


Implementing an Open Source Conferencing System for Distance Education

Philip J. Lunsford II

As distance education becomes more popular, the need in educational institutions for inexpensive long-distance voice and teleconferencing services grows. Students need to communicate with each other and with instructors. Additionally, the availability of broadband connections to the Internet for students is increasing and the cost is decreasing. This article reviews a case study of implementing an open source system of an Internet protocol-based private branch exchange (PBX) along with software-based phones that allow students and instructors to communicate via voice or both voice and video. Since the system is based on open-source software, all of the software is free and the hardware requirements are minimal. The Asterisk-based phone system TrixBos was implemented on a VMWare virtual

machine as the PBX and the software-based phones X-Lite, Ekiga, and Idefisk were used by the instructors and students. A review of open-source licensing, along with the rationale for choosing the particular software distributions, is given. Additionally the topology of the implementation is reviewed along with the configuration challenges found. Finally, a review of security and policy considerations is given.



INTRODUCTION

Organizations all over the world are expanding their Internet presence. Whether educational institutions seeking to broaden their reach and increase revenue by offering online courses, companies trying to offer telecommuting capability, or organizations seeking

to provide online access to the handicapped, online collaboration and conferencing is becoming popular across all industry sectors.

Community colleges, notoriously on shoestring budgets, have begun offering online classes for general education and preparatory courses. In order to complete the online coursework, a collaboration tool is needed. E-mail is not real time and does not offer the consistent interactive "feel" online students need, but can be provided with such tools as conference calls, video and text chat. Schools constantly deal with a segment of their students that does not perform well. These students fall behind and their grades seldom recover. Rural school districts could use distance learning tools to help mediocre students shore up course work after school without the students needing to be on campus. Another educational segment, the growing homeschooling community, continues to struggle with the question of socialization. Pundits such as Mike Jerry, expert author at Ezine, teacher organizations, and civic groups and even the popular TV psychologist Dr. Phil, feel that homeschooled children's isolation causes disconnect later in life (Delp, 2006; Jerry, n.d.). The usual and optimal solution is to schedule play times and field trips with other home school groups, but at home there is still a disconnect. Homeschool groups need a solution for families and children to keep in step with each other and provide needed socialization. The handicapped community has always struggled with accessibility to work, school and other social settings. Providing this equal access is often costly, time consuming, and complicated. These organizations require a cost effective and quick solution in a setting where the handicapped need access to various people on a consistent basis and always feel connected to their office or classroom. Finally, universities all over the world are expanding their online course offerings and yet the available tools lag behind. According to an

MSNBC report, 3.2 million students took online courses in the fall of 2005. That accounts for one in six higher education students (Pope, 2006). These institutions seldom initially consider the inherent challenges and the tools these students and instructors need. Consequently, the credibility of these universities can be on the line. Although the solution can be complex given the exact needs, the vastly expanding Internet is a mature infrastructure. Solutions that utilize the Web for voice and video calls are becoming more popular. Since the Internet itself does not care if information is data, voice, or video, it does not enforce a limit on the number calls, unlike circuit switched technology. The only limitation is either a service provider enforcement of total bandwidth usage or the limitation of a private network's technology. Although some rural areas do not have broadband, the outlook on its expansion is good since, as of 2003, it was available to 80% to 90% of Americans, according to Mehlman (2003). Mehlman also cites the Jupiter Research estimate that 56% of households will be using broadband by 2007.

Most universities and colleges have broadband Internet access, so the challenge is at the private home where the monthly price of broadband continues to thwart low-income earners. Although the problem is clearly at the residential end because of location or cost, the good news is the advent of not only expanding DSL and cable networks, but also technology such as broadband data modems and WiMax, which are putting downward pressure on monthly service costs. Data modems simply use the digital 3G technology of cell phones but offer broadband services usually by the use of either a PCMCIA card for laptops or PCI card for regular PCs. This is an attractive and comparable (\$60 per month) solution in many rural areas that have cellular coverage but no DSL or cable. WiMax is the new standard IEEE-801.16 that extends broadband

services via radio signal to businesses and residences up to 30 miles in non-line of sight scenarios. This comparable service (\$50 to \$100 per month) is now in residential trials and is expected to instantly expand deployed broadband services to rural areas by 30 miles (Kay, 2003). Of course, satellite broadband access is available to businesses and residences as part of satellite TV service at comparable prices.

