Design of a distance education, 2+2 articulated, IT curriculum in Linux System Administration

Alessio Gaspar, William Armitage, Naomi Boyer
University of South Florida (USF),
3334 Winter Lake Rd, 33803, Lakeland, Florida (USA)
softice@lakeland.usf.edu

ABSTRACT
This paper discusses the design of an articulated program between Polk Community College (PCC) and the University of South Florida Lakeland (USFL) Department of Information Technology (IT). This blended program (distance-education and face-to-face) is meant to train "advanced Linux technicians" to both meet the workforce demands and have the potential to engage in graduate academic pursuits. We discuss the curricular and technical challenges of this endeavor with a focus on how this work leverages previous results obtained with a NSF-sponsored grant.

1. INTRODUCTION

The nation-wide need for qualified Linux technicians, developers and system administrators is a driving force for projects enabling Associate in Science (AS) students to further their education by obtaining a Bachelor in Science. To this end, the University of South Florida Lakeland (USFL) Department of Information Technology has recently built an AS to BSAS (Bachelor of Science in Applied Sciences) concentration in IT aimed at AS degree holders as an alternative to the BSIT, which is most appropriate for students who have earned an AA (Associate in Arts) degree. This paper discusses recent efforts to further complete this BSAS offering by focusing on a specific application domain and technology; Linux system administration. This endeavor encompasses objectives from both the Linux-focused technical courses taught at Polk Community College (PCC) and the advanced topics developed in the ACM's national IT model curriculum. Our mission is to provide, through USF's AS to BSAS transition, a bridge to further develop Linux technicians' training and expand their professional and academic horizons. As in other computing disciplines, the educational challenge is to not only provide them with a comprehensive technical training on current technologies but also to make it so that they can successfully adapt to an always unpredictable technological future.

Our objective is to integrate the Linux system administration tracks offered at community colleges (e.g. PCC with which we partner on this pilot project), with the contents of the ACM/IEEE IT model curriculum **[6]**. In order to broaden our impact, we will integrate the new and existing courses in a distance education framework. We will also address technological issues stemming from the need to enable students to work from home on virtual machines on which they can experiment with system administration tasks. To sum up we want to;

(1) Offer a Linux system administration distance-education track of courses at PCC compatible with leading professional certifications. Today, these are only available as face-to-face offerings.

(2) Offer a Linux system administration distance-education track of courses for USF BSIT majors to take as electives for a technical specialization. These will be taught at an accelerated pace by leveraging other IT core courses' contents. To our knowledge, such a track has not yet been implemented in a 4 year IT curricula nor available with distance education options.

(3) Offer the above track as part of our IT Professional Certificate, thus enabling returning students to benefit from similar educational opportunities.

(4) Offer community college AS graduates the possibility of completing their education with core IT courses leading to a BSAS with USFL. More specifically, our intent is to offer a BSAS with a focus on Linux system administration that culminates with an "advanced topics" course, which will bring students to learn from state of the art, in field, research and development.

The next section will present both PCC and USF IT department curricular structures and discuss how they support the above objectives. Section 3 will detail courses to be developed in our IT department and how they will be used. Section 4 will discuss technical requirements to support these new offerings in a distance education framework. Section 5 will conclude by discussing ongoing and future work.

2. EXISTING CURRICULAR STRUCTURES

PCC offers, through the PCCNET track, many courses aimed at producing ready-to-operate AS graduates with a focus on Microsoft, Cisco and Linux technologies. Of particular interest to us, is a series of 5 courses which cover the fundamentals of Linux system administration professional certifications. These courses are;

CGS2760 Fundamentals of Linux: practical introduction to Linux (installation, desktop, shell, command line, file systems, editors, processes, network configuration, services & security…).

CTS1321 Linux System Administration I: basic Linux networked server administration skills (updates, monitoring, software installation, init scripts, networking, network services…).

CTS2322 Linux System Administration II: advanced Linux networked server administration skills (networking, security, backup & recovery, shell scripting, software compilation…)

CTS2301 Linux Network Services: setup and management of a Linux network (DNS, DHCP, printing services, Samba, mail server, proxy, tomcat, network traffic monitoring…).

CTS2311 Linux Security: securing Linux environments (network vulnerabilities, services, cryptography, firewalls, Application Level Gateways, VPNs, Intrusion Detection systems.

These courses constitute a strong technical foundation for an AS degree while preparing students to meet industry standards and needs. Our aim is to further develop these skills through an upper-level offering.

