



Adding Big Blue to the Palette: Plans, Projects, and Problems

Martha R. Sites, Associate Director

mrs@virginia.edu

Alice G. Howard, Manager, User Services

agh@virginia.edu

Jayne P. Ashworth, Systems Analyst, User Services

jpa@virginia.edu

Janet C. Sakell, Senior Programmer/Analyst, User Services

jcs@virginia.edu

Sherry A. Heitchew, Programmer/Analyst, User Services

sah@virginia.edu

Academic Computing Center

Gilmer Hall

University of Virginia

Charlottesville, Virginia 22903

(804)924-3731

Institutional Profile:

Enrollment (Fall, 1988): 17,610 Faculty: 1,604 Staff: 7,432

Undergraduates: 11,197 Graduates: 4,298 Professional Schools: 1,703

Continuing Education: 412

Functions: Instruction: 52%, Research: 47%, Administration: 1%

Facility Personnel: 62, FTE: 60, Part-time: 2

The Academic Computing Center provides instructional and research support to faculty, students and staff of the University of Virginia through a variety of computer resources (mainframe, micro and minicomputers) and services (including consulting, user training, professional programming and data entry).

A PLANNING PROCESS OVERVIEW

The Beginning

During the '85-'86 academic year, the University of Virginia's Computer Policy Committee assessed the "future of mainframe computing" at the university and recommended that the Academic Computing Center acquire an IBM-compatible system (3090 class) in addition to the systems already supported. The IBM-compatible machine was to supplement our CDC 180/855 which was approaching saturation, to facilitate the flow of information and software between UVA faculty and their colleagues at other "peer" institutions (mostly using IBM mainframes), and to enable us to offer the most current releases of various software applications. Also, researchers had long been frustrated by the CDC's lack of virtual memory and database management tools.

This recommendation was combined with the perceived need to supply some local ongoing vector-computing capabilities and the result was a decision to acquire an IBM 3090 (entry level) with vector facility during the '87-'88 academic year.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.
Copyright 1989 ACM 0-89791-330-2/89/1000-0411 \$1.50

Guesstimates and Negotiations

From the summer of 1986 into the spring of 1987, the Center's management speculated about the implications of this decision -- there were many, many questions and almost no immediate or easy answers. Our physical space was cramped and our existing staff stretched thin over multitudinous systems already -- where could we find enough room for all those blue boxes? Could we get additional staff positions, and if we did, where could we fit more desks? We would have to train staff (Systems, User Services and Operations), select and install software, write local documentation and prepare to teach Short Courses. Suppose we had to do it with existing space and staff -- how would we find room, adjust priorities, reduce support for existing systems and services? And what were the likely needs and characteristics of the potential users of this new system? Over the course of the following two years, we answered most (but not all) of these questions.

Meanwhile, the Director entered into extensive negotiations with IBM on the terms and conditions of a Research Special Bid, while maneuvering his way through the labyrinth of state-mandated paperwork for such a large procurement and a sole-source justification for an IBM 3090-150E with vector processor, which was approved by the state's Department of Information Technology in October, 1987.

Planning, Preparation and Installation

By May, 1987, we had a planning outline with the following section headings: Time Tables, Renovations, Hardware, Software, User Support, Staff Support, Accounting, Site Visits, and Miscellaneous. We also had answers to some of our questions: we were to embark on some renovations --turning office space and a user area into a new machine room, and creating two new office areas (Systems and User Services) from space acquired from a neighboring department. We would adjust priorities and assignments for our existing staff -- hence our CDC would go into maintenance mode (no Short Courses, minimal software upgrades, etc.), and our Primes would get reduced support.

In October, 1987, the Director appointed an IBM Working Group to coordinate the installation of the 3090. The group consisted of the Director, an Associate Director (chairperson), the Systems and Operations Manager, the User Services Manager and Large Systems Support Coordinator, and IBM representatives. Bi-weekly meetings were begun in November, 1987 and continued until June, 1988.