Because of budgetary constraints and the expense of paid conferencing solutions, open-source products are an excellent option. A vital aspect of open-source is the extensive network of online forums and communities made up of software developers, integrators, and network administrators. Note that not all open source products have extensive communities, but many do. The open source software removes the cost barrier to market entry so that educational organizations can avoid having to put a “down payment” on a solution without a clear outcome.

Many equate free software to open source. Although an open source package may be free, all free software is not open source. Such is the case of VMWare Server which is a free software package that emulates individual computers within a physical computer. This software is free but is not open source; therefore, users are not able to modify and repackage the code, creating a new emulator. Open source means users have access to the software source code, and it is usually issued under a license called the GNU General Public License (GPL) (Su, 2006). With the GPL license, anyone can use or modify the program. The caveat is that if any software is sold that is based on the original software, then the newly-derived software must be released under the GPL license, thus making the modified source code available to anyone free of charge. What normally does not come with free open source software is support in the form of a helpdesk or auto-

matic updates. However, there are many online forums in the open source community where users can ask questions and often the documentation for application is extensive and available online. The business model for companies dealing with open source software is based on selling services such as integration, training, support, and updates.

H.323 AND SIP

Internet communication has eclipsed circuit-based systems in which channels, usually of 64kbps, were assigned to subscribers. ISDN combined two of those channels for a total of 128kbps, not enough bandwidth for an IP Web-based video conference. When choosing a multimedia solution, we had to decide if we preferred or required a specific protocol. The two choices for protocols that coordinate the creation of a phone or videoconference call are the older H.323 and the newer SIP. The H.323 protocol was originally designed for video conference setup using switched circuit connections. It has been updated to include packet switched IP networks. SIP is simpler and truly a multimedia protocol because it is agnostic of the underlying transport protocol (Stegh, 2006). So, a SIP server could set up a text chat or whiteboard call session between two SIP clients without the server caring about the protocol. H.323 must support and negotiate the protocol before clients are allowed to collaborate. Although the industry has continued deployments of H.323 systems, service providers often deploy SIP capability in tandem for eventual cut-over. This is because SIP is gaining ground due to its relative simplicity and by becoming the 3GPP approved standard (Stegh, 2006). Although there are H.323 open source tools, the server and client side of most new collaboration software uses SIP as its session protocol.

CASE STUDY

After several weeks of study which consisted of learning about Linux, the latest trends in voice over IP, operating system compatibility, and the available open source products, we selected Trixbox 1.2 as the server for the East Carolina University lab trials. We considered other software, such as Webhuddle, which runs on a Web server, and Netmeeting, with the Linux server acting as ILS (Internet Locator Server). Ekiga was not considered as a server solution, as it does not offer call control and registration. It is certainly usable as a client in our solution and has been tested as such. One caveat is the Windows open source distribution of Ekiga that we tested was unstable and crashed the client computer in our tests. Webhuddle, on the other hand, is very stable and offers all the control, whiteboarding, chat, and desktop control. It also does not require a software download, but the client side must download and start Java to run the Webhuddle applet. Although it is a relatively "light" applet, it does require an initialization time, up to 1 minute, depending on the client computer's processor. Webhuddle has less community forum support and the installation was complex and confusing. We chose to install it on Windows running within VMWare. This worked, but the lack of call control and administrative capability caused us to abandon WebHuddle. In trying to find a peer-to-peer solution that would not require a server, we configured Netmeeting on a Windows machine side and Ekiga on a Linux client and realized the two were not compatible. We were left with Trixbox 1.2 and Asterisk system available at www.trixbox.org. This installation went smoothly. With complete real-time control of the audio and video codecs, conference room monitoring, and the capability to define hundreds of user extensions, we selected Trixbox as our final solution. Trixbox is an audio and video conferencing system based on Asterisk, which is a free Public Branch Exchange (PBX). Trixbox is

