

X3T9.2/88-97

ENDL

August 2, 1988

Mr. John Lohmeyer
NCR
MS 21
3718 N. Rock Road
Wichita
KN 67226

Dear John,

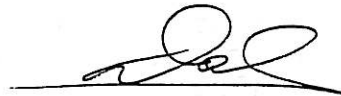
I have been giving a lot of thought to SPASTIC, and have come to the conclusion that we may be spending far more committee time and engineering effort on this problem than is justified.

We have accepted the requirement to do electronic address assignment at face value. Attached is an extract from the latest issue of the ENDL Letter which contains some of my thoughts on the matter.

I am guilty of having liked the idea before Gerry brought it into committee, because it is a capability I always thought would be useful. After wrestling with the actual configuration problems that it creates I am less than enamored with the concept.

It does not deserve a separate agenda item, but could we spend a little time at the next meeting in discussion on whether we really think we need to add this capability to SCSI?

Yours sincerely,



I. Dal Allan

ENDL Letter

July 1988 Happenings

Copyright © 1988 ENDL Incorporated. Reproduction without permission is prohibited.

Is it Needed?

The most disturbing thing about this whole exercise is that the arguments in favor of SPASM may be specious. Gerry's justification was based on IBM mounting its disks on skids which slip into place in the chassis. Reaching jumpers or dipswitches to set an address is impossible because sheet metal for the skid surrounds the disk, and obscures much of the PCB.

At Boise, Gerry stated that internal devices do not need the additional 10-pin bus, as the addresses are set at the factory. This means that IBM now has exactly the same configuration as Apple, with external attachment boxes needing address assignments. SCSI box manufacturers have already solved the problem of harnessing the address lines from a thumbwheel. The Mac world seems to be doing fine with thumbwheels. Surely IBM could buy a few to see how other companies have solved the address assignment problem.

If IBM really wanted to be innovative, and solve the problem in a friendly way it could be done with cables. How about building a special connector for a bulkhead with some extra pins? A blue connector is address 0, a red one is 1 and so on.

All a user would have to do to make sure a system is configured correctly is verify that the connector plugging into each attachment box was a different color. What a fashion opportunity - color-coordinated cables and connectors (Red-White-and-Blue for patriots). As the extra pins are not conductors in the cable, they will not affect the internal harnessing.

The other justification for SPASM is dynamic address re-assignment to give priority to the high performance devices. It is questionable if this is needed on a PS/2 but if it is, then a simpler solution is to use synonyms. IPI-3 included synonyms for software, so the framework exists and a simple adaptation to the SCSI command format is all that is needed.

Gotcha, you might think. What if all the addresses are assigned? No problem. The initiator can assign its address as a synonym to a target to free up one address and after assigning all the other synonyms, give the target which has the initiator's address the last synonym. NOTE: It is a little known quirk that SCSI works quite well with duplicate addresses as long as one is the initiator.

Appealing though dynamic address assignment may be, it is a messy problem that offers no neat and tidy solution. One has to question the cost:benefit ratio of going through the exercise. The problem has been with SCSI since Day 1 and it has done nothing to slow down industry growth. Users have had no problem accepting the need for devices to have different addresses set so that the interface will work.

The arguments on user-ignorance and user-friendliness are empty rhetoric unless one believes PS/2 purchasers are innately less competent and more simple-minded than Mac users (what a line for Apple).