


Honeywell

cc: CT Clingen
FJ Corbató
NI Morris
JH Saltzer ✓
SH Webber
TH Van Vleck

DATE: MARCH 26 1971
TO: W. J. BURNER
FROM: J. W. GINTELL
SUBJECT: SECONDARY STORAGE PLANNING FOR MULTICS



Summary of Recommendations:

A number of issues are discussed here; all of these issues require some action on your part in order to reach solutions to problem areas. The details are discussed in the body of the document; however, for convenience I list the significant action items here:

1. You should request that an HPC be moved from GIOC-B to GIOC-A so that we can run both DSU-270 channels simultaneously. We should be able to complete the software implementation for this during April; consequently the need for this change is quite immediate. This should result in a noticeable improvement in system capacity/performance.
2. You should initiate further studies on disk usage needs for the latter half of 1971 and 1972, factoring in the various pricing and administrative policies discussed here in order to determine the best set of actions. Your feedback and recommendations will allow us to begin planning and implementation of the appropriate support software.

There are three problem areas which affect disk planning for Multics: performance, capacity and administration. Although there is some interaction among these areas it is convenient to consider each individually. The following is a fairly detailed discussion of the issues leading to these recommendations.

Performance

The current Multics static multi-level scheme arranges the contents of secondary storage so that the most frequently used pages are on the fastest device. We have plans to improve this allocation facility with a dynamic page multi-level capability that will insure that the drum gets used to its extreme capacity. We also plan to improve our static multi-level to insure an optimal balance between the use of the DSU-270's and the 2314.

One of the most critical performance bottlenecks in Multics is the disk channel capacity on the DSU-270. The effect of this bottleneck is different for reads than for writes. For example, although the DSU-270 average latency should be about 50 milliseconds, we frequently observe 100 millisecond latency due to channel queuing on reads. Queued reads contribute to multi-programming idle time which prevents processes from running. This idle time ranges between 5 and 15% of the total time. Since Multics queues writes behind reads and processes don't wait for writes to complete, the direct effect of queued writes is to tie down core. We do not compute the average amount of core tied down by this but it does range between 0 and 30 pages.

The recommendation here is that we implement 2-channel 270 software as soon as possible and move an HPC from the Development machine to the Service machine. The second channel can be switched via the peripheral switch so that 270's can be used on the Development machine when needed there. We will make the second 270 channel dynamically reconfigurable.

At a later date if the 2314 channel capacity becomes a bottleneck, we should consider adding a second channel and the appropriate software.

Finally, there are a few more tricks up our sleeve which can be considered for the future.

- Make a very smart 270 dim which can start a data transfer operation in the middle of a page.
- Implement a 2314 dim which optimizes seeks.
- Put multiple copies of frequently used data on the 270's to minimize latency (while sacrificing storage space).
- Investigate the use of the second Librafile (probably impossible due to the fact that each Librafile transfers at a speed greater than one half memory speed).

Capacity

As the number of users grows and the work load of users increases the need for secondary storage increases. There are three solutions to this problem:

1. Administratively "solve" the problem.
2. Add more disk storage.
3. Add the capability of storing some segments off-line.

The administrative solutions are discussed later.

The second solution is technically simple. Your proposal is to add a second 2314. In order for Multics to use the second 2314 we need only to change the configuration cards. There is an interaction between this and performance in that as the ratio of 2314 to DSU-270 usage increases there will be a gradual slowdown. Our multi-level scheme does its best to minimize this effect.

Probably the most serious degradation will be due not to the reduction from 15 to 10 DSU-270 files but the reduction from three to two electronics units; this reduction will affect the channel capacity. If it is at all possible, I would recommend a DSU-270 configuration that includes three electronics units instead of two.

The third solution requires software development that we feel is very important to do this year. A facility could be built which will allow the system to remove certain segments from the hierarchy, placing them on tape. A record of this action would be maintained in the file system. An easy-to-use facility to retrieve these segments will be provided. This facility would allow a smaller amount of on-line storage to be used for a given user community.

Administration

A serious administrative problem is that when there is insufficient storage capacity it is difficult to allocate the resources in a manner which maximizes dollars and minimizes the hassle encountered by both the users and the administrators. Since it is desirable to sell all the disk space available, and since usage grows in small increments while storage capacity grows in large increments, this administrative problem will probably always exist.

In order to lessen this problem, some or all of the following tasks could be adopted:

1. Adjust the charging policy to force users to use less disk space because it is very expensive.
2. Institute a policy to charge for unused quota to induce users to request only what they need at the moment.
3. Give the users a cost reduction if they specify segments to be stored off-line even if the system chooses to store them on-line. This will tend to create a large amount of space eligible for removal by the system and thus be usable by other users.

4. Reorganize the project directories so that the System Administrator will need to interface with a smaller number of project administrators. (e.g. all of the Project MAC projects could have sub-directories under a single MAC directory. The individual Project MAC leaders would all interface with a single Project MAC person who would interface with the System Administrator).
5. Oversell disk quota based on statistics obtained while running the system. If the system crashes because it runs out of space, it is to be considered a bug which we will rapidly repair.

Summary

A number of steps have been recommended in this document. One requires immediate action; the others require some additional study prior to implementation. The list of steps is summarized below:

1. Move an HPC from GIOC-B to GIOC-A and implement 2-channel DSU-270 software.
2. Arrange for another 2314 if the additional space will be needed during 1972.
3. Implement the off-line storage facility.
4. Implement all of the suggested administrative procedures and policies.
 - adjust the charging policy
 - charge for unused quota.
 - give a price reduction for off-line storage.
 - reorganize the project directories.
 - oversell quota.

I suggest that we arrange a meeting with the next several weeks to discuss these issues so that we can get to work on the development as soon as possible.


JOHN W. GINTELL

3-26-71

Wes:

This reference might be useful

John

