

MEMORANDUM

To: Joan Tierney

From: David Wozny

Date: 23.3.96

Joan,

I would like to confirm in writing my dissatisfaction that the IT manager has decided not to consider my revised proposal for overcoming the server storage space problems at LG. I set forth my reasons below.

Original Proposal

The original proposal that I put forward consisted of a hierarchical storage management system (HSM). Utilising the 'on-line' disk space on our servers (approx. 3GB), a 'near-line' optical jukebox providing 19GB capacity for data meeting specified migration criteria and an 'off-line' storage facility comprising of an auto-loader DAT drive to provide disaster recovery capability. This proposal WOULD solve the existing storage management problems that have at LG. However, it is a very 'technically innovative' solution involving a significant amount of ground breaking technology, such as the Octopus replication tool functioning with phantom files, using SCSI switches to enable changeover of the jukebox from one server to the other in the event of a disaster, use of the HSM software on NT which is very new and is currently the only software of this type on the market and installation of HSM agents on PC clients which could cause memory conflicts.

None of these problems are I believe insurmountable, however, there are 2 limitations that will always exist:

1. The on-line space available is limited to 3GB (this includes software) and I feel that we would need to be very stringent in our migration criteria in order to continually observe this on-line limit. The amount of on-line data space will reduce as the LAN evolves, software which is yet to be server installed such as Visual Basic, ASR, AutoSketch, etc. will eat away at the available on-line space. This software space problem is at least as significant as the user data storage problem.

Note: HSM is often quoted as a solution to the '80/20 problem' of file access, an HSM implementation at LG of roughly 2GB on-line data to 20GB of near-line data would constitute a 90/10 problem, which is an unlikely scenario.

2. The jukebox disks are not RAID controlled and hence disk failures in the jukebox would require rebuilds.

Revised Proposal

Subsequent to the original proposal being agreed and the necessary purchase orders raised, I became aware that high-capacity SCSI disks have reduced in price massively since I last considered them (at the time the servers were originally specified).

With the availability of inexpensive high-capacity SCSI disk, I now believe that replacing the existing 1GB disks with 4.3GB disks is a better way of addressing server space problems on the servers than installing optical jukeboxes in a HSM architecture. The benefits are thus:

1. The replication technology is proven;
2. The disks will be RAID 5 controlled;
3. Installation will be swift;
4. We understand the technology fully;
5. Access to disk will be rapid, faster than current access!
6. There are no software overheads;
7. Disaster recovery plans / testing will be the same as currently used;
8. Extra disk will enable more flexibility in server configuration;
9. Re-architecting the servers to allow for a mirrored operating system will be much easier to achieve.

Comparison

Although for a similar cost, only 10.6GB data will initially be available compared to 20+GB with a HSM, all of that space will be fully usable i.e. not constrained by a 3GB bottleneck. 10.6GB should be enough capacity for at least another 12 months, considering that 3.0GB has been adequate for 2 years' worth of growth (and the initial loading of up to 5 years of historical data). In the event that extra space should be required, adding extra space is simply a case of adding an extra disk in each server at a cost of about £2,000.

Conclusion

I believe that at this moment in time, choosing to go forward with the HSM proposal would be the wrong decision, bearing in mind that a better solution exists. It is an unproven technology on Windows NT and there could be many problems in fully installing it, increasing disk space is a relatively simple task which will provide a far more flexible system, with better resilience and a more rapid access time.

David M. Wozny