produced by Digium, which originally designed Asterisk. Digium has made available a full Linux distribution of Trixbox as an ISO image at www.trixbox.org. An ISO image allows an installation CD to be created that includes a complete installation. Trixbox is composed of several software packages that have been integrated, including a webmin interface for web browser based administration. There are several other components which make up this system. Asterisk is at the system's core and provides the capability to handle SIP VoIP calls. In its basic form, Asterisk can be fully customized using the configuration files in the directory `/etc/asterisk`. This system will handle inbound and outbound calls and know how to route them, which codec to use, and if video is supported. It includes hundreds of customization options. Included in Trixbox is SugarCRM which is an integrated contact and calendaring tool that users can log into, define contacts, schedule meetings and call into a meeting directly from a browser. Also included is the A2Billing management interface used for optional billing management. Most educational institutions do not need this feature. The Flash Operator Panel is a real-time call operator control panel which displays in the administrator's browser each configured extension and if it is presently on the phone. Web Meet Me Conference is the administrator's global view of the conferencing status for the entire system. It allows for a real-time status of all conference rooms and parties. FreePBX is a browser administration front end for Asterisk. It provides the administrator the ability to define and configure the system from a web browser. The CDR reporting tool is another administration tool for call tracking and monitoring. FreePBX also includes Config Edit designed to provide browser editing of Asterisk flat configuration files. We have found this version of Config Edit causes problems and has actually deleted all contents of the vital `sip.conf` file, requiring a

complete system reinstall. Thanks to VMware which saves snapshots of the virtual machine, a system restore was trivial and it was back up and operational within minutes. Finally, this distribution comes with the CentOS operating system which is a very stable early Enterprise linux distribution. This version does not require an operating license and does not come with a support license.

The ECU pilot implementation found that available free SIP softphone clients are not yet feature rich, but some basic audio and video SIP phones are quite mature. Those with many features are not yet stable, such as Ekiga for windows which supports file transfer and whiteboard features. In our tests, the Ekiga client crashed the Windows XP host laptop several times. We ended up using the free softphones X-lite and Idefisk (www.soft20.com and www.tucows.com).

Our final solution included the use of the virtual machine emulator package VMware; thus, we were limited to the list of supported Linux distributions for VMware Server. We settled on Ubuntu 6.06 because of the extensive free support. The individual in charge of this prototype installation did not have extensive Linux administrator experience. Firstly, we needed to understand the ISO installation process and the administrator's duties such as enabling root access and changing the root password. VMware's Ubuntu installation is not trivial, but once installed it is quite stable and easily operated, especially from the Ubuntu desktop. The VMware virtual machine is easily created and defaults to a 3 GB virtual hard disk. This virtual appliance shares the host system's RAM, so the amount of physical memory on the host machine is critical. The ECU pilot server runs on 2 GB of

