Document No. X3T9.2/91-013 R03

Date : November 18, 1992

To : John Lohmeyer, Chairman, X3T9.2

From : Gary R. Stephens
IBM Corporation
D64A/040-1
9000 S. Rita Rd.
Tucson, Az 85719
(602) 799-2246
INTERNET: GSTEPHEN@TUCVM2.VNET.IBM.COM

Subject: SCSI-3 Generic Packetized Protocol (GPP)

This document is the third draft of the SCSI-3 Generic Packetized Protocol (GPP). SCSI-3 Generic Packetized Protocol is a SCSI logical protocol suitable for operation in a Fiber Channel environment as well as in other environments.

This document functions as a Fiber Channel FC-4 document mapping SCSI-3 I/O Processes to the Generic Packetized Protocol and then mapping GPP to the Fiber Channel. Two normative annexes describe information packet transfers using the SCSI-3 Parallel Interface.

Section 3 contains a glossary. Terminology is used consistently in all sections based on the terms. All terms used beyond Section 3 are defined in this glossary, or at first use. The glossary contains terms from Fiber Channel. These terms do not alter or replace Fiber Channel definitions. Rather, they add a SCSI sense to the definitions.

Section 4 contains information on the construction of information packets as required by GPP.

Section 5 contains information on I/O process management as required by GPP.

Section 6 has the Fiber Channel options selection menu. A brief description of Fiber Channel options, where appropriate to SCSI-3 Generic Packetized Protocol, is included. All requirements of Fiber Channel reside in the appropriate standards. Since Fiber Channel is aimed at several audiences it contains options not required in a SCSI-3 GPP environment.

Section 7 specifies how SCSI-3 Generic Packetized Protocol uses the services of Fiber Channel to develop its exchange protocol. SCSI-3 Generic Packetized Protocol is a user of Fiber Channel and is not required to, nor is it wise to, use all functions available in Fiber Channel. To maintain its interface independence, GPP uses a minimum of FC-PH functions which are not generally available in other interfaces. The exchange protocol is used to transport information packets over the physical serial medium of Fiber Channel.

Annex A specifies the mapping of SCSI-3 Generic Packetized Protocol onto SCSI-3 Parallel Interface.

Annex B specifies the information packet transfer protocol used with the SCSI-3 Parallel Interface.

Gary R. Stephens

Intentionally Left Blank

**NOTE: Copies of this document may be purchased from:
Global Engineering Documents, 2805 McGaw, Irvine, CA 92714
(800) 854-7179 or (714) 261-1455.**

**Working draft proposed
American National Standard
for Information Systems -**

**SCSI-3 Generic Pacietized Protocol
(GPP)**

November 18, 1992

Secretariat

Computer and Business Equipment Manufