had been identified where E3 client stations can print to the UniPrint system and retrieve their output using the same steps an MS Windows system would. This was accomplished by modifying the KDE Kprint dialog application to include the ability to name each print job before it was sent to the Pharos server. The Library can continue to provide in-house paper output but now utilize the ASU managed print solution using Pharos.

Automatic Morning Startup

Control of power has been an objective for TS&D within the E3 design. The Library needs a way to power up and down workstations, preferably using an easily re-programmable schedule. This has been accomplished by using network interface cards and motherboards that support the Wake-On-LAN feature from Intel's Wired for Management specifications. By using a database of asset MAC addresses for each workstation, coupled with server schedule entries and scripts, the server can send a wakeup signal to a specific client workstation. The server also can insert a scheduled item into the remote client schedule system. This item can designate the time to execute a shutdown/poweroff command.

Linux Server File Access

The E3 design for servers uses separate strong-password and accounts stored locally on each server. To minimize security risks, a method for server file, directory, and console access needed to be identified. TS&D is using Secure Shell (SSH) as the means for remote access to servers. In addition, off-site access is further secured by using a VPN to each server.

Staffing for the New Environment

Since E3 reflects a drastic operational and environmental change, technology staff utilization and capacity had to be analyzed. It has been discovered that the old model design required a full-time tech to support the servicing, management, updates, and replacements of public computing hardware and software. The new E3 environment virtually eliminates the overhead that makes the tech necessary. All machines do not require any local image loading; this is automated. All machines do not have a third party warranty service contract. All servicing is performed in-house using inventoried parts. All functions that previously required high technical expertise in the deployment and upkeep of a traditional fat-client have been reduced to the level where a student employee can perform the service. Administrative records keeping the tech performed has been delegated to Library Administration where there is less risk of information redundancy, and less risk for error. Staffing required by E3 needs extremely knowledgeable personnel who have a clear understanding of enterprise system integration, computer programming, computing theory, current technology industry events, the Linux operating system, and project management. In addition, the staff must be flexible to fulfill a "Jack of all trades" role in systems support, yet have primary attention to the components that make up E3. The E3 design using Linux reduces the day-to-day staff involvement performing software fixes and maintenance. The result has been a reduction of one fulltime staff person and a realignment of the remaining two full-time positions.

ATA Hardware RAID 5 on Linux

ATA RAID is a challenge under Linux. Linux supports software RAID from the kernel. However, hardware RAID requires either unusual setup procedures and/or custom drivers. RAID using SCSI hard drives has been common for many years. The cost of SCSI drives is higher than ATA IDE drives. The benefit in using SCSI drives comes with the constant throughput exchange. The use of IDE drives in a RAID array is beneficial if the cost needs to remain low and the constant throughput requirements are not necessary. The Library does not need maximum throughput, nor a constant data rate for its servers. Therefore, ATA RAID is the identified solution. This solution created a challenge for TS&D. TS&D created custom installation procedures for RAID 1 as well as a custom compiled driver for RAID 5 installations. At the end of three weeks, a stable method of implementing RAID 1 and RAID 5 on all servers had been created. Upon further OS updates, it became apparent that the custom driver method for the RAID cards is not conducive for our rapid support environment. TS&D decided to replace all Promise SX6000 RAID cards with proven 3Ware Escalade RAID 5 cards. The result was a stable hardware RAID without any need to modify, nor rebuild the kernel.

Client Memory Management

Because the client workstations do not have a hard drive, all files for the workstation exists in RAM. All workstations have 1 gigabyte of DRAM storage. The Linux operating system can exist comfortably within a RAMdisk of this size. Where the challenge occurs is in the application sets loaded and their memory utilization. To enable a desktop interface over the X graphics server, we loaded KDE (K Desktop Environment) for our user interface. We enabled two virtual desktops and applied a custom theme and style for the GUI. To keep an application from using up all available RAM, a disk quota monitor checks to see if there is enough memory to execute an application. Too little memory and the application will not

load. The Web browser has been configured to use less than 5MB for session caching purposes. The temporary storage RAMdisk partition is set to a manageable size to allow for temporary storage but not affect the overall system resource needs. All unnecessary operating system files have been removed from the served RAMdisk image. When a user logs out from the workstation, a cleanup process takes place freshening available memory. The X server also had to be configured so that it will not crash in the event of 100% memory utilization. We do not expect any incidents of total memory usage. All applications run compactly using 100K to 12MB of memory. The filesystem used for the /usr tree and some others is CRAMFS (compressed RAM filesystem). This filesystem is in read-only mode and is compressed. A slight consumption of OS overhead takes place whenever a resource is requested from the filesystem due to the need to decompress in real-time. This is not a real burden on performance and the gain in reducing partition size with over a 2:1 ratio allows us to maximize the availability of our limited 1GB RAM for physical application memory. Only those files that need read/write access exist in a non-compressed partition.

