Document Number: T10/98-158r0
Date: May 4, 1998
To: T10 Members
From: James McGrath, Quantum Corporation
Subject: Communicating with New 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 second of three from Quantum addressing the problem of expanders and protocol changes. The first dealt with the basic issue of detecting old expanders. This one 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, and 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 – during the establishment of any nexus, the devices exchange a message that even a simple expander can snoop and use to configure itself for subsequent operation. Note that while the information is packaged in a message, the ideal implementation of this feature is as a change in the low level signaling (e.g. toggling a signal on the SCSI bus). A message is the closest implementation we could find that would introduce the least amount of backward compatibility issues.

This proposal will focus on messages for old expanders that need to know about double edge clocking/CRC and higher speeds.

Proposal

Before the identify message is sent at the start of the nexus/connection (e.g. BSY asserted, MSG IN or MSG OUT phase, byte is the first byte of the message phase), a message byte with bit 7 = 0 and bit 6 = 1 is sent. This message has bits 4 and 5 set to 0, and the lower nibble set as follows:

Bit	7	6	5	4	3	2	1	0	
Value	0	1	0	0	X	0	0	0	Asynchronous
Value	0	1	0	0	0	X	X	X	Single Edge Clocking
Value	0	1	0	0	1	X	X	X	Double Edge Clocking

For bits 0-2, the maximum safe reception speed (in MT/s) corresponding to the value is as follows:

Bit	7	6	5	4	3	2	1	0	MT/s
Value	0	1	0	0	X	0	0	1	5
Value	0	1	0	0	X	0	1	0	10
Value	0	1	0	0	X	0	1	1	20
Value	0	1	0	0	X	1	0	0	40
Value	0	1	0	0	X	1	0	1	80
Value	0	1	0	0	X	1	1	0	160
Value	0	1	0	0	X	1	1	1	320

This byte is optional, but is intended to be easily snooped by expanders, and so should always be sent on every connection unless it is know by a higher level means that no expanders exist in the data path. Note that sending it every connection greatly simplifies

the job for the expander designer.

The expander should repeat this message (like it would any other signal transitions).

The device at the other end of the connection may compare this message to the expected transfer parameters and, it different, end the command with a CHECK CONDITION status. It is recommended that these devices captures and store the byte for diagnostic purposes.

Extensions (e.g. Packetized SCSI)

This proposal is designed to work in a non-packetized SCSI environment, but it could also be made to work in a packetized environment. The rule would be the same – that the first information on the data lines sent during a nexus is the expander message. But now this might occur in a packet in or out phase as well as a message in or out phase.

Indeed, the whole point of this proposal is to allow expanders to adapt to a wide variety of future protocol conditions. Extensions are obvious – the expander should be designed to continue to detect a sequence of two byte messages (each message having a different, but well defined first byte), each of which could address an expander capability.

Two byte messages were chosen rather than extended messages since they minimize protocol overhead.

The devices use of these messages should be controlled via a mode page setting or as part of the protocol negotiation process. This allows the messages to be avoided in those cases where they are not needed, reducing protocol overhead.

Conclusion

New expanders should have this capability added in order to allow them to “keep up with protocol changes” and avoid future implementations that can easily fail when new protocols are introduced. Note that seeing an unknown two-byte message gives the expander the information it needs to, for instance, deliberately induce a detectable failure if it desires. Naturally if it can perform the adjustments required by the message, then it can do so.