

The Analysis and Maintenance of a Robust, Highly-Available Computer Systems Laboratory

Overview:

A well designed computer systems environment requires the simultaneous solution of a large variety of conflicting problems. One of the best examples is that of the balance between security and functionality. If users were allowed to run an unlimited amount of processes and use all available RAM, then those users that need the maximum available resources would be more efficient, however, those who abuse their privileges would cause extreme detriment to the general service availability.

This project is an exploration of one possible environment that meets the criteria for a "robust" and "highly-available" laboratory, while still providing the students who work in the lab with all of the required facilities. The first goal was to determine exactly what those criteria are, and exactly what "required facilities" entails. This is perhaps the most difficult portion of the project, due to the complex issues involved, and it must be executed for each lab that is created in order to best fit the systems design to the needs of the students and staff.



Requirements:

The requirements determined for the TJHSST Computer Systems Laboratory are as follows, in order of decreasing priority:

- The lab should provide sufficient hardware for each student to have his or her own computer during his or her classes in the lab, as well as providing a certain number of spares should any particular computer be in need of repair.
- The lab should provide sufficient computational resources on each machine, as well as on the network as a whole, for students in classes such as *Computational Physics*, and *Super-Computer Applications* to learn how to write software in a true high-speed networked multi-processor environment.
- The lab should provide the software such as compilers and interpreters for as many differing languages as is practical, with the express purpose of assisting students in *Artificial Intelligence*, and *Comparative Computer Languages*, as well as providing necessary services for seniors working on their technology projects.
- The lab should enable lab administrators to easily and efficiently manage users, software, servers, roles, and activities without undue and error-prone repetition. The administrators should also be able to delegate such privileges as they see fit to responsible assistants, and the systems should be simple enough for each rising class of admins to understand and operate.
- The lab should provide a simple, secure, networked environment such that each student can use their own personal environment independent of what system they are connected to. This also includes providing access to other information-technology services within the school.
- The lab should minimize its equipments and software costs, both initial purchases and maintenance over time.

Software:

The TJHSST Computer Systems Laboratory's requirements for robustness and stability, as well as the requirement for maximum software support and the desire to minimize costs where possible, are all direct factors in the decision to utilize Linux. Linux is a completely free and Open-Source operating system that has more software support than any other POSIX-compliant system available. It is extremely robust and requires little day-to-day maintenance other than software updates. The kernel and applications are constantly being maintained and improved on a day-to-day basis.

Of the available Linux distributions, the Computer Systems Lab uses Debian <<http://www.debian.org/>>, primarily due to the ease and flexibility of configuration, as well as the constant software updates available. The Debian distribution is one of the largest available distributions which, in addition to being non-profit, is dedicated to ensuring that all software distributed is free of legal complications.

Workstations:

One of the primary requirements for the lab is that each student has his or her own computer with sufficient resources to run demanding graphics and computationally intensive tasks, and yet it is on this requirement that the lab most frequently falls short. The lab possesses a collection of 19 2.8GHz Pentium 4 workstations with 256MB RAM each and 11 2.4GHz Pentium 4 workstations with 512MB RAM each. The workstations all have aging GeForce2 graphics cards which are barely sufficient for the 3d work that many students wish to undertake, and the small quantities of RAM are insufficient for any but the smallest data-sets for the *Computational Physics* classes. As 30 workstations is insufficient to support more than one class per period, there are additionally 12 ancient workstations with 1.8GHz Athlon CPUs or dual 800MHz Celeron CPUs. The primary cause of the resource shortages is the lack of a hardware replacement and renewal line-item in the school budget.

Computer Systems Lab Workstation Administration	
Debian Software	Custom Software
<ul style="list-style-type: none"> • Perl: A powerful all-purpose scripting language • SystemImager: Simplistic console-based ghosting tools that use an rsync backend. • SNMP: An advanced remote system-monitoring tool. 	<ul style="list-style-type: none"> • sysconf/syspref: Workstation autoconfiguration tools. • sysimage: SystemImager-based software that provide beta-testing and remote ghosting capabilities. • tj-kpkg: Kernel source tree management tools.