Our BSIT curriculum comprises core courses and five department-approved electives leading to a BS in IT. The former provide students with a solid foundation in IT (programming, data bases, operating systems, networking, etc.) which is key to guarantee their life-long professional development. The latter completes this core with five department-approved electives which can be picked from related or complementary disciplines. These allow students to choose between gaining a technical specialization or a multi-disciplinary profile. Our IT department also offers two IT certificates which allow returning students to update their IT knowledge from either a technical or management perspectives. Through these programs, our IT department has been serving a non-traditional student population composed of many full time workers favoring evening-time classes. This leads us to believe that a distance education program would allow us to better serve our student population.

USF also offers a unique AS to BSAS transition which allows students holding an AS degree to complete their education in order to get a Bachelor degree in Applied Sciences. Unlike other Floridian

institutions, this degree also allows them to continue toward graduate courses. Our IT department already offers such a BSAS in IT, which allows us to expand the horizon of AS graduates and better articulate our offering with these of the surrounding community colleges. The next section will describe the courses, which we plan on offering in the context of a BSAS in IT, focused on Linux system administration thus allowing AS students with previous knowledge in Linux technologies to complete their education and gain further specialization in this field.

3. CURRICULAR ASPECTS

This section details the courses which will allow us to achieve the objectives outlined in section 1. We focus on objectives (2), (3) and (4) from the perspective of the new contents developed in our IT program. The following diagram depicts a simplified workflow of the tracks to be offered.

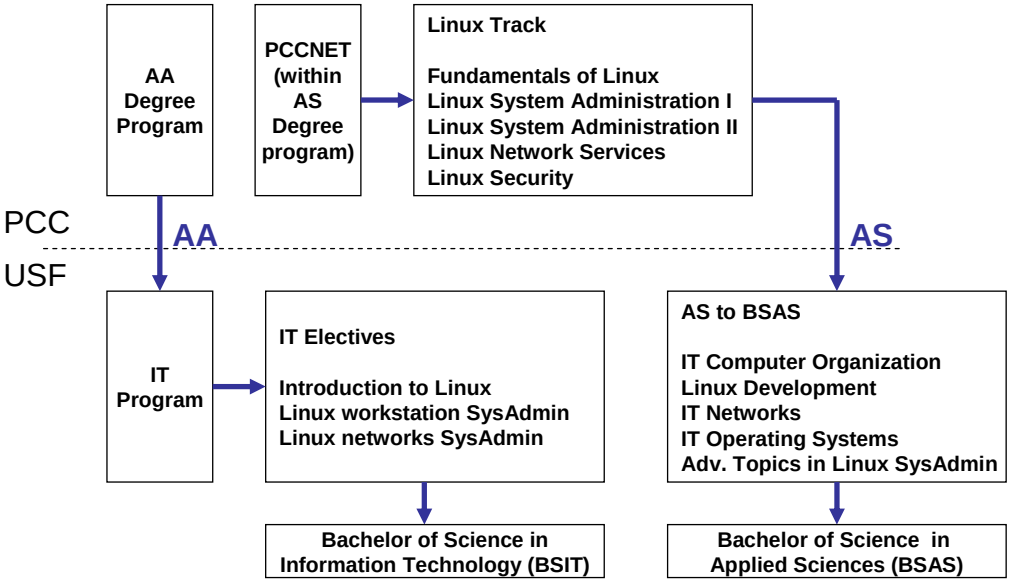

Objectives (2) & (3): IT electives & professional certificates

Our first objective is to offer to the USF IT students, who didn't take Linux courses at their community college, an opportunity to acquire system administration skills as a specialization of their IT Bachelor. To this end, learning units meant for integration in a distance education platform will be developed using leading professional certifications as guidelines. These learning units will then be offered by PCC but also recombined and re-designed to constitute three IT elective courses to be offered as an "accelerated pace" version of the PCC track to our IT majors. To allow for this accelerated pace, the content will be redesigned to rely more heavily on the knowledge already provided by other core IT courses. It is not our intent to replicate PCC offerings, but rather to offer a well tailored pathway for IT majors to cover similar material as a specialization of their degree. The three electives developed will be; Introduction to Linux systems (equiv. to CGS2760), Linux workstations system administration (CTS1321 + CTS2322 + CTS2311 as appropriate), Linux networks system administration (CTS2322 + CTS2301 + CTS2311 as appropriate). Besides IT majors, these courses will also be offered as part of an IT professional certificate. Once again, the learning units will be recombined differently and adapted to account for the difference in objectives of these professional certificates. By making these units available for others to adapt and adopt, we plan on contributing to the IT education community.

