Document Number: T10/98-157r0
Date: May 4, 1998
To: T10 Members
From: James McGrath, Quantum Corporation
Subject: Testing for Old Expanders

Background

Recent conversations with customers has identified a potential problem with any modification of the SCSI protocol (double edge clocking, CRC, packetized). Expanders in current systems that are invisible to the initiator and target but could, through lack of knowledge about changes in the SCSI protocol, prevent new SCSI features from operating.

There is no way to make an old expander "new protocol aware." And as long as expanders are implemented using some pretty strange techniques to allow them to work invisibly on the SCSI bus, any new protocol changes can introduce a very difficult to diagnose system failure on systems using old expanders.

This proposal is the first of three from Quantum addressing the problem of expanders and protocol changes. This one deals with the basic issue of detecting old expanders. Another introduces the ability for targets and initiators to give expanders some simple information to allow them to work reliably. The final introduces a technique to allow devices to get the equivalent of inquiry information from expanders.

Throughout we are assuming that a design constraint on expanders is that they are simple signal repeaters, with a minimal degree of intelligence. Obviously SCSI bridges could solve the problem of new expanders adapting to changes in the SCSI protocol, but discussion with vendors has indicated a need/desire for less expansive and simple expanders as well as SCSI bridges.

Proposal Development Process

Quantum has developed these proposals in consultation with some companies currently selling or using expanders – Adaptec, Compaq, Dec, Symbios Logic. Their consultation should not be taken to imply agreement, although I think there is a good degree of agreement on the general issues and their resolution if not in the details of the implementation.

Pending feedback on this proposal, Quantum will submit a revision that specifically addressing the changes needed in the appropriate draft standards.

General Solution

The general solution is really quite simple – after negotiating a change in protocol, the devices perform an act that should lead to a failure if an old expander is present. Inducing the failure as part of an "extended" negotiation process allows the devices to back off and renegotiate to a protocol that, eventually, will not fail. These failures could be reported to the user with some sort of sense code/key – that is not included in this proposal at this time for simplicity, but could be added in the next revision if the group desired it.

This proposal will focus on detecting old expanders that might fail under a double edge clocking/CRC protocol.

Implementation Issues – Host/Command approach

There are two extreme methods of implementation that have been proposed – we would appreciate the group's feedback so that a better proposal can be drafted. Quantum would desire a single proposal so as to simplify design and testing.

One technique is simple in concept, but operates at a very high system level. Basically, the host performs a WRITE (best would be WRITE BUFFER) command using the existing, SPI-2 compatible legacy protocol. It then negotiations to a new, SPI-N level protocol (e.g. double edge clocking/CRC in data phase) and issues a READ (best would be a READ BUFFER) command.

A short command timeout is used to catch and old expanders that misbehave so badly that they hang the bus (we do not think any recent expander will behave that badly). Otherwise a data comparison is done. Note that in particular an expander that does not understand double edge clocking will block every other word from the sender to the receiver. Thus a data pattern of 1,2,3,4,5,6 will be read back as 1,1,3,3,5,5. It will also generate a CRC error.

While this approach has the benefit of requiring no protocol changes (it uses existing commands), it also operates at a very high protocol level (commands) while the problem is a physical/protocol one. Specifically, the implementation must insure that no intervening SCSI activity takes place between the steps of the test, which may be difficult to do in some environments. The implementation also places the burden on testing on the initiator, with the target playing a passive role.

Implementation Issues – Target/Protocol approach

A second implementation does the exact opposite – it utilizes new protocol that could take place totally at the hardware level, and focuses of using the resources of the target. This approach minimizes the impact on the host, while the first maximizes the impact.

After the last MSG exchange of the new the target goes into the DATA IN phase. The target transfers a series of test patterns to the initiator, who then "echos" them back to the target in a follow-up DATA OUT phase.

This tests the transceivers both transmitting and receiving data, but places the burden of intelligence for this process on the target, not the initiator. This matches the asymmetric resources of initiator chips and target devices by placing an asymmetric burden heaviest on the device with the most resources

Under no circumstances should any of this test data be passed on to the higher level protocol. Indeed, as far as the host is concerned this data transfer never took place. The devices always assume that a data transfer after the successful negotiation of new protocol is for testing the environment for things that might make the negotiation invalid (like old expanders). Once the target changes to COMMAND or STATUS phase, then all subsequent data phases during that connection should be treated as normal (i.e. host accessible) data transfers.

Note that if the next logical phase for the target would have been a data phase (unusual, but possible), the target simply disconnects and later reconnects. This introduces a performance loss, but only during protocol negotiations (which occurs very seldomly). If the target was not allowed to disconnect (this is legal, but negotiating after COMMAND phase for new protocols on a data transfer command and not allowing disconnects has to be rare), then the target simply goes into STATUS phase and indicated an error.

Conclusion

Old expanders can de detected, and should not pose an obstacle to further evolving the SCSI protocol. But the details of detection should be developed with attention paid as to where in the protocol stack it resides and which device should carry the burden on implementation.