MultiCampus *Parcel*: A Cooperative Approach To Computer Laboratory Management

Gregory M. Dick
University of Pittsburgh at Johnstown

Abstract

Maintaining student computing laboratories presents special problems to the laboratory - network manager. Tasks associated with managing student accounts and PC disk files consume large amounts of staff time. At the initiation of this project, commercial automated solutions to these problems were not available. Small campuses do not have the resources to develop automated tools to solve these problems. This paper discusses how *technology enabled cooperation* at the University of Pittsburgh has supported the disbursement of locally developed automated management tools to the University's four regional campuses.

Introduction

Providing student computing laboratories is a challenging task. In addition to the normal issues involved in managing a networked computing system, situations unique to the student lab environment exacerbate the management problem. Included among these are: (1) the issues surrounding the management of a large number of student accounts and (2) the maintenance of a stable disk configuration on a large number of client PCs. These problems could be effectively addressed only by building customized tools. Small campuses typically do not have the resources to construct such tools. Staffs are small and individual staff members tend to be generalists. (i.e., the small size of the staff does not permit the luxury of specialists.) Thus, special customized tools to assist in system management are typically not available.

The purpose of this paper is to discuss how *technology enabled cooperation* at the University of Pittsburgh has allowed the University's four regional campuses to share the fruits of development efforts conducted by Computing and Information Services at the Pittsburgh Campus.

Discussion

The Problem

Student laboratories present a particularly challenging set of problems to the network manager. In addition to the normal issues involved in managing a network, situations unique to the student lab must be addressed. Account management and disk file configuration management can be particularly troublesome.

Account management at even a small college like the University of Pittsburgh at Johnstown (UPJ) can be troublesome. Each year new accounts must be created for approximately 900

freshmen and transfer students. While not a difficult process, account creation is time consuming.

Student labs must present a consistent login environment each time a student "boots" a desktop PC (i.e., the PC must be returned to a known state regardless of possible modifications caused by the previous user). New applications and updated versions of hardware drivers and operating system components must be installed on each PC. Frequently used application software may be loaded on the local disk of a client PC, rather than served from the network, to enhance performance. Each of these situations contributes to the problem of disk file configuration management. UPJ makes more than 100 desktop computers available to its students. Maintaining the disks on these machines, like account creation, is a labor intensive task.

In 1993, when this project was initiated, commercial tools to automate solutions to both the account management and disk configuration problems were not available. The only approach to an automated solution was the creation of "in-house" customized software. This was not a viable option. The Academic Computing staff at Pitt-Johnstown included two network/laboratory managers. By necessity, both were generalists who devoted their talents to a wide variety of tasks. The small size of the staff, and the breadth of demands on its members, precluded the development of specialized management tools. Accordingly, much time was devoted to manual efforts to address these situations.

Solution: Prerequisites

In order to institute the use of automated management tools at Pitt-Johnstown, both the automated tools, and a means to deliver them must be established. Both of these prerequisites were satisfied.

Prior to 1993, computer labs at UPJ were isolated LANs. In 1993 these LANs were connected to PittNet, the University's WAN which is connected to the Internet. Thus, communications between the University's central servers and the regional campus servers was made possible. This technological step set the stage for the shared management of computing laboratories across the University's five campuses.

The University's Pittsburgh campus is an order of magnitude larger than the Johnstown campus. Within Computing and Information Services (CIS) resides a relatively large number of computing professionals. This technical resource has sufficient specialized expertise to develop in-house automated solutions to both the file management and account management problems described above. Furthermore, the size of the problem, 600 desktop clients, is sufficient to make the required investment cost effective. Given this environment, CIS developed the required solutions.

Solution: Disk File Configuration Management: *parcel*

A development team within CIS designed and implemented a client PC disk file configuration management tool which is named *parcel. Parcel* supports the centralized management of a large

number of similar, but not necessarily identical, network clients. A set of configuration files, which reside on the server, specify the arrangement of directories and files which should exist on each client. When the configuration of a set of clients must be changed, the appropriate files on the server are adjusted and *parcel* is executed on the clients. A simplified description of the *parcel* system follows.

Parcel is a client - server disk file configuration management tool developed at the University of Pittsburgh. On the *parcel* server reside the configuration files, operating system files, and application files. The configuration files describe the required client directory tree structure and the contents of those directories. The executable *parcel* client resides on the client PC. *Parcel*, client executable and server file collections, maintains the contents of the client disk.

