



PROBLEMS IN IMPLEMENTING AND PROCESSING COMPUTER CHARGING SCHEMES

By Michael F Morris
Smart, Inc.
1919 Freedom Lane
Falls Church, VA 22043

It is important to point out at the beginning of this presentation that we have strayed quite far from the titled topic of our workshop -- "Pricing Computer Services." This makes my task much easier because I'm not at all sure what "service" we get from computers and "pricing" is seldom related in any economic sense with the cost of production. Here we have really been discussing "Charging for Computer Resource Usage." I will stay with the topic as we've been discussing it rather than with the topic as I thought it should be. To make the distinction clear between pricing services and charging for resource usage I will relate a very simple story from a recent newspaper.

There is a service called "Date-a-Dog" operated in the Chicago area. The operator of this service accepts up to about ten attributes that the owners of female dogs want for their dog's potential mates. An automated search of the "Date-a-Dog" files produces a list of the four males that most closely match the specified traits. This list is supplied to the client for \$50. This is a price for a service.

The price is not connected in any apparent way to the cost of using the data processing equipment. Nor is the service that is provided one that would usually be thought of as a computer service. In this simple example, the most basic card processing equipment could perform the file search and a stand-alone printer or even a clerk could convert the four selected cards into the prospective breeder's list. None-the-less, the client is paying an acceptable (to him) price for a service that would be difficult to duplicate without the central processing facility. The data processing cost to generate each list would only be of concern if it gets close to the price of the service to the client. And most of us here could probably imagine how very small the computer resource cost for such a simple task would be. My guess would be in the \$1 to \$5 range unless the "Date-a-Dog" file was extremely large.

This is really quite a typical case of pricing a computer service, but the topic here has become computer cost recovery. So no more about "Date-a-Dog."

Cost recovery, or its intra-organizational counterpart -- cost allocation, has been the accepted method of paying for computers since their inception. The only complicating factor has been the increasing number of accounts over which these costs must be distributed and the number of accounts that can incur costs simultaneously. At the outset of our technology, the computer had only one "user." This evolved into "time-sharing" in its earliest definition: each user took over the computer completely for some portion of each day, week, or month, and paid the percentage of costs according to the percentage of time that he used. In general, this was

accomplished after the fact as a bookkeeping exercise by the overall organization that "owned" the computer. The basis for these cost allocations was manual accounting logs showing the times that each user took over and relinquished the facility. Even when outside users came into the early computer centers, this simple manual approach was sufficient for cost recovery. There were no difficulties in implementing such charging mechanisms in the technical sense. The only difficulty was to insure that each new user remembered to log in and log out on the time accounting sheet.

The first technical problems with charging for computer time began when the idea of multiprogramming became reality. The reason that this created a real problem was that time in the system was no longer a reliable measure of the amount of computer resources used by any one program. That is, a program with the same input-process-output load could require different times in the system when run with different companion programs in a multiprogramming jobstream. This was a very difficult problem for two or three years until the idea of treating resources in two categories was developed in the early 1960s. These two categories were shared and dedicated. Dedicated resources (like tapes, printers, etc.) were easily treated just as the total system had been treated previously because time continued to have a fixed meaning relative to the use of such devices. Shared resources (like processors and in some cases memory) were tracked separately for the time and quantity of resource used to service each program. By 1970 nearly every computer supplier and many software houses offered some form of accounting package that produced resource usage data in both of these categories on a program-by-program basis and in such a manner that repeated runs of the same program in different mixes would yield a "repeatable" charge.

(This repeatability issue still seems to raise some discussion in the performance evaluation arena that is no longer of much interest outside this arena. I take "repeatable" to mean that re-runs of the same program would produce charges that varied by something less than 5%. Customers don't complain about this kind of variation -- as long as it can work in both directions. And all of the successful commercial accounting packages that I know of are well within this range. As a point of passing curiosity, it is interesting that the only people who still seem concerned with this repeatability issue are those connected with installations that haven't yet begun to bill their users. When such installations do begin to charge for their computer usage, the discussion of repeatability is seldom heard again.)

