# Overview and Challenges

The original intention of the project was to use Usermode Linux to provide each student with a personalized, remotely-accessible Linux server. Shortly into the project, however, it became quite evident that UML was not yet up to this task.

## *Shortcomings of UML*

While UML appeared promising in it's approach (each subsystem ran as a single process under the main host's kernel), in practice it lacked features critical to this project: It was not stable under Ubuntu 6.06 (the initial OS installed on the server), there was no easy way to limit how many instances were running. There was no way to 'lock' a given amount of memory to a process. That is, while the server had 8 Gigabytes of RAM, and I wanted to allocate 128 MB (or 0.128 GB) of memory to each instance, the system would allow hundreds of virtual machines to start. This resulted in severe memory starvation, and abysmal performance.

The virtual machines would randomly halt under student use, frustrating students and tying up resources needlessly.

## *Xen*

Further research revealed another possible approach: Xen. Xen is a paravirtualization technology similar to VMWare. Unlike VMWare, however, it is free and open source. Switching to Xen, however, was no small task.

Using Xen would require a reformat and reinstall of the operating system on the server. While Xen is nominally supported under the installed version of Ubuntu, the time necessary to get it running was prohibitive. Fedora Core 6, however, shipped with heavy Xen integration. In fact, getting the initial test systems up and running under Fedora Core 6 took only about two hours.

Once the dual-Opteron server had been reformatted and Fedora Core 6 had been installed, the customization began.

## *Customization*

Out of the box, FC6's integration is really designed for a single user running a couple of small images. There are a variety of convenient graphical tools for creating new virtual machines (hereafter called 'images'), managing those machines, and starting and stopping them. They were, however, all inherently manual—that is, they required human intervention at many, many points. This would not be appropriate for Hennepin Technical College, as the server would be expected to be available twenty-four hours a day, seven days a week, for at least the length of each semester's classes.

Indeed, the primary push for this project was to make Linux systems more available to students.

Furthermore, the machine needed to be secured, to prevent break-ins by persons inside and outside of the school.

The first hurdle to overcome was the initial configuration of student images. The images needed to be substantially similar, with an administrative account accessible to the student, a separate administrative account accessible only by the instructor, and a full load of Linux server software. Furthermore, each machine needed a unique IP address (which would always be the same).

Configuration of the initial image was relatively simple: it was created via Fedora's graphical tool, and final configuration was completely by logging into the running image. Appropriate software was installed and configured manually. Network settings were adjusted, so that the machine could more easily grab it's hostname and IP address via DHCP. The machine was then shut down, and copied into a master file. Indeed, the primary advantage of the xen approach is that the entire system image is stored in a single, 8 GB file. To re-image a machine, that file is copied onto a target partition. Generally, that takes three to four minutes. Each student is assigned their own logical partition. For this configuration, LVM was used to create a Volume Group which spanned three 500 GB SATA drives.

Virtual machine definitions were created which limited each image to 256MB of RAM (it did not prove possible to run Fedora Core 6 stably in less memory), assigned a MAC address, and set the image to boot to a character-only interface. X-Windows was simply not wanted for these machines. Also, the machine definitions control which partition is assigned to which machine.

Since each image consumes 256MB of RAM, and the master server itself has only 8 GB of RAM,

all student machines could not be running simultaneously. In fact, by spring semester, some 85 student machines were configured. Were they all to be running at once, over 21 GB of RAM would be required. Students could not be expected to shut down their machines when finished with them – while most undoubtedly would do so, circumstances might prevent their doing so. Perhaps their internet connection would drop, and they would lose their connection to their image. In this case, the image would have to be automatically shut down (or allow reconnection). It was conceivable that a student would try to keep their image running continuously, to do processing or run some other job. In this case, some safeguards needed to be in place.

A script was written which compared the list of running machines with the list of currently logged-in users. If a machine was found to be running when a user wasn't logged in, it was shut down. This script ran every ten minutes.

As part of the initial image creation, the student's file system is imaged, their username on the Xen server is created, their password is set to an initial value, and entries are created for their system in both the DNS server and the DHCP server. This allows their machine to be addressed by an alphanumeric name as well as an IP address.

User accounts were set to execute a shell script on login. This script gave them the choice of changing their password or connecting to their machines. The connection portion of their script determines if their machine is running yet. If it is, it connects to the machine using the SSH protocol (version 2). If the machine is not running, it boots their image, directing those boot messages to their session. The student sees exactly what they would see if they had physically powered on a Linux box while sitting at the console.

## Rollout

By the beginning of spring break, the system was refined enough that it was cut loose to Linux Admin 3 students and Linux Admin 2 students. They were given the chance to earn extra credit points by compromising security on the virtual machines or the master server. They performed admirably, finding several subtle ways of compromising both performance and security on the master server. Since one of

the clauses of the assignment was that they had to tell the instructor HOW they'd broken in, the holes were quickly closed.

The system was then turned over the Linux Admin 1 students. They found no remaining problems in the system, and were able (with sufficient coaching) to log into the system and use it.

## Final analysis

At the end of the semester, logfiles were analyzed to determine how many students were using their virtual servers. All of the Linux 3 students had used the machines on and off. About forty percent of Linux 2 students (a total of 11 students) used their servers regularly. Some used them for testing out configurations which would later be used in class. Two of them used the virtual machine as a sandbox – trying things they dared not do on their normal machine. One student decided to use it to crunch numbers, twenty-four hours a day.

Among the Linux 1 students, nine of them used the system regularly, practicing creation of user accounts, and generally familiarizing themselves with the environment. Out of eighty-six total Linux students, 43 students used the system semi-regularly. All students had logged onto the system at least twice. In the original write-up, a secondary measure of success had been defined as 30-60% utilization. A full fifty-percent of all linux students used the server to a reasonable degree.

The primary measure of success had been the ability to host a server for each student. This has been completely satisfied. Not only does the system have sufficient diskspace to host upwards of 150 images, up to twenty-five of them can be active simultaneously – allowing every student in a class to work concurrently.

## Future

Beginning with fall semester 2007, the server will be fully integrated into all Linux courses. This will be used not to accelerate the pace of the learning, but rather to offer expanded extra-credit and to give students much more practice time.

In Spring of 2008, the server is tentatively scheduled to be opened up to Network Essentials and database design students.

The scripts which have been developed are currently licensed under the GPL (Gnu Public License), and have been distributed to a handful of individuals. Attached to this writeup is a binary file commonly called a 'tarball'. It contains a copy of all scripts developed for this project, along with a sample virtual machine configuration file.

This writeup will be used as the basis for the whitepaper proposed in the initial Award application. Expected completion time for the whitepaper is Christmas 2007.