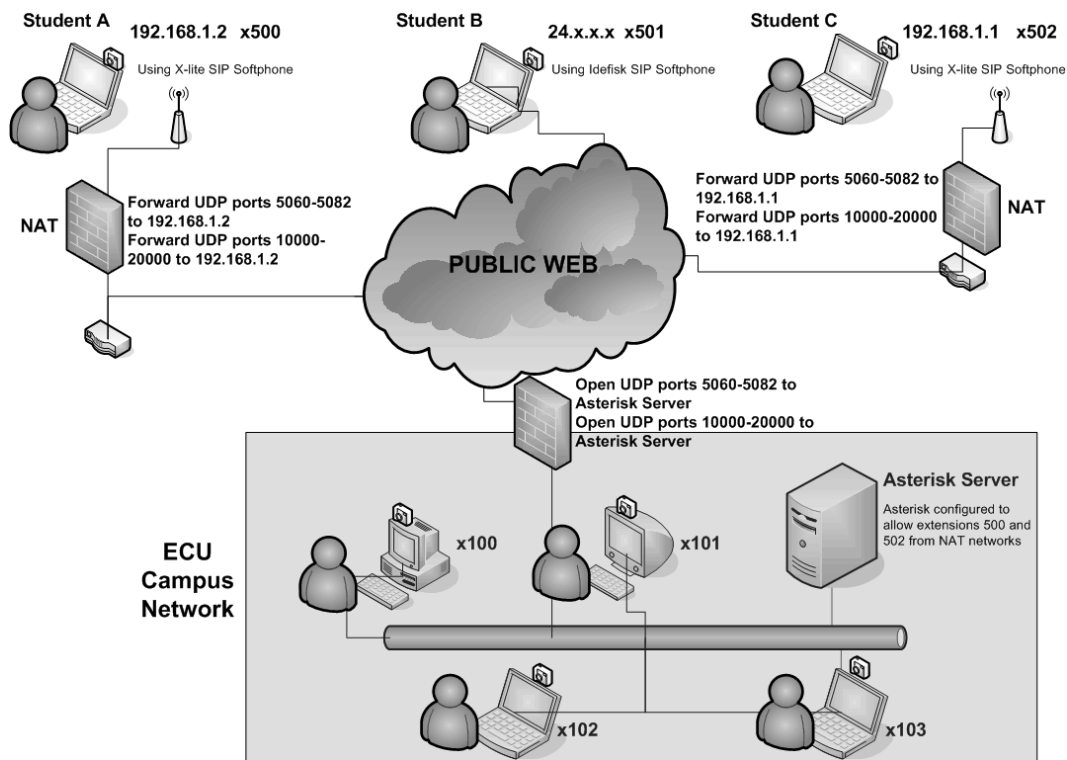


Figure 1. Topology of the Tribox system implementation.

RAM, of which it has dedicated 128MB to the Trixbox virtual machine. It should be noted that VMware provides a solution that is easier to manage and configure, but is not required. The Trixbox server could be set up on a machine running native Ubuntu and would require less RAM.

Upon installation of the Trixbox system, the dynamic DHCP IP address gave us full connectivity to the network. Upon bootup we had a console printout which showed us the http path to use for Web access. After exploring this system via the browser, we got a feel for the administrator's access and changed the system login password. Unfortunately, the connectivity to the Trixbox machine was lost without warning. After much investigation, we saw the CentOS distribution of the `ifcfg-eth0` comes with DHCP turned on, but something in the machine makes it renegotiate the IP address periodically. Therefore, if we changed the IP address to static or even rebooted the machine, it renegotiated a DHCP address and all connectivity was lost, although data is not affected. Though the cause of the forced DHCP request was never diagnosed, the workaround was to change the TrixBox `/etc/sysconfig/network-scripts/ifcfg-eth0` to `PROTO=none` and define a static IP address.

At the core of this system is Asterisk, which has two main concepts: extensions, and routes. These concepts make up the dial plan. The dial plan idea is that every call in or out of the system must be handled and routed to a client device, whether a SIP desk phone, SIP softphone, pager, cell phone, or public switched telephone network (PSTN) (i.e., legacy phone). In our installation, we did not integrate our Asterisk installation with the PSTN network. An extension is comprised of a number (e.g. 100) and a secret (e.g. 12345) for authentication and registration. There are inbound and outbound routes. There can be many of these and many combinations of them based on the different types of client devices and routing. In our simple case, we

have only SIP and IAX2 as possible clients. No PSTN connectivity is required, since the ECU implementation includes the need for audio and video. If desired, a PSTN to SIP VoIP call architecture is achievable by simply purchasing a low cost monthly service from a VoIP gateway service provider. This provider assigns a traditional 10-digit telephone number to the Asterisk server's IP address and does the translation upon receiving a PSTN call. The central Asterisk IVR attendant module, unless otherwise configured, will answer the call translated from PSTN to VoIP from the gateway and the PSTN client will hear a number of options and the chance to connect to an extension via the key pad.

Our implementation's Trixbox server sits on the same LAN as its physical host Ubuntu machine. It resides behind a firewall but not behind a NAT device (more about understanding NAT can be found at <http://www.pcsupportadvisor.com>). With the advent of home routers and wireless access points, the problem of NAT comes in at the client side. When a SIP softphone tries to register with the server, the signaling packets are from the private side of the network, most likely from a 192.168.1.X address. The router replaces the source address by its own external address and the Asterisk server sees the external address of the client as the SIP client, which is incorrect. Asterisk must be told to ignore this address, register the client and continue the SIP call setup until the call is active.

Because firewalls and routers close TCP and UDP ports as part of their design in order to work effectively, we had to reconfigure the firewalls in front of the Trixbox server to open UDP port ranges 5060 to 5082 and 10000 to 20000. The lower range is used for signaling and the higher range for voice call content, otherwise called the bearer traffic. The alternative to opening a range of ports at each end of the connection is to use a Virtual Private Network

(more about VPN can be found at <http://compnetworking.about.com/od/vpn>). One advantage of a VPN is that of secure communication. The complication is that each user must install and configure the client, which can become complex and difficult to support. Portability is also an issue, since the VPN package and SIP phone must be installed on every end device. So, if a user moves to another computer, the setup and configuration must be duplicated for both the VPN and the phone. As part of our solution, we have downloaded the free SIP phone Idefisk. This perfect solution for non-computer savvy students is an executable softphone copied to a USB drive along with its customized configuration file for that particular user. The user only needs to plug in the USB flash drive, click on the Idefisk icon and the client registers automatically with the Asterisk, authenticating with the preconfigured extension.