After soliciting suggestions and advice from the group, the Associate Director issued version 1.0 of our UVA ACC 3090 Installation Master Plan in December, 1987. Related installation tasks were grouped into categories, staff members were assigned responsibilities for each task, and target dates (for starting and completing each task) were specified. As the group continued to meet and reality impinged on our good intentions, the Master Plan was revised and rearranged -- responsibilities were re-assigned and target dates adjusted (mostly forward). Our initial target dates for the completion of the machine room renovations and the installation of the hardware and operating system (VM/CMS) were the last week in April, 1988. The actual completion of these tasks took place in late May and June --by then we had a name for our new system (watson) - - and our first "guinea-pig" users (a very small number) who were allowed to start using the system in an on-your-own, unsupported fashion in July, 1988.

During May the installation phase was nearing completion, but there were many implementation tasks and decisions that still needed consideration and coordination. A new ACC-IBM group was formed to take on these functions, as the initial work group disbanded in June. This new group is larger and more fluid (often consisting of 6-7 User Services staff, 5-6 Systems staff, 1-2 IBM representatives and the Associate Director/Chairperson), and continues to meet regularly with an agenda that has encompassed: the initial system configuration, accounting, user aids, database support, staff and user education, software selection and installation, local documentation, communications, and magnetic tapes, plus "bugs" and user problems.

In reflecting back upon this process, the formation of the two working groups, together with periodically soliciting input from the affected managers and members, did perform a major role in providing the leadership, enthusiasm and staff participation needed to accomplish our goals.

Some areas in which we anticipated difficulties turned out to be manageable (e.g. the initial constraints on physical space and staff, getting familiar with a new vendor's mores and conventions, etc). There were some surprises in terms of the IBM's capabilities (or lack thereof) which continue to sap much more of our resources than we anticipated (e.g. accounting and magnetic tapes), plus our local IBM office is not really fluent in VM/CMS and this has been a drawback. Some differences in the Center's staff became more obvious and problematic during this installation and implementation process -- we have a fairly autonomous organization, with real differences in leadership styles and group priorities, in addition to a staff who are accustomed to organizing their own work and who are resistant to deadlines and more direct supervision. There was not consensus and/or unanimous enthusiasm for this project among the staff, and this contributed to a sense (perhaps more perceived than real) of "foot-dragging" upon occasion. Finally, the continuing efforts to try to guess at the characteristics and needs of potential (but unknown) users has contributed more uncertainties to this endeavor than we anticipated.

PROJECTS: AND THE LIST GOES ON

Education, Training and Documentation

We found that education, staff and user training, and document production all went hand-in-hand -- intertwining most days, tangling others! We found out very quickly how inseparable they are.

Teaching the Teachers: Staff Training

When we began to address the issue of staff and user training, we realized that we did not have enough information to know where to begin. We decided to try to obtain some of this information from other academic computing centers at other universities. We used the *1987 Directory of Computer Facilities in Higher Education* (edited by Charles H. Warlick and published by The University of Texas at Austin Computation Center in cooperation with Seminars for Academic Computing) to compile a list of universities whose undergraduate enrollment was greater than 10,000 and who had academic computing centers operating IBM 30xx or 43xx machines running VM/CMS. This list of 68 candidates was winnowed to 14 sites that seemed to be most like our own. Our Director sent letters to these facilities asking each of them to send us a "CARE" package that would include a list of non-IBM software provided on their IBM system, descriptions of relevant introductory courses they offered (with copies of course materials, if possible), and copies of any locally written primers and/or user manuals. We received responses from 8 locations. We cataloged the material and completed an evaluation for each "CARE" package.

After we had reviewed the packages, we decided that the fastest way for us to learn more was to have someone else teach us, and that the best way to do this was to visit other sites. From the schools that sent us packages, we selected three sites that were geographically close. We contacted them about a site visit and compiled a list of questions to ask. On each visit, we sent a team of User Services and Systems programmers, who prepared a written report and then participated in a question-and-answer session with the members of the installation team. We found these trips to be enlightening, educational, and inexpensive. We learned from these trips of the lack of support for a good user and magnetic tape interface. We also learned about SELF-TEACH, an interactive tutorial for the IBM editor XEDIT; this has proved to be an excellent product. In the process we established some useful contacts who have continued to help us when needed.

