1.0 Introduction

While there has been an interest in campus-wide information systems and videotext applications at the University of Pennsylvania for quite some time, it was in the Summer of 1991 that the University's Department of Data Communications and Computing Services (DCCS) led a formal campus-wide effort to determine the best approach for implementing a campus-wide information system (CWIS) for Penn. Penn's CWIS project team used a requirements-based process to determine the best approach to implement the system. One of the most important criteria was that it should support the decentralized organizational structure at Penn, permitting a wide range of information providers throughout the University to post their information to the system with minimal effort, training, or prerequisite computing expertise. A fundamental decision to be made was whether we should develop our own CWIS, utilize a public-domain system developed at another university (use it "as is" or modify it to meet our needs), or purchase a commercially available product. After a brief investigation of available systems, we narrowed our choices to three public-domain systems (Cornell's CUINFO, MIT's TechInfo, and Princeton's PNN) and a fourth system, which would be developed through an in-house effort. These choices were, in part, prompted by our knowledge that the developers of these three systems were pioneers in introducing CWIS systems to the Internet community (through conferences such as EDUCOM and CAUSE), and they were also collaborating on the development of a new protocol, CWIS-P, which would permit the "seamless" sharing of information between heterogeneous CWIS systems.

The result of these deliberations was a decision to use MIT's CWIS, TechInfo. [1] A critically important feature of the system design is its client/server architecture, which is consistent with the strategic direction for computing at Penn as articulated by the Vice Provost for Information Systems and Computing, Dr. Peter Patton. This client/server architecture enables us to distribute the responsibility for posting information on the system, and it reduces the potential for performance bottlenecks, which occur in traditional terminal to host systems, by utilizing the computing power of end-users' desktop workstations. Macintosh and VT100 implementations of the client software are the only ones currently available, but others are forthcoming. We made our decision while realizing that some effort would be necessary to customize TechInfo for our environment.

2.0 Implementation

DCCS engineers, Linda Murphy and Jerzy Sliwinski, customized TechInfo, as follows: o Linda changed the way the system is configured to permit automatic start-up via "crontab," a UNIX daemon. This provides unattended recovery from a system crash, minimizing the downtime that follows any such event. o Linda modified some text strings in the UNIX code, as required. For instance, during start up, the message

"PennInfo initializing" replaces "TechInfo initializing." o Linda wrote a program to capture local weather information, several times a day, by running the NNTP weather daemon. The source of the weather information is the University of Michigan, which permits noncommercial use. o Linda wrote several programs to generate usage statistics from system log files. Thus, we are able to tell information providers how often their documents are read (we have no way to determine by whom, but we would not want to do so because of privacy considerations). We also have access to gross usage statistics for the PennInfo system.

+ Page 7 +

 o Linda also wrote programs to generate both a document index and a series of keyword lists for each information provider's files. o Jerzy used ResEdit, a Macintosh programming tool/resource editor, to customize text strings, create new icons (such as the PennInfo icon that is used to fire up the PennInfo application from the end-user's Macintosh), and set certain default configurations for use in the Macintosh environment. The document, Site Specific Macintosh TechInfo Customization, provided by MIT, was useful in this process. The Penn-specific version, PennInfo, was successfully launched as a supported service in November 1991, and it has since grown to include approximately 3,000 documents, posted by over 70 information providers from various schools and offices at Penn. Recruitment of new information providers remains a critical activity that is necessary for the continued success of PennInfo as a network service that is widely used by members of the Penn community.

3.0 Description

PennInfo is a menu-driven system that includes information on a wide variety of topics, ranging from the University Calendar of Events to information about counseling and support services for staff and students. The first-time reader may want to traverse up and down the hierarchy of menus and read many of the documents available in the system. However, prior to moving through the menu hierarchy, an "Outline" feature is provided to indicate the type of documents that are available under the various menu headings. This enables the user to see a "road map" of the document structure before actually traversing it. A "Path" feature tells the user how he or she got to a given document in the hierarchy, which is a useful "navigational" tool. One of TechInfo's strengths is that, compared to other CWIS systems, it provides many more "traversal" aids to users. Help screens are also available to assist the novice user.

+ Page 8 +

While traversal through the hierarchy of menus and documents is useful in giving the user a sense of the breadth and depth of the information content of PennInfo, an easier way to find information on any given topic is to use the keyword search feature, which produces a dynamically generated menu of all items (e.g., folders, documents, or other menus) that have been assigned that keyword by any information provider. Since the base TechInfo software does not currently include a full-text search engine, the effectiveness of search efforts by end-users is limited by the thoroughness and experience of the many information providers who are responsible for deciding how many and which keywords are assigned to the various documents. Recognizing that choosing the right set of keywords for a document is an art, we have worked with experts from Penn's University Library to train information providers on appropriate techniques. An e-mail feedback mechanism has also been set up to provide us with the details of unsuccessful searches by end-users.

4.0 Technical Features

One of the design goals of TechInfo was to provide ubiquitous access from any computer. Thus, access from dial-up, Telnet, UNIX character-based (curses), and Macintosh environments is provided. PennInfo runs on a dedicated 24 MIP DECstation 5000/200 with two (RZ57) 1.3 gigabyte disk drives. The operating system is Ultrix v4.2. We also run a customized version of MIT's VT100 (curses) "client" software program to provide access for Telnet-based connections to the PennInfo server. A "point and click" interface is also provided via MIT's Macintosh "client" software package, which we have customized for the Penn environment and distribute to the growing number of end-users that use Ethernet-connected Macintoshes as their desktop devices. A key benefit of this Macintosh software is that it enables the user to print documents on locally attached LAN (AppleTalk) printers.