A companion "problem" that is sometimes raised when discussing charge-back in multiprogram environments is that some portion of the computer system is used simply to allow multiprogramming (among other things) to take place. This is the "system overhead." I state this as a, quote-problem-unquote, because this is only a tiny part of the total overhead, and most of us have no problem taking care of all the rest of the overhead without calling it a problem. There was a time long ago when the cost of the computing equipment was the major part of the total cost of the computer installation. This hasn't been the case for several years,

Now, the largest part of the installation's cost is for people. Most current estimates show that about 75% of costs are directly allocable to personnel and that the equipment and supporting software is in the 10 to 15% range. Many installations actually spend more each month for such expendables as paper and cards than for all of the computer equipment and its control programs. I can't see allocation of system overhead as any more of a "problem" than that of recovering the cost of personnel, or, for that matter, of the electric bill.

The spectrum of overhead recovery schemes is extremely wide. The simplest is to place a fixed charge on every run based on the expected total number of runs that will recover all overhead costs. In practice, there may have to be two or more fixed charges based on such things as type of user (batch or remote), run priority required, user department, and so forth. We have been shown a very rich and extremely interesting example of multiple levels of overheads and cost recoveries in a whole-sale - retail environment in Einar Stefferud's earlier discussion. I will give a very simplistic example of a way to develop overhead allocation.

Suppose that examination of available accounting data shows that during a month an average of 6000 batch runs and 4000 terminal sessions occur. If total overhead costs were observed to be \$24,000 in an average month, then a charge could be placed on each batch run of \$3 and on each terminal session of \$2 to recover the \$24,000 overhead plus \$2,000 either as profit or as a hedge against variations in runs and sessions from the established averages. A case could be made that this would be unfair as some terminals may only be used a few times each month and not even pay the equipment cost of these terminals. Supposing a total of 40 terminals, this overhead allocation could be modified to charge \$100 per month per terminal plus \$1 per session. The overhead would then be distributed "unfairly" in favor of terminals that were heavily used. More analysis might suggest a "bulk-discount" for terminals that are heavily used and a "surcharge" for lightly used terminals... and so forth.

The distribution of overhead can be as simple or as complex as suits each individual installation. What is important here is that no particular attention need be given to the specific components of overhead -- whether it be for people, lights, heat, or operating systems -- it is all treated as overhead. I would also argue the same way for allocating the costs of special purpose software packages and equipment, even including the tools of our trade -- software and hardware monitors, accounting packages, simulators, etc. What I'm trying to point out is the the "problem" of system overhead is only a problem if we insist. And that if we do insist, we're worrying about an insignificant portion of total overhead.

By now I expect that everyone here disagrees with at least part of what I've said with regard to how simple charging for resource usage is, so let me try to convince you by pointing out that there are now many companies that routinely sell time on computer systems to all comers. These companies somehow survive in business by allocating all of their costs along with an appropriate share for resources used to each of their customers. Most of these companies have now found that they can even arrive at the end of each of their accounting periods with a profit. I know of none of these companies that considers implementation and processing of their charging method as a problem. It is a simple business necessity.

Furthermore, these companies must be able to present a bill for system usage immediately at the end of any user's run or session. In these environments the bill preparation must be minimized as it takes place in a period when the user who incurred the costs is finished using the system but won't "go away" to make room for another user until the bill has been prepared. This insures that, even if nothing else on the system is done efficiently, the bill preparation will be. As long as there are profit-oriented companies that are routinely using computer charging schemes successfully, I don't see how we can argue that the problems in using such schemes are too difficult for us to solve.

It seems to me that most of us exist outside the competitive "time-for-sale" environment. We are largely from the internal computer centers of the private sector (business or university) or from government computer groups (federal, state, or local, or government-funded non-profit corporations). In many of these environments there certainly are problems associated with implementing charging schemes for computer services. But these are still not technical problems. They might better be described as management decision or even political problems. The basis for these problems is the reluctance of both the "sellers" (computer center managers) and the "buyers" (computer users) of our installations. It also seems that it is the rule rather than the exception that our installations do not now charge the users for computer service.