Objective (4): BSAS specialization

The second part of our IT department's curricular and course development effort will be focused on the AS to BSAS transition. More specifically, we want to offer a set of 5 courses which can be taken by AS students who followed a community college's Linux track in order to (1) prepare them to serve as "advanced Linux technicians" in the workforce, (2) arm them with the conceptual toolkit necessary for their continuous professional development and (3) allow them to pursue graduate courses at USF. This track can be taken in 2 semesters, students without preliminary Linux training will have to take the above-mentioned electives as pre-requisites. The BSAS-specific courses will be;

> COP 4610 Operating systems: concepts, algorithms and implementation of OS components (e.g. Linux kernel, processes, parallelism, memory management, resource allocation and file systems).
>
> CDA 3103 Computer organization: computer engineering topics usually covered in discrete courses on digital design, computer architecture and computer organization.
>
> EEL 4782 Networks: concepts of computer networks; Physical and logical structures, physical media, circuit switching, data flow, high-level protocols, and the ISO model are discussed.

Further, two other courses will focus on developing Linux system administration skills;

> Linux Programming: combines the relevant contents of the "program design" and "data structures" pre-requisites to COP 4610. Students are taught the C language and revisit common data structures. The focus is on low level and technical aspects of C programming meant to strengthen their skills and understanding of the link with OS topics. This course results from ongoing departmental work on the role of C in modern computing curricula **[18, 19]**.
>
> Advanced topics in Linux system administration: this course is the corner stone of this track, concluding our Linux system administration AS to BSAS transition by introducing students to advanced technologies and methodologies; virtualization, clustering and provisioning techniques, state of the art research led in configuration engines and autonomic computing.

4. TECHNICAL ASPECTS

Our intent is not only to develop an articulated curricular track in Linux system administration between PCC and USF but also to innovate in its implementation. This section discusses the challenges in providing distance-education students with remote access to their own Linux virtual machines.

Ongoing research: SOFTICE project

The USF IT department has been host, for the past 3 years, to a NSF-sponsored project, SOFTICE **[1]**. Its objectives were to provide an inexpensive hardware and software platform to facilitate the teaching of some IT and CS courses. More specifically, we focused on improving the pedagogy and classroom management of those courses, which traditionally require students to have privileged access to their workstation in order to engage in meaningful hands-on activities. For instance, in operating systems and networking courses, students are often engaged in activities which can be rather disruptive to the hosting network and classrooms. By using clustering and virtualization open source technologies, we provided a solution to facilitate the deployment of such laboratories **[3]**. This work resulted in a load-balancing scalable cluster infrastructure which can host an arbitrary amount of virtual machines using recycled classroom PCs and serve access to them over the internet to any platform. This infrastructure was leveraged to develop pedagogically innovative labs for operating systems **[4, 20]** and networking courses. Only recently, we started investigating how such a platform could benefit Linux system

administration courses and, beyond classroom-based offerings, how it could support a distance education offering. The following sections will discuss the technical challenges associated with such an offering.

Remotely accessible virtual machines

When the SOFTICE project was initiated, many institutions supported system-oriented courses with "custom" classroom management solutions to allow student fully privileged access to machines they could experiment with. Since then, virtual machines became the *de facto* standard in classroom management for such courses. With SOFTICE, we tried to stay one step ahead by allowing students to access their VMs over the internet and from any platform. No need for constantly moving around virtual hard drive images **[16]**, for keeping all the students VMs up and running constantly **[9],** or even for limiting the number of simultaneous connections despite a costly hardware investment **[17]**. The students simply login to our cluster and span entire virtual networks on demand. To achieve this, we relied on the User Mode Linux virtual machine project **[10]** and the Warewulf clustering toolkit **[11]**. In a technological field, things evolve quickly and a wealth of new technologies emerged as good candidates to replace the currently deployed ones. On the virtualization front, the Kernel Virtual Machine (KVM) is the next mainstream (i.e. incorporated in the Linux kernel sources) virtualization solution **[12]**. It relies partly on tools developed by QEMU **[13]** and would allow us to expand our infrastructure so that it can remotely serve VMs running any operating system. These new technologies, which leverage the virtualization extensions of recent AMD and INTEL processors, open new perspectives for us to improve at the supporting infrastructure levels. From the clustering perspective, new possibilities are offered by stateless and statefull provisioning toolkits **[15]** which might allow us to deploy and administrate on the fly both cluster nodes and virtual machines with the same technology.