+ Page 9 +

5.0 Operational Support

We can provide operational support for PennInfo with existing staff because of their depth of experience in the UNIX and TCP/IP environments. The DCCS Operations staff also creates POI (Provider of Information) accounts, with password control, for the many information providers that have joined the PennInfo program. Over 70 individuals from about 30 offices (called "Sources" in TechInfo jargon) participate as information providers. End-user support has been provided primarily by existing staff at Penn's Computing Resource Center, but demand has been minimal. We distribute a Quick Reference Card, which explains the various features of PennInfo. Training and ongoing support of over 70 information providers, however, is a very time-consuming undertaking in Penn's decentralized environment, and it requires a dedicated staff person. Our PennInfo Administrator, Gayle Belford, trains these persons in the use of the TechInfo "Provider" software (which, at the moment, requires the use of a TCP/IP-connected Macintosh), and she also provides guidance in the areas of planning the menu structure, screen presentation, and keyword selection techniques. This care and feeding of providers is facilitated by the use of an e-mail mailing list, which is augmented with occasional, well attended meetings of the POI SIG (Provider of Information Special Interest Group). As with many UNIX-based systems, there have been problems related to the number of simultaneous users that can be supported with PennInfo. Specifically, the number of "pseudo-ttys" that can exist on an Ultrix system, as shipped by the vendor, defaults to 32. The number of "open file descriptors" per process, on Ultrix systems, likewise defaults to 64. To resolve these problems, Linda Murphy of our engineering staff has made some modifications that effectively permit up to 170 simultaneous accesses to the system. These changes are particularly useful during course registration periods each semester, when many students use PennInfo to obtain course information posted by the Registrar's Office. Another important aspect of the PennInfo program is the need for kiosks to provide access to those in the Penn community that do not have connections to the network. We have accomplished this by providing free connections to qualifying offices (preferably located in public-access areas with high traffic). Additional work remains to be done, however, to develop an integrated PennInfo Kiosk application for this purpose. We have been working with MIT on this matter.

+ Page 10 +

5.0 Future Plans

We have recently upgraded to version 3.1 of MIT's server (along with version 3.9 of the Mac client), and we are excited about new features that have either been released or will be forthcoming. Of particular interest are: o Worldwide TechInfo, which gives the user access to other systems on the Internet that use the TechInfo protocol. o GIF support for TechInfo desktop clients with graphics capability, which will permit the display of GIF images such as campus maps. o Built-in full-text search capability using the WAIS search engine. o A "What's New" function that will allow the end-user to enter a date and find all documents that have been created or modified since the date entered. o An X-TechInfo application (minus provider functions) and a DOS/Windows client software package. o A kiosk-mode function. o A feature, to be available on the VT100 client only, that will enable users to automatically send a PennInfo document to anyone via e-mail. Because of the tremendous Internet-wide success of the University of Minnesota's CWIS and "navigational" system, Gopher, we have recently adopted a "best of both worlds" strategy and set up a central Penn Gopher server (gopher.upenn.edu) along with two bidirectional gateways, which were developed by DCCS engineer, Linda Murphy. [2] These new additions to the PennNet infrastructure effectively enable end-users at Penn to use the CWIS client software of their choice to get to both PennInfo and worldwide Gopher information.

+ Page 11 +

6.0 Conclusion

In summary, the PennInfo project has been very successful, both as a key network service to the Penn community and as a working partnership between the developers at MIT and the project team at Penn. We invite readers to browse through PennInfo and to enjoy reading the many interesting documents found there, including Zen and the Art of the Internet, The BITNET List of Lists, and the CIA World Fact Book. [3] We look forward to getting feedback and comments from our users.

Notes

1. Information and documents about MIT's TechInfo are available via FTP from MIT's server, net-dist.mit.edu, in the /pub/techinfo directory.

2. To access Penn's central Gopher server, Telnet to gopher.upenn.edu (or "point" your Gopher client to gopher.upenn.edu, port 70). DCCS has put up some of the better Gopher clients (preconfigured to point to our Gopher server where possible) on our FTP server, ftp.upenn.edu, in directory pub/gopher.

3. To access PennInfo, Telnet to penninfo.upenn.edu. If you have the MacTCP software (developed and sold by Apple Computer, Inc.), use FTP to obtain pub/mac/penninfo.hqx from ftp.upenn.edu.

Acknowledgements

The author wishes to express thanks to Gayle Belford, Linda Murphy, Jerzy Sliwinski, and Dan Updegrove of DCCS, and to Steve Neiterman Wade of MIT for reviewing and providing useful suggestions for this article.

+ Page 12 +

About the Author

Alfred C. D'Souza has led the efforts to select, deploy, and provide ongoing support for Penn's PennInfo CWIS. His address is: Alfred C. D'Souza, Director of Program Management, Data Communications and

Computing Services, University of Pennsylvania, 3401 Walnut Street, Suite 221A, Philadelphia, PA 19104. Internet: DSOUZA@DCCS.UPENN.EDU.