Other options that we pursued for our own staff training included: working on UVA's Education School's IBM 4361 (they had limited software and were running a different level of the

operating system, so this was of limited use); reviewing videotapes (the ones we saw were not very useful); having our local IBM representatives arrange presentations for our staff (we seemed to miscommunicate what was needed and the sessions did not meet expectations), and attending a variety of conferences and classes. Our Systems people found the annual VMWorkshop conference to be worthwhile. Some of our User Services staff attended helpful IBM-offered training sessions on the use of SQL/DS (a database system) and the use of the Vector Facility. Our User Services manager attended a conference sponsored by ACIS (IBM's academic support group) and SHARE (the IBM large systems user group conference). Both of these conferences had much relevant material.

Looking back on staff training, there was so much to learn that finding a better way to learn earlier in the process would have helped. Given the quality of the IBM-offered courses some of the staff attended, it might have been useful and effective to have sent people to IBM-offered introductory courses, too. Under ideal conditions, we would have had formal training to make us literate early in the process, more time to become experts and iron out implementation details, and time to complete the documents and courses before allowing any user access. In our real setting, we were able to learn and get the job done within the parameters we had.

Teaching the Helpers: Consultant Training

We shared the information we learned with our student consultants as quickly as we could absorb it and give it to them. We adapted a training program used on another computer to give them a chance to experience use of the IBM first-hand. We also had them learn the editor by using SELF-TEACH and learn CMS by working through short course notes from a site we had visited.

Teaching the Masses: User Training

We tried to give our users the information they needed as we learned it. We had several articles in our Newsletter before the machine came on-site. After the machine arrived, we continued to keep our users updated about its status with the news programs on our other computer systems. When the IBM was made available to our initial group from the user community (after only a month of staff training and development), we kept the users abreast of changes by implementing a news facility on the IBM that was similar to the facilities on our other computers. Our first attempts at an on-line news facility using VM/SP's CP LOGMSG facility were not very successful. This facility had no editing capabilities -- changing a misspelling meant retyping an entire line -- and the messages kept disappearing. We finally developed our own news system using an XEDIT profile, the FILELIST command, and a batch file. This has worked beautifully for us and the system is easy to maintain.

When drafts of the introductory documents were complete, we used them as the basis for our "beginners" short course, which we first taught in the fall of 1988. This hands-on course was offered in two sessions with each session lasting two hours. The first session targeted terminal emulation and basic CMS commands. The use of the ASCII terminals in the IBM environment proved to be as difficult for our users as we expected it would be, and this part of the presentation was quite lengthy. The second session focused on the use of XEDIT and CP. Since we were still learning and the system was still changing daily when the courses were offered, our deliveries were often challenging.

Despite our preparations, there were difficulties when our first users began their work. These difficulties were caused because the users, who were chosen because they had previous IBM experience and would need little help, had questions that were frequently beyond the knowledge of any of our staff. This "trial by fire" time caused us all to learn more quickly. Sharing problems and solutions via electronic mail helped to increase not only our knowledge base but also to decrease the response time for problems. This sharing was simplified by creating a mail alias that allowed messages to be sent to the entire installation team with the use of one address regardless of the system used to originate the mail.

The Mountain of Manuals: Vendor Documentation

The Center was warned that IBM provided documentation for the 3090 in "great" quantity - "amazing" would be a better adjective! Since February 1988, the coordinator of the IBM document library has filed between eight and nine hundred manuals and bulletins or enough material to fill 5 four-drawer file cabinets and a dozen shelves. Over 200 manuals were returned to IBM. And the manuals continue to arrive. Initial preparations for the library, which included allocating a small room and a person to manage it, helped somewhat; dealing with this influx of documentation in reality was a massive chore. The sheer volume was intimidating in the beginning and there were real questions as to which manuals were really necessary.

It was decided early on to assign one staff member responsibility for the "IBM library." This library coordinator receives all manual shipments from IBM, files the material, and then enters the data about each manual into the library's database via a locally-written and maintained dBase III program. While it is still being fine-tuned, the database software allows searches by IBM order number (very important), keyword and index. Results of these searches may be printed for future reference. We have also developed a "check-in/check-out" method for staff to try to insure that manuals can be located at any time; this has been somewhat successful.