Administration:

The Computer Systems Lab workstations are a collection of systems of widely varying construction, especially the older ones that have been repaired time and time again. Due to the heterogeneous composition and the multitude of requirements, a secure, efficient method of simultaneously managing more than 50 independent systems was needed. To this end, we used Perl, a powerful scripting language, to write a set of scripts totaling over 4000 lines of code that run on both clients and servers to automate the process of managing the clients. The software includes "sysconf", a tool which automatically detects the hardware in a workstation and reconfigures all the software to run optimally, "sysimage", a complex wrapper around SystemImager that allows us to remotely ghost clients, optionally with a beta version of the workstation image, and "tj-kpkg", which quickly and efficiently recompiles new kernel binaries and associated modules for multiple architectures simultaneously.

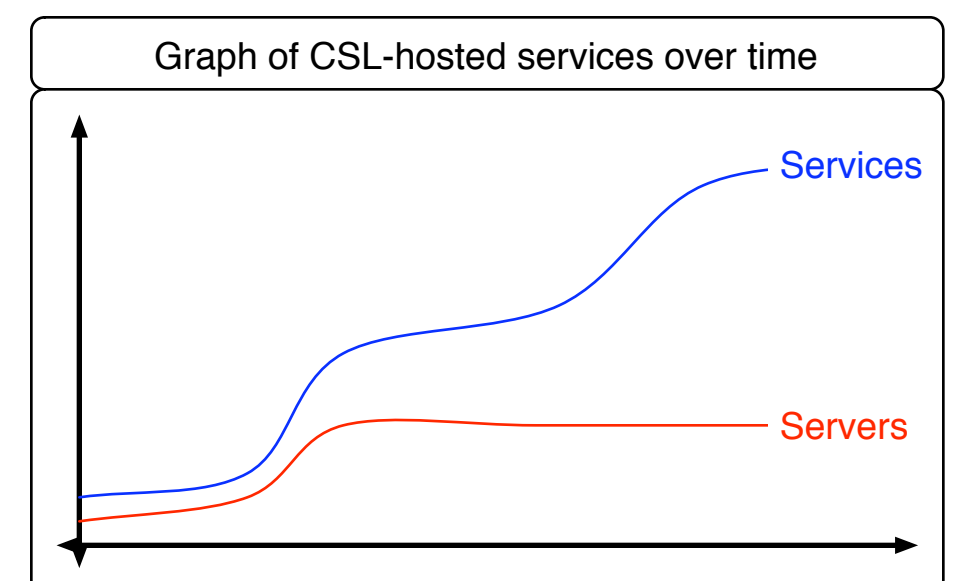
New Computer Systems Lab Service Layout	
king	emperor
Authentication:	Kerberos
User/System Data:	OpenLDAP
File Services:	OpenAFS on DRBD
Web Authentication:	WebKDC/WebAuth
robustus	macaroni
Domain Name Resolution:	BIND 9
Client Network Autoconf:	DHCP 3
Student Intranet:	Apache 2 & PHP 5
chinstrap	cronos
Mail Transfer (SMTP):	Postfix
Mail Access (IMAP/POP):	Courier
Mail Storage:	Ext3 on DRBD
adelie	rockhopper
Student websites:	Apache 2 & PHP 4
Student remote access:	Secure Shell

Services:

The Computer Systems Laboratory does not rely on other portions of the school for its services, instead it houses more than ten servers, two clusters, and a Cray super-computer. The converse, however, is not true. Due to the available server-space and CPU-time, as well as the high stability of the servers, the CSL is frequently used to host services for other portions of the school's IT network.

The best examples of this are the school's web server and the student Intranet, both of which operate out of the Computer Systems Lab. They designed and operated by the student administrators at a high level of efficiency. Recently, however, the student Intranet has become too old and outmoded to properly serve the school, and so a new system, called *Iodine*, is being designed to complement a major services upgrade within the Computer Systems Lab.

One of the largest problems in the Computer Systems Lab is the amalgamation of large quantities of services on the same few pieces of hardware. This creates an increased probability that a single system failure will cripple a huge proportion of the lab, or even of the school itself.



Conclusion:

While the Computer Systems Laboratory is well maintained by its staff of student administrators, many of whom work more hours than some part-time jobs, it receives insufficient funding to properly handle all of the challenges that it must meet. Due to lack of hardware, the ratio of services to servers is climbing at an unsustainable rate, despite the fact that many of those services are being redesigned to better accommodate the additional limitations.

The upcoming renovation of the CSL's authentication and authorization services, in combination with the new student Intranet (*Iodine*), is a much needed upgrade to core school services, and it should be accompanied by a corresponding upgrade to the servers that will host it.