Disk file configuration files are maintained as a set of three hierarchically related directory structures. Taken collectively, they completely specify the contents and structure of the client disk. The directory structures are:

- The system directory which describes files which are to be located on all PCs within the Pitt student laboratory system. This directory specifies the base operating system and related files.
- The enclave directory which describes files which will reside on a set of identical clients. This might be all the PCs in a specific lab which were purchased at the same time and share identical hardware configurations. The enclave directory specifies files unique to this collection of machines. e.g., The newest lab is equipped with high resolution 17" monitors, 8x CD-ROM drives, and advanced audio cards. The enclave for this lab will include the drivers for these devices.
- The host directory which describes files destined to be loaded to a single machine. Typically, the host directory is used to modify the configuration of one (or a few) specially equipped clients in a lab. e.g., One PC in each lab is customarily equipped with an optical scanner. It is the host directory for this individual machine which includes the driver, and associated application software, for the scanner.

Parcel uses these three directories to insure that the client disk is up to date. The process is:

- *Parcel* client retrieves the three directories (system, enclave, and host) from the server. The three directories are merged into a single unambiguous description of the required disk structure.
- *Parcel* client compares the file structure resident on the local disk with the required file structure as specified by the *parcel* directories. If differences are detected, they are corrected by either deleting local files or downloading the required files from the server.

The directories are interpreted in a hierarchical fashion. Host takes precedence over enclave which takes precedence over system. Thus, the system directory may specify all of the standard drivers and the enclave need specify only exceptions to the standard (rather than specifying all drivers in the several enclave areas - which would become overly repetitive).

Parcel file comparisons are scheduled according to one of three algorithms. File comparisons may be scheduled to occur on a regular timed basis e.g., every eight hours. Alternately, if the

client environment is relatively secure, file comparisons occur only if the client is out of synchronism with respect to the server. This approach is the norm. Each time the configuration on either the server or the client is updated, the time is recorded. At "boot time", if a client "update time" is earlier the server "update time", then the server has been updated since the last *parcel* of the client. This situation forces a file comparison and update of those files which are out of date. Finally, a *parcel* file comparison can be forced if client disk corruption is suspected.

A typical example of the operation of *parcel* is the installation of a new Windows application. Such an installation generally requires that Windows configuration and/or initialization files be modified. This information is distributed to all *parcel* clients by the following process:

- At the *parcel* server:
  - the modified Windows files are created on the server
  - the server configuration "update time" is changed.
- At each *parcel* client:
  - when the client is "booted" the server configuration "update time" is compared to the client configuration "update time" - the absence of synchronism is detected and a file comparison is initiated
  - the file comparison detects the modified files and updates the client appropriately
  - the client configuration "update time" is changed preventing subsequent file comparisons until triggered by another change at the server.

Thus, the required file modifications are accomplished on all client PCs by making the modifications on the *parcel* server and causing (via the "update time" compare mechanism) all clients to duplicate those changes on their local disk.

Parcel has been found to be an effective solution to the problem of disk file configuration management on a large number of networked desktop PCs.

Solution: Account Management: *PittAuth*

In order to grant controlled access to student laboratories, the University must create several thousand student Netware accounts each year. While the actual account creation process is straightforward, the sheer volume required caused this situation to place a large demand on staff resources. CIS sought an automated solution; *PittAuth*, an automatic Netware account creation tool, is the result.

The University of Pittsburgh makes several different computing environments available to its students. In addition to desktop computers, students have access to both UNIX and VMS time-sharing systems. This variety of platforms has, in part, motivated the use of MIT Kerberos as the core of its user validation system. All systems (UNIX, VMS, and Netware) authenticate login attempts by querying the Kerberos server. This allows a student to use a single username - password pair to access all public computing resources.

Kerberos is also central to *PittAuth*. The function of *PittAuth* is fundamentally simple. When a user attempts to login to a Netware server, *PittAuth* is executed. It:

- queries the Kerberos server to determine if the user is valid
- if the user is valid, and no Netware account exists, then one is automatically created.

In this fashion the process of Netware account creation is entirely automated. The first time that at user performs a login to a Netware server an account is created. During subsequent logins, *PittAuth* only authenticates the user as valid. This system has been found to be an effective and reliable solution to the problem of student Netware account creation.

Solution: MultiCampus Implementation

The University of Pittsburgh includes five campuses. The largest of the regional campuses is located in Johnstown, seventy miles east of Pittsburgh. The University of Pittsburgh at Johnstown enrolls approximately 3000 undergraduate students. 1993 marked the beginning of a project whose goal was to establish closer levels of cooperation between the Computing organizations at Johnstown and Pittsburgh.

The first technology step in this process was the "porting" of the laboratory environment developed by CIS for the Pittsburgh campus to Johnstown. The necessary technology prerequisite (i.e., a common network backbone) was established. The process continued with the installation of a *Parcel* server and the appropriate client software modifications. The entire process went smoothly and was deemed a striking success.