Scalable virtual machines hosting

The first technical innovation of our infrastructure was to allow students to access VMs remotely, even over a relatively slow internet connection, and create/destroy them on-demand. The second main innovation focused on how these virtual machines were to be hosted and managed. While hosting all these VMs on a single server or dedicating fixed physical machines to some of our students' VMs **[9]** would be possible, we went one step further by hosting them on a load-balancing cluster of recycled classroom PCs. From the student and the system administration point of view, this cluster appears as a single machine through the use of the Warewulf clustering toolkit **[11]** which allows us to manage an arbitrary number of cluster nodes with a single machine. The outcome is that one can easily scale up the cluster's computing power without scaling up the manpower needed to administrate it. The transparency from the students' perspective is obtained through TCP/IP load balancing software such as the IPVS project and ultra-monkey. This software allows incoming connection to be load balanced over the entire cluster. Just like for the virtualization technologies, new technical possibilities have emerged lately. From the clustering perspective, Single System Image (SSI) clustering has matured into a commercial offering **[14]** which might enable us to overcome compatibility bugs between UML and Open SSI which prevented us from using the latter. The advantages of SSI clustering are tremendous in terms of facilitating the system administration of the cluster even further as well as in terms of migrating the virtual machines from node to node. Should SSI technologies still pose kernel-level incompatibilities issues with our virtual machines, we now have newly developed live VM migrations alternatives.

5. Discussion & Future Work

This paper presented a curricular design and supporting technologies for the development of a Linux system administration track articulating a transition between PCC and the USF BSAS in IT. While this paper focused on curricular aspects, we are also working on developing innovative pedagogical strategies to facilitate active learning even in a distance education setting. Also, the successful integration, in our supporting infrastructure, of virtualization with load-balancing clustering (Warewulf or SSI) constitutes by itself a non-trivial achievement which can both serve as prototype for industrial applications and for educational institutions.

6. ACKNOWLEDGMENTS

This material is based in part upon work supported by the National Science Foundation under award number 0410696. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

7. REFERENCES

[1] softice project's wiki, http://softice.lakeland.usf.edu

[2] A. Gaspar, S. Langevin, The case for Case Study Assignments in Undergraduate Operating Systems Courses, EIAE 2006, IEEE International Conference on Engineering Education, Instructional Technology, Assessment and E-Learning, December 4-14, 2006.

[3] A. Gaspar, S. Langevin, W. Armitage, Inexpensive, Scalable and Transparent Open Source Classroom Management for Linux-based Laboratories, EIAE 2006, IEEE International Conference on Engineering Education, Instructional Technology, Assessment and E-Learning, December 4-14, 2006

[4] A. Gaspar, S. Langevin, New approaches for Linux-based undergraduate operating system concepts laboratories, EIAE 2006, IEEE International Conference on Engineering Education, Instructional Technology, Assessment and E-Learning, December 4-14, 2006.

[5] A. Gaspar, S. Langevin, Active learning in introductory programming courses through student-led "live coding" and test-driven pair programming, **submitted** EISTA 2007, Education and Information Systems, Technologies and Applications, July 12-15, Orlando, FL

[6] ACM SIGITE IT model curriculum, available at http://www.acm.org/education/curricula.html

[7] CC 2001 Curriculum Guidelines for Undergraduate Degree Programs in Computer Science, ACM / IEEE, available at http://www.acm.org/education/curricula.html

[8] Stockman, M., Nyland, J., Weed, W., Centrally-Stored and Delivered Virtual Machines in the Networking/System Administration Lab, ACM SIGITE Newsletter, volume 2 issue 2, 2005

[9] Benjamin Anderson, Thomas E. Daniels; Xen Worlds: Xen and the Art of Computer Engineering Education; In Proceedings of 2006 ASEE Annual Conference and Exposition; Chicago, IL; June 2006.

[10] UML, User Mode Linux, http://user-mode-linux.sourceforge.net/

[11] Warewulf clustering toolkit, http://warewulf-cluster.org/

[12] KVM, Kernel Virtual Machine, http://kvm.qumranet.com/kvmwiki

[13] QEMU, http://fabrice.bellard.free.fr/qemu/

[14] Scyld clusterware, http://penguincomputing.com/

[15] Cobbler and Kuan, http://cobbler.et.redhat.com/

[16] Stockman, M., Nyland, J., Weed, W., Centrally-Stored and Delivered Virtual Machines in the Networking/System Administration Lab, ACM SIGITE Newsletter, volume 2 issue 2, 2005

[17] Border, C., The Development and Deployment of a Multi-User, Remote Access Virtualization System for Networking, Security, and System Administration Classes. ACM SIGCSE 2007 conference, Special Interest Group in Computer Science Education.

[18] Gaspar, A., Eijnioui, A., Boyer, N., The role of the C language in modern computing curricula: Part 1 – survey analysis. Submitted to CCSC-SE 2007.

[19] Gaspar, A., Eijnioui, A., Boyer, N., The role of the C language in modern computing curricula: Part 1 – beyond survey results. Submitted to CCSC-SE 2007.

[20] A. Gaspar, C. Godwin, Root-kits & Loadable Kernel Modules: Exploiting the Linux Kernel for Fun and (Educational) Profit, The Journal of Computing Sciences in Colleges, CCSC (Consortium for Computing Sciences in Colleges), Vol 22, Number 2, December 2006, pp. 244-250