Printing and Saving Documents

It became apparent early on after deploying the new workstations that users were not able to print or save various document types they were viewing. One issue had to do with the behavior of the Blackboard Portal Server sending MIME document types as an application type and not a document type. TS&D manually coded Mozilla to recognize the dot-3 extensions instead of the MIME type in order for the appropriate reader application to launch. Printing using the Microsoft document readers rendered blank pages. This was an issue in the way WINE interfaced with the LPRng print subsystem. TS&D was able to fix this problem. Printing from Adobe Acrobat presented us with the toughest problem. Adobe provides a native Linux PDF viewer. Unfortunately it is one whole version number behind their Windows and OS X offerings. Another issue with Adobe Acrobat Viewer for Linux is that it uses standard X/Motif graphics libraries for widgets and controls. When opening a file dialog box, it presents all files and directories, including hidden dot files and directories. This is a cluttered mess for the end user to navigate. In addition, the Adobe viewer does not allow a user to save a document to removable media. The application wants the user to be a super user (root). Even if a user is root it still does not work. We looked at using xPDF as a viewer but it would not represent some level 3 postscript documents correctly. We investigated the KDE KghostView application. This is a visually and functionally pleasing application. Unfortunately, while the screen representation of a PDF is correct, printing the PDF sometimes resulted in alignment problems and layout discrepancies. We opted to implement the native Adobe Acrobat viewer and set it to default save to the user's Home directory in a subdirectory named Downloads. This was an empty directory so the user did not view all of the hidden Linux files which resides in their root Home directory. The user then can open through the Desktop their Home/Downloads directory and drag/drop their PDF file onto their respective removable media or networked drive. We anticipate KghostView to improve over time whereby it would be a suitable alternative to the closed-source Adobe viewer.

Floppy Disk Removal

In order to keep the use of removable media as simple as possible, users should be able to insert and remove both floppy and zip disks at will. This is standard in a Windows environment. Under Linux, normally a disk must be "mounted" to be used and "un-mounted" before the media is physically removed. Many options were examined ranging from buttons on the desktop to press to perform the mounting/un-mounting actions to educating the user on how to properly mount/un-mount drives. None were preferable. TS&D ended up using a kernel patch called "AutoMounter" to solve this problem. AutoMounter constantly mounts and un-mounts the drives in real-time. This way a user can place a floppy disk in the floppy drive and it will automatically mount and become available. When the user saves to the drive, the cache is immediately flushed to the drive. This way the user can remove the disk with assurance that the data has completely saved to the disk before it is removed. The same behavior takes place with the zip drive. AutoMounter does not behave well with WINE. Therefore, a user must insert the removable media before launching a WINE run application in order for the drive to be available within the Windows application. Our next version of our public machine image will include the 2.6 Linux kernel and a replacement for automounter called submount. Unlike automounter which must be compiled into the kernel, submount is a kernel module. This will allow us to take advantage of other storage devices such as USB disks and Compact Flash.

III Circ Java Client

Innovative Interfaces Inc. uses a Java-based client for access to the Circulation Module. III currently only distributes the jar files and installer for MS Windows operating systems. Since the client uses Sun Microsystem's Java Runtime Environment to execute the jar files, theoretically, the Circulation Module client should run on any platform which allows JRE applications. TS&D tested this theory by setting up a Linux-based client, loading the Linux version of the JRE, and copying over the jar files from a MS Windows machine that currently had the III client already installed. The results were promising. The Circulation Module looked and behaved almost exactly the same as it's MS Windows equivalent. This test was necessary to confirm or deny the option of implementing Linux-based circulation clients in the future. III might release a Linux RPM-based installer in the near future which will ease the implementation.

Timeout Screensaver

A useful tool for public workstations is the ability to have the machine logout after a set period of inactivity. Unlike MS Windows which has a screensaver that performs this function, there is not a similar screensaver for KDE. TS&D had to write its own screensaver to perform this function. After 10 minutes of no activity registering from either the keyboard or mouse, the screen flashes into a large 10 second countdown warning the user that it will log off unless the keyboard is pressed or mouse is moved.

New Microsoft Office Formats

Microsoft has not published a free viewer application for any Office formats beyond Office 2000. Some files created in the newer Office applications do not load or display correctly in the Microsoft Office viewer applications. This causes WINE to freeze the associated application viewer and virtually renders the user's session useless. TS&D added a button to the kicker menu at the bottom of the desktop which a user can click on to terminate the stalled Windows application. This action does not harm any other currently running application and allows the user to continue with their session. In time, OpenOffice.org may add the ability to read these newer formats in a "read-only" mode (our library does not allow production software on the client workstations). At such time, we would then evaluate OpenOffice.org as a possible solution to MS file viewing.