Success at Pitt-Johnstown prompted investigations from the other regional campuses. They determined that the use of *Parcel* and *PittAuth* to manage their labs was advantageous and proceeded with similar installations. At this time, all five campuses of the University operate servers which, at the base level, are identical. These servers are cooperatively managed by the staffs in Pittsburgh and at the regional campuses. The base configuration is built by CIS and distributed to the regional via T-1 data links during periods of light network traffic. Regional servers can be customized to respond to specialized needs. This might include the installation of local applications which are not available from the common software suite, or the installation of special purpose hardware (e.g., DSP daughter boards) and both the drivers and applications software required to support it.

Solution: the Benefits

The establishment of a cooperative approach to laboratory management has yielded several benefits. These include:

- The obvious, i.e., small campuses are able to take advantage of highly effective, custom-built management tools whose development is beyond their resources. This has had the effect of significantly reducing the amount of staff resource devoted to server and account management tasks. Thus, more time is available to devote to user consulting and training.
- A single software suite is shared across all five campuses of the University. In Johnstown, the number of applications grew from a few dozen to more than 150. Software licensing and purchase is negotiated centrally, thus campuses as small as a few hundred students share in the economies of scale available to a large university.

- The computing environments at all campuses are essentially identical. When students travel to another campus, they are immediately able to use computing facilities. In addition, CIS has developed a set of instructional resources structured around the specific Pitt environment. These resources are applicable to the regional campuses and their use removes the necessity to develop similar resources. "How to" documents, and others, drafted by CIS in Pittsburgh are available in the labs in Johnstown. "Quick Start" courses, again developed by CIS, have been offered at Johnstown and other regional campuses.
- Overall levels of cooperation between the several campuses has increased significantly. This promotes a more productive environment. The successful completion of the *parcel* project led to the establishment of the Five Campus Committee. This group includes members of the computing staffs from each campus. The committee meets monthly to address issues of common interest. Through this mechanism the regional campuses are able to impact decisions concerning the future of the Pitt computing environment. Relations between the computing staffs at the several campuses are improved. Staff members see themselves as members of a multi-campus team which broadly serves the same public and deals with the same problems. Cooperative problem solving lead to faster and better solutions.

Looking Toward the Future

This project was instituted in 1993 when effective and robust commercial tools to support automated account and disk file maintenance did not exist. The only solution was custom-built tools. The University is now planning for the next major step forward in the development of its computing environment. This initiative will include the merging of student and faculty computing into a shared environment in which all lab software will be available at faculty desktops. Many details concerning this project have not yet been determined, but it is apparent that commercial tools to support automated management are available and continue to improve. Both Netware and Windows NT, the prime candidates for the network operating system, include tools which address this issue.

As we move Pitt's computing environment forward the answers to several detail oriented questions is unclear. One item, however, is very clear. The principle of shared management of public computer laboratories across the University's several campuses will continue.

Summary

This discussion has, by necessity, emphasized technology. The technology of the software tools *parcel* and *PittAuth*. The network technology required to deliver those tools to distant sites. This emphasis, by word count, is unfortunate. The real success story in this project is one of human cooperation. This cooperation may be enabled by technology, but it is the deliberate decision on the part of a number of individuals to put aside regional differences and cooperate for the betterment of all!

Acknowledgments

A large number of people contributed to the success of this project. I hesitate to mention individuals for fear of omitting others. Nevertheless, I am compelled to acknowledge:

- The Staff of Pitt's Computing and Information Services:
  - Dan Temple - whose leadership and spirit of cooperation fostered the initiation of this project.
  - A wide variety of technical staff members, several of whom have moved on to other position both within and outside the University - it was they who implemented - in technology - the cooperation we sought. Notable within this group are: Kip Ruefle - who continues now to contribute to the project and has insured that my technical statements in this paper are correct, and David Schatz, Peter Donovan, and Robin Ruefle - without whose constant support the project would not have succeeded.
- The Staff of Academic Computing at Pitt-Johnstown - it was these individuals who chose to relinquish exclusive control over local labs so that the students might benefit:
  - Craig Gresko
  - Rob Eckenrod.

GREGORY M. DICK is Associate Professor of Electrical Engineering Technology at the University of Pittsburgh at Johnstown (UPJ). His academic interests include Digital Signal Processing, Digital Control Systems and Computer Applications to Engineering Education. He currently serves as the Director of Academic Computing at UPJ and has spearheaded an effort to foster greater cooperation between the University's five campuses.