

To: T10 Membership
From: Gene Milligan, Seagate Technology <gene_milligan@notes.seagate.com>
 George Penokie, IBM <gop@ibm.us.com>
 Lawrence J. Lamers, Adaptec, Inc. <ljlamers@ieee.org>
Subject: Packet Protocol Extensions
Date: Thursday, March 04, 1999

This proposal presents extensions to the existing packet protocol that address concerns about the granularity of the CRC interval and the bus turnaround effect on write operations. The extensions work as a compatible extension of the existing packet protocol.

CRC Interval

The proposal is to add a CRC interval to the L_Q Information Unit. This extension defines use of CRC intervals in the associated Data Information Unit. The CRC interval specifies the number of bytes transferred before a 4-byte CRC is transferred. This allows the granularity of the CRC interval to be less than a Data Information Unit size, (e.g., a sector size for disk drives), so that reasonable Data Information Unit sizes (e.g., 2048, 4096) can be preserved. This also keeps the overhead of L_Q Information Unit in line with expectations.

A Data Information Unit contains zero or more CRC intervals and may contain a runt CRC interval if specified by the data length. This allows the data transfer to be validated in smaller chunks so that it can be moved to memory or media without buffering the whole Data Information Unit. A CRC interval of zero indicates that the CRC occurs at the end of the Data Information Unit. The last four bytes of a packet shall be CRC bytes. The CRC Interval shall only be valid for Data Information Unit transfers.

The proposal adds a 4-byte extension to the L_Q information unit, increasing the size to 24 bytes. Two bytes are used for CRC Interval and two bytes are reserved. The CRC Interval field contains the length in 2-byte words of the number of bytes that are sent before a CRC is transferred including the pad bytes if any.

In addition a CIPad bit is added to byte 12 of the L_Q information unit to indicate whether the interval has zero or two bytes of padding.

The Padbytes field applies to the CRC at the end of the Data Information Unit to allow for odd byte length transfers.

Note: The Data Length description needs to be modified to exclude all CRC bytes and include all pad bytes.

Bus turn around on write operations

Streaming write operations allow a target to stream Data Information Units on a write operation to eliminate the bus turnaround at the end of each Data Information Unit by allowing a single L_Q Information Unit to transfer multiple Data Information Units. A new Type is added to Table 39 for Streaming Data. (Code value at editor's discretion.)

All the Data Information Units sent during a single stream shall be of the same size. The stream shall only be stopped on Data Information Unit boundaries by a phase change.

The Data Information Unit stream ends when the initiator detects a phase change. If an initiator detects a REQ after the last transfer of a Data Information Unit and the DATA OUT phase continues to exist the next Data Information Unit shall be transferred.

If the initiator is not going to transfer the next Data Information Unit it requests a disconnect by asserting the ATN at least a negotiated offset value number of words prior to the end of a Data Information Unit and continues to transfer the current Data Information Unit until complete. In the event of a Data Information Unit size smaller than the negotiated offset value the target is allowed to continue the stream operation across two Data Information Units but not three. A one Data Information Unit stream is allowed.

The offset count is not required to be zero at the end of a Data Information Unit if a subsequent Data Information Unit follows for the same I_T_L_Q nexus.

Streaming read operations

Streaming read operations operate in the same manner as streaming write operations except that they are required to change to a MESSAGE IN, MESSAGE OUT (if ATN is asserted) or BUS FREE phase prior to sending the Status Information Unit.