When no charging mechanism has existed in the past it is easy to understand why users would be reluctant to start paying for computer usage. Why would anyone want to pay for something that has been free? Many arguments can be made around this question in a computer setting, but the basic economic issue that discards all answers about the benevolence of some groups of computer users is made with the old analogous question: "Who would buy a cow if milk was free?" (And I refuse to view cows as status symbols.)

The reluctance of "sellers" is not quite so easy to describe. If a charging mechanism is begun in a previously free computer system some users may elect to not use the computer, some may elect to use the computer less, and some may elect to use a different computer that is cheaper. None would elect to use more of the computer simply because it now has a charge for its usage. (The charges, of course, must be paid in real money or some other negotiable barter instrument; not in certificates that are only good at the captive data center. Charging and paying in anything but money ultimately has the same effect as not charging at all.) Knowing that buyers may elect to use less of the existing system than the "seller" normally provides, will quite naturally cause the seller to be concerned that he has more computer capacity than the group will support. The predictable result would be that the computer system may either get smaller or disappear entirely if it must operate in a free market. Hence, the prospective seller's reluctance is based on the fear that his domain may either shrink or disappear if a charging scheme is introduced in his previously "free" environment.

It would seem to be worth any reasonable cost that might be associated with implementing a charging mechanism from a performance evaluation viewpoint, simply to encourage users to remove unsuitable jobs from the system

and to force computer center management to plan their installation's capacity against the real income production of the capacity. The cost for computer charging allocation packages varies from a low of about \$2500 as a one-time leasing charge in smaller installations to a high near \$30,000 plus annual internal and external costs of as much as \$50,000 for a package that does nearly everything a software monitor will do and produces bills in the process. Most of the widely used billing packages are in the \$5000 to \$10,000 one-time charge range plus a \$500 to \$1000 annual maintenance fee with little or no internal usage costs. (See the December 1975 issue of EDP Performance Review, published by Philip Howard in Phoenix, Arizona, for a list of usable accounting packages that are commercially available.) At the relatively low prices for these packages, it seems unwise to re-invent billing packages in-house. In addition to their low prices, the processing overhead associated with use of most commercial accounting packages is on the order of 1% or less. This is truly in the "noise level" at every installation that I've come in contact with.

To summarize, I hope that I've made the point that I don't believe there are any significant problems standing in the way of implementing or processing a charging scheme at most computer installations. I also believe that all multi-user computer installations should be on a pay-as-you-go basis.

Had we really been concerned with pricing computer services here, rather than charging for computer usage, my conclusion would probably have been that we don't know enough yet about what a "service" is from a computer to effectively formulate pricing policies. But my approach would have been an economic one -- that is, the price would have to be paid in real money to allow the user to buy the service wherever the user chooses. The few examples of pricing computer services that I have seen are like the "Date-a-Dog" example at the beginning of this presentation or cases where pay checks are supplied from time cards or salary tables at some flat monthly rate per account. In these situations the price of the service shows no real connection with the cost of computer resources used to produce the service.

The interesting point in the difference between these two approaches to us as performance evaluators is that when the charges are based on resource usage, the user should be interested in installation performance. Yet if we priced computer service, the computer center management should be interested in performance. In other words, there is little economic incentive for computer management to improve performance if income is produced based on equipment usage. No matter how badly the customer uses the computer, he will pay for all resources consumed. In priced services approaches, the customer would have no incentive to use computer resources conservatively but computer management would want to sell as many services as possible. This would shift the concern about performance of the installation directly to computer center management where it reasonably belongs.

In closing, I would suggest that we have the right workshop title for a SIGMETRICS event of real importance: Pricing Computer Services. But I don't think we have the right emphasis at this convening of the workshop. Charging for use of computer resources is just not a problem anymore.