In the Asterisk deployment, the solution designer must budget for the likely used bandwidth versus the available bandwidth. The ECU network has abundant bandwidth, so we have chosen the high-quality G.711 audio codec, which reproduces noncompressed voice content, therefore excellent quality. After installation and the short ramp-up to understanding the system's architecture and basic concepts, we were able to demonstrate several scenarios. First, we remotely configured two extensions via the Web server's Web tool, one for a student and one for an instructor. Then, we defined a conference room to start at 6:30 PM, a student password, and an administrator password. The student called into the conference room defined extension (900) using the X-lite softphone from his home network behind a NAT router, and was prompted for a password. The system recognized the student login and played a .wav file as background music while waiting for the instructor. The instructor called into the same extension from the ECU network on X-lite, but logged in with the call

administrator password, instantly recognized as such. The .wav file stopped playing and instantly both parties were allowed to talk. In the second scenario, one student using X-lite on a data modem from an Internet café directly called the other student's extension. This student invoked the Idefisk softphone from a wireless home network with NAT. By plugging the preloaded USB key with the Idefisk executable and configuration files into a simple laptop never before used for conferencing, the student was able to just click on the .exe, start the program and make and receive calls after a quick moment of authentication with the Trixbox server at the ECU campus.

FINAL CONSIDERATIONS AND CONCLUSION

As in any new system rollout, due diligence must be made to protect areas which only administrators should access. Trixbox has already done this by password-protecting the administration area. The Asterisk CentOS server itself should also be password-protected in the traditional Linux sense. A user policy should also be disseminated so the system is not abused. Our installation accomplished this through locally developed tutorial Web pages which also covered all aspects of SIP client configuration and troubleshooting from the students' perspective. Each instructor can be given a range of extensions and passwords that can then be assigned to students. Those extensions can be recycled at the end of the course. This policy also prevents any personal abuse of the system once a course is finished. The instructor community should be trained first on the limited administration and day-to-day best practices. The Asterisk server and its network should be protected from known and potential exploits. We only found one published Asterisk vulnerability, where the server SCCP "Skinny" port 2000/tcp allows an attacker to execute code on that port as

root without authentication. Our fix was to disable the Skinny module initialization at Asterisk boot time. The CentOS firewall should not be disabled as it does not hinder server operation. Because of the firewall and NAT challenges and the Linux knowledge ramp-up, anyone considering the implementation and rollout of TrixBox should certainly have at least some basic Linux administrator and network security skills.

Although a few careful post-installation modifications are required, Trixbox 1.2 proves to be a free, quick, efficient, and customizable conferencing solution. However, the lack of stable and feature rich SIP video softphones somewhat dilutes the power of this solution. We are confident more free and open source clients will become available.

REFERENCES

Delp, V. (2006). *Dr. Phil doesn't get homeschooling*. Retrieved December 12, 2006, from [http://](http://homeschooling.families.com/blog/dr-phil-doesnt-get-homeschooling)

homeschooling.families.com/blog/dr-phil-doesnt-get-homeschooling

Jerry, M. (n.d). *Homeschooling socialization*. Retrieved December 12, 2006, from <http://ezinearticles.com/?Homeschooling-Socialization:-The-Greatest-Conflict-of-Homeschooling&id=342922>

Kay, R. (2003). *Quick study: WiMax*. Retrieved December 6, 2006, from <http://www.computerworld.com/mobiletopics/mobile/story/0,10801,87555,00.html>

Mehlman, B. P. (2003, April). *The broadband revolution*. Speech presented at NECA-NARUC Conference, Washington, DC. Retrieved December 4, 2006, from http://www.technology.gov/Speeches/BPM_030428_broadband-rev.htm

Pope, J. (2006). *Sharp increase in online-course enrollment*. Retrieved December 4, 2006, from <http://www.msnbc.msn.com/id/15625558/>

Stegh, C. (2006, July). *How is SIP winning favor over H.323?*. Retrieved December 8, 2006, from http://searchvoip.techtarget.com/originalContent/0,289142,sid66_gci1198457,00.html

Su, J. (2006). *GNU general public license*. Retrieved November 28, 2006, from <http://www.gnu.org/copyleft/gpl.html>

→ **Trixbox**

→ **Webhuddle**

→ **Netmeeting**