TO: Multics Administrative Distribution

FROM: F. J. Corbató

DATE: July 29, 1971

SUBJECT: A proposal for a New Hardware Maintenance Policy

INTRODUCTION

A new maintenance policy is proposed suitable for computer utilities or time-sharing systems where the system has sufficient redundancy that it is capable of operating even though not all the equipment is functional. The proposal is meant to be constructive and offer a direction for effective change since there will be an increasing number of time-sharing systems of this type in the future. What has become clear with contemporary operation of time-sharing systems is that current hardware policies are woefully inadequate and mired in the past. In fact, most of the current maintenance policies for hardware are merely a carry-over from the old punched card equipment days when there was only a feeble notion of a system and individual modules had to be up but a small number of hours per month.

The key failures which this proposal attempts to correct are:

1) Currently equipment maintenance is not viewed on a system-wide basis. Typically the outage of a single key unit such as a disk channel or a high-performance drum can cripple an entire system and affect system performance by far more then the monthly rental of the down unit.

2) There is at present no effective way for either the user or the manufacturer to place a cost on reliability. Because there is no financial penalty except for catastrophic unreliability, the manufacturer does not know whether or not to have spare units. Similarly, for the customer to complain effectively about inadequate system availability, he has only the drastic recourse of not paying rent.

3) Currently the Field Engineering (FE) organization has a principal incentive of reducing costs. From the vantage point of a customer, it would appear, no one ever judges FE performance with regard to making the system available.

4) As systems are built deliberately with redundancy, they are able to limp or operate at less than full capacity and the incentive to keep the entire system up is weakened. This has always been a minor problem when an installation has had multiple tape drives and found that one unit is always being repaired. With the advent of pools of larger units such as CPU's and memory modules, the impact is far more serious.

5) When there are system crashes due to equipment malfunction, the damage to the computing service which the customer is offering to his users is not properly reflected back to the FE's. Specifically the confusion and time spent reestablishing the system after a hardware failure, today, is only absorbed by the customer

6) The use of large hardware/software systems can be compared to

a marriage between the customer and a manufacturer where it is not feasible for every misunderstanding or dissatisfaction to be handled by the threat of divorce. Of course, the ultimate recourse for a dissatisfied customer is to switch to a competitive supplier, but this threat cannot be idly made since it literally takes years for the customer and his users to properly exploit any large-scale, sophisticated, on-line system. Moreover, the corresponding user dependence on the system requires at least one or two years to properly phase out usage without causing severe disruption of important functions. As a result reliability and system availability will play a major role in customer attitude towards the manufacturer. The manufacturer has a deep interest in doing an outstanding job of system maintenance.

THE BASIC SCHEME

The key to the proposed scheme is to make the monthly system rental paid be based on the accumulation of "system availability" units. As will be defined below, the maximum units which can be acquired is 24 per day. If on the average over a month more than a threshold of, say, 16 units per day are made available by the manufacturer, the customer pays full rental. If less than the threshold is realized, the rent paid is proportional to the ratio of realized units to the threshold value.

System availability units are computed by a straightforward formula which

takes into account the different services which are being offered as well as the shift involved and the hours of availability. Specifically the formula when applied to M. I. T. would be the following integral over time:

$$\text{SA units} = \int_{\text{day}} E \times w_{\text{shift}} \, dt$$

where

E = effectiveness of the combined Service and Development systems

w_{shift} = shift weight factor

In interpreting the above formula the notion of effectiveness refers to the usefulness that the combination of equipment would potentially have regardless of whether or not the customer chooses to use the full Service system or the Development system or chooses to partition the equipment into further subsystems. The effectiveness of the system is unity when all equipment is operational. In general, the system **effectiveness is the** product of the corresponding "limp" factors of each device or feature that is inoperative.

To illustrate these ideas more concretely, the remainder of this memo will discuss an implementation of this scheme as it might be applied to the current M. I. T. installation. The equipment is currently used by M.I.T. in two major ways either as a Service system or as a Development system (with a correspondingly crippled Service system). Both Service and Development systems are considered important service offerings. Consequently, even though not all equipment is needed for each pattern

of use, any unit which is down detracts from the combination of potential services and therefore, should lower the overall system effectiveness factor. The Service system is defined as containing:

2 CPU's

3 memory units (128K words each)

GIOC -A

Drum

1 channel x [3x (electronic unit + 5 DSU270)]

DSU170

Interlace

2 card readers

2 printers

2 punches

1 channel x 6 magnetic tape units

The Development system is defined as containing:

1 CPU

1 memory unit (128K words)

GIOC-B

Drum

1 card reader

1 printer

1 punch

1 channel x 4 magnetic tape units

A possible schedule of limp factors for the M.I.T. complex of equipment might be:

Device	Units Down				
	1	2	3	4	5
CPU	.5	0			
Memory unit	.7	.2	0		
Drum	.7	0			
GIOC-A (Service Sys. only)	.3				
GIOC-B (Dev. Sys. only)	.7				
DSU170	0				
DSU270 units	.95	.90	.85	.80	
DSU270 (elec. + 5 units)	.7	.3	0		
DSU270 channels	.8	0			
Interlace	.9				
Tape drives	.8	.6	.4	.2	0
Card readers	.7	0			
Punches	.7	0			
Printers	.7	0			

Peripherals are too drastic, time down should be taken into accn. How about 75% lines?

The above limp factors normally apply whether the particular modules either are down or are up but being used by the FE's. An exception is made, however, whenever a module is used exclusively in one particular system. If such a module is down, the limp factor still applies (so that there will be an incentive to fix the module); if the module is up but being used by FE's, the limp factor does <u>not</u> apply since giving the module back to the customer will not improve any of the services being offered. (This exception will occur most often with GIOC-B and the second drum on the Development system.)

SHIFT WEIGHT FACTOR

To account for differences in shift values, a schedule of weighting factors is arbitrarily specified:

Time	Weight Factor
0000-0800	0.5
0800-1800	1.4
1800-2400	1.0

The above schedule applies to normal week days, Monday through Friday, inclusive; weekends and holidays are given a uniform weight of 1. These weight factors are normalized so that a maximum of 24 system availability units are possible each day. (Although not specified in this example, in other environments it might be desireable to make the weight factors be relatively reduced for holidays and weekends.)

CRASHES

As indicated above, system availability units are based on the potential use of a system. Clearly crashes (and their consequences) detract from such use. In the case of the present Multics Development system, it is proposed that crashes be ignored since the damage they produce is not too severe. However, Service system crashes are of crucial concern. Whenever a Service system crash or shutdown occurs for equipment reasons, the time required to restore the system to normal operation is to be deducted from the Service system up time. (One could introduce an additional penalty for a crash, but, it will not be done in this proposal.) If the cause of the crash is unknown or disputed, then only half the restoration time is deducted. The reasoning behind the partial loss is to give the FE's incentive to determine the true cause of the crash. The analysis

of crashes is to be done by a team containing a representative from the manufacturer, the operations staff, and the software staff.

EXAMPLES

To see the effect of the above formula, several cases will be given:

a) Partial Service system in use; Development system used by FE's; first shift.

$E = (.5)(.7) = .35$, reflecting the loss of a cpu and memory module due to FE use.

Hence

$SA\ units = (.35)(1.4)\Delta t = (.49)\Delta t$

b) Full Service system in use; Development machine unavailable (GIOC-B down); second shift.

$E = .7$ Entire Service system available but GIOC-B down

Hence

$SA\ units = (.7)\Delta t$

c) Partial Service system; software checkout on Development system; 2 tape drives down, holiday.

Because the customer has the entire spectrum of services potentially available, even though not all are used, there is only the degradation from the shortage of tape drives. The two down drives give

$E = .6$

so that

$$\text{SA units} = (.6)\Delta t$$

With the above illustrations of the formula, it is informative to apply it to a couple of days of site operation. The records of Tuesday, June 29, and Wednesday, June 30, 1971, have been quickly analyzed to determine what impact the above system availability measures would indicate. This is a small sample but it does give some feel as to how the formula applies. There was considerable service disruption during these days and in addition the second drum was down the entire time so that the Development system operated at 1/3 effectiveness. The record of June 29, had the usual morning FE take-over of the system as well as three other several hour periods of FE use of the development machine. The result of this chopped up schedule, was that only a gross figure of 12.2 system availability units were accumulated. In addition, there were two crashes: one of certain hardware cause and one of questionable cause. After correction for the crashes, the net units for the day were 11.0. The principal reason this figure is low is that as mentioned above the development system drum was down the entire 24 hour period. If the drum had not been down, the net system availability units would have been 15.7 units, slightly below the full-rent threshold. (Presumably this day will average out with better ones.) Similarly, on June 30, there was much less FE utilization of the development machine with only one extra period being used. However,

the development machine drum remained down; the net figure after taking into account three hardware caused crashes and one crash of unknown cause was 12.1 system availability units. If the development machine drum had been operational, the net would have been 17.4 system availability units.

These examples indicate that the formula appears to recognize realistically system availability difficulties and resultant performance drop-offs. There is the further advantage that a single measure of merit is developed each day and is, therefore, suitable for summaries, performance reviews and analysis.

CONCLUSIONS

From the above it is concluded that the proposed mechanism should not only provide an equitable means of measuring system availability but at the same time provide a better incentive scheme for manufacturers to use in maintaining reliability. In addition various issues, such as how many spare units to have on a site are given a rational cost basis. A possible objection to the scheme is that it will introduce too much uncertainty for a manufacturer who is bringing out a new line of equipment. It would appear that the proper solution to that aspect is that the manufacturer should obtain insurance if he is either not able to insure himself or to anticipate reliability. This situation is not very different from that which an airline faces having to obtain insurance for a new line of airplanes. Finally, it should be realized that the purpose of the proposal is not to save the customer money. (Money is only used as an attention getting device.) Instead the objective is to develop a framework where everybody is working for the same goals, namely, reliable, satisfactory performance of the overall system.