Innovation Dissemination

Throughout the entire project, activity reports and detailed documentation (totaling 184 pages) were provided through various communication outlets. The ASU University Libraries LIST group was given the project plan as well as updates to inform and inspire their staff. ASU West IT was provided with a hardcopy binder containing all project information with frequent updates and additions sent weekly for insertion or replacement with the intent that the project might open the door to possibilities for classroom integration and cost containment. The ASU West Library Website provides a link to the complete project description and all current documentation for the general public to view. Various library and open source related listgroups were notified at the start of the project as well as updates when appropriate. The Library Technology Support and Development department created the ASU West Linux User Group to facilitate the sharing of knowledge amongst the ASU community. Article submission queries were initiated to library and academic journals. Presentation submissions were made for upcoming 2004 national library and higher education conferences detailing the E3 project. Announcements were made on various PR websites including linuxpr.com, oss4lib.org, and lisnews.com. The Library has an informational Website which has available for public download all custom applications and processes used in the E3 model and is also listed at Google. Bill Lewis, CIO of ASU was also informed on the project and its outcomes. We have received many kudos from other libraries, academic institutions, and library-related companies on the success of this migration.

Measurements

The measurements used to determine success for E3 are the following:

Survey

All public workstations allow the user to provide feedback through an electronic survey instrument. For Fall 2003, the Library had 78,699 ASU and 9,770 non-ASU accesses accounting for a survey response of

188 users. Sixty-five percent of respondents found the new system comfortable or very comfortable. Only 13% were not comfortable with the new system.

TCO [Total Cost of Ownership]

The Library used a simplified formula for measuring TCO:

TCO=Equipment+Staffing+Incidentals+Training

Equipment-- Equipment includes the cost of purchasing all necessary software, licenses, and hardware.

Staffing-- Staffing includes all hourly and salary costs associated with staffing the support of the item.

Incidentals-- Incidentals include peripherals, replacements, and supply costs for maintaining the item.

Training-- Training includes formal training, reading materials, conferences, etc. in support of the item.

The TCO savings assumes staff are dedicated solely to the item measured. This really is not the case and therefore realized savings is actually higher. It is also assumed that Incidentals will not be equal between the two TCO design models. Since TS&D has control over the quality of parts and workmanship in the second model, the second model may also have an additional lowering of TCO. TS&D was able to eliminate 1FTE Windows hardware/software support personnel.

	Year-1 non-depreciated
Old Design	
FT Staff	\$100,000.00
Student Staff	\$23,250.00
Workstations	\$145,600.00
Servers	\$36,000.00
Software	\$8,790.00
Peripherals/Supplies	\$1,500.00
Support Material	\$1,200.00
	\$316,340.00
TCO	\$3,071.26
New Design	
FT Staff	\$109,000.00
Student Staff	\$23,250.00
Workstations	\$51,140.00
Servers	\$31,717.00
Software	\$850.00
Peripherals/Supplies	\$1,500.00
Support Material	\$1,200.00
	\$218,657.00
TCO	\$2,082.45
TCO OLD - per unit	\$3,071.26
TCO NEW - per unit	\$2,082.45
Savings per unit	\$988.81
Total Savings	\$97,683.00

Strategic Plan 2004-2008

All aspects of the E3 architecture and project had to be in congruence with the Library Strategic Plan.

Budget

The E3 model breaks the dependency on a 3-year replacement cycle. The Library accommodate virtually unlimited hardware types and configurations utilizing the inherent features of Linux and the boot image services. Since TS&D supports hardware replacements in-house, the Library is not required to pay for warranty/service agreements. In addition, the hardware can be replaced as needed in small or large quantities over short or extended periods of time using the funding available and need requirements. E3 has allowed for a 50% reduction in capital equipment spending to meet future available funding.

Trouble Tickets

E3 has resulted in a drastic reduction in reported systems and software problems. Previous fiscal years averaged 372 reported problems. The new environment has generate 31 trouble tickets during its first year- a 91% reduction.

Conclusion

Established in 1984, ASU West is a young institution. The Library was founded on a truly innovative model that provided access to materials for our students and faculty instead of amassing onsite collections, which were the traditional, and expensive, model in libraries at that time. Over the years, we have served our students and faculty by initiating services and programs that are recognized as unique to libraries, and adding value to teaching and learning.

We have consistently worked to create and maintain a technology environment that supports our capacity to experiment and grow in ways that enhance our services and meet emerging user needs. By 1998, however, limited by Microsoft architecture and systems, we had lost the leading edge that had become our trademark.

The implementation of E3 has propelled us into the forefront of library technology applications. The most important result is the capacity to be responsive to the technology and research needs of our users with the resulting impact on their studies and research. Equally important, we will be more effective stewards of our equipment and time, implementing new technology that will not adversely affect our budget.