As the library has grown (an original estimate of four file cabinets to hold all materials was quickly "revised"), so too has the need to determine what manuals are or will be most useful. This has been both an organizational and economic concern. The technical publications were reviewed by three upper-level staff people and only those that suited our needs were retained. In consultation with IBM, other universities with IBM equipment, and experienced IBM users, a staff member investigated and made recommendations for individual manuals (a set of five manuals including CP, CMS and the Editor), "communal" manuals (including most of the language manuals), and manuals to be sold by the Computer Store. Since IBM manuals are expensive, we have tried to include as much as possible in our local documentation and we encourage users to "preview" manuals belonging to staff and consultants before purchasing their own sets.

Tunneling through the Mountain: Local Documentation

The writing of local documentation paralleled the education and training process. The Center has always provided locally-written materials for the user community; these materials range from a Newsletter to one-page "how-to's." The Center has no technical writing group and so over the years documentation has picked up individual and sometimes peculiar flavors. In the Spring of 1987, a committee had been formed to consider guidelines for style, presentation and content, and to offer some ideas for reorganizing documentation. It was decided that we would "beta-test" these recommendations on our new IBM documentation -- a clean slate upon which to work. We began with the "CARE" packages, the University of Delaware's Academic Computing Services' *VM/CMS Handbook*, and existing Center documentation for our other computer systems. Four User Services staff members formed the "IBM Documentation Group": we devised a list of necessary introductory documents based on what we provided for our other systems and then pulled from the "CARE" packages and the Handbook things we thought worked well and schemes for adapting them to suit our needs. These four staff members, for the most part, wrote and reviewed all the introductory IBM documentation. We established style rules (some times in conflict with what an individual had previously used), a mandatory text formatter (again, some had to learn a new product), and target dates for reviewing documents. For six months, this group was immersed in every aspect of document production; it was an exhausting process but it produced some of what we feel is the Center's most cohesive documentation. It is an easy set of documents for the user to consult; commands are presented in the same manner even in different documents, there is a common organizational scheme, and the documents mesh. At least one member of the original group reviews all new IBM documentation by new authors thus maintaining this continuity.

Among the newly implemented documentation ideas has been the concept of a "meta-doc." This is a short document itself but it provides pointers to all other IBM documentation; users need only read this, make a list of what documentation best suits their needs, and gather that set of materials. While benefitting users, it saves some printing costs for the Center as users do not feel the need to "take one of everything." All IBM documentation has been printed on paper with a

new masthead that is the same color and style as our Newsletter, giving Center documentation one "look." We include appendices in most documents explaining our style conventions and try to avoid revising multiple documents by referencing other documents wherever possible (and making the revisions to only that one document).

Things would have been easier had the documentation standards been in place before we started the IBM documents; attempting to implement the standards and write documents was tough. The results were worthwhile, though, not only in terms of quality documentation but in the improvement in writing and reviewing skills for staff and an increased awareness (not just for the original four) of how important good local documentation is to users. Some would argue that losing the "individuality" of a document is bad; in our case, it was the only way to have the documentation out in time for our users and has made it much easier to go back and revise documentation you may not have originally written. No one can tell where the original ended and the revision began. Revision has been a constant process with these documents and without consistent style guidelines, they would be unusable at this point! There are also differing opinions about the role of outside documentation; did using it help or hurt us? The issue was complicated because we were imposing new style conventions; we were also having to adapt the tone of the presentations to blend with the Center's overall tone; and we were frequently going from SCRIPT to troff, not a trivial conversion. Despite these problems, we would not have been able to publish the documentation as quickly as we did without using this outside documentation, so we appreciate the other sites who were willing to share not only output, but also source for much of the text. As we begin our second year of writing "how-to's" for the IBM, most of our original choices are standing us in good stead and producing and reviewing these documents is easier than writing for some of our other systems.

SOFTWARE SELECTION

Picking and Choosing

As we began to look at software, we reviewed the software lists we received from other sites as well as software usage on our other computers -- we tried to determine what was popular and why. We also looked for software that would fill needs that users might have that were not met on other machines. After we had this list, we called the vendors for pricing and licensing information. When all this information had been received, we determined that the purchase of SPSS-X, SAS, and IMSL would meet the needs of most of our initial users and would fit within our budget. Other products that we purchased were XLAT (an IBM FORTRAN '66 to FORTRAN '77 translation tool that would also help convert some programs from our Control Data mainframe to IBM), a version of ProComm that would help ASCII microcomputer users emulate an IBM 3270 terminal, LaTeX (a formatter that is rapidly becoming the "standard" formatter at our site), and SELF-TEACH (an interactive XEDIT tutorial). We also purchased VMtape which provides a better user interface to magnetic tape usage than the standard IBM system does. This is a useful product but we are still having difficulties in meshing it with existing tape management systems on our other machines. After these problems have been resolved, we will make this product the standard user-magnetic tape system. We also acquired the Rice Mailer (a public domain mail product) because of its compatibility with the other mail systems we support and because we thought it offered a better user interface than the IBM mail product.

Because of the complexities and capabilities of database systems, we considered database software separately from other software. We had a team of User Services staff contact 12 other sites by telephone and electronic mail to ask what products they were using for database support and why. Most were unable to quantitatively assess the packages installed at their sites but were able to provide us with substantial qualitative information -- how responsive the vendor was, quality of documentation, and user-friendliness. Many sites use DBMS packages for administrative functions only or support a product only out of consideration for a small, but vocal user group. The process of evaluating our database management needs and priorities continues with faculty involvement; due to this and other factors, we have opted to use the SQL/DS software provided by IBM for the present time.

In retrospect, we seem to have met the needs of our users in the software area. As they begin to gain more expertise and are able to share more easily with colleagues at other institutions, we may hear more requests for other software. We have been pleased with the speed with which many of the larger products run on the IBM. Graphics output has been a problem in our ASCII, asynchronous environment but we are finding ways to cope with this. In the same vein, high quality print output is now a reality. The capabilities of the database package continue to excite us.

THE SECOND YEAR: OUTSTANDING ISSUES

The first anniversary of the 3090's arrival has come and gone and we still have many issues left to tackle. Accounting is "in progress"; accounting software for our other systems was written and is maintained in-house and bringing the IBM into the accounting fold has not been an easy task. On our other systems it is possible to capture accounting information on a user-by-user basis; the 3090 has not been as cooperative. With only two tape drives and two cartridge drives, users have besieged us on a daily basis with requests for more than the default four cylinders of disk space for file storage. Disk allocation will become more problematic as we install more software, databases, and more users in the future. Communication problems -- 3270 emulation -- have yet to be completely ironed out. We do not have IBM terminals. While IBM's 3270 emulation has worked since the beginning, in our public labs we could not change the default colors so most of the IBM screens were "invisible" on the majority of microcomputers. The emulation package we first chose to use would not work on these microcomputers either. While awaiting source for that product so that we could make our own "bug" fixes, NCSA released a 3270 emulation for their PC telnet package. A year later we have a working 3270 emulation but still are dealing with key mappings for the various other terminals and communication packages in use at the University. Procedures for magnetic tape usage are still under development. VMTAPE is not a panacea for tape problems -- there have been user problems, staff problems, and operations problems stemming from our lack of expertise and its installation. VMTAPE's inability to read unlabelled, multi-volume tape sets dramatically affects our ability to move large publicly available datasets (ICPSR and CRSP, for example) to the 3090. We have supported unlabelled tape usage on our other systems for many years (our Primes cannot write standard labelled tapes) and moving to the standard labels preferred by an IBM system will affect both Center staff and our user community and must be addressed.

We hope that by the time the 3090's second anniversary rolls around, we will have resolved these major outstanding issues. As with any computer system, there are likely to be some problems until the day the machine is unplugged. But as staff and user expertise and understanding of the system grow, so too should our ability to deal efficiently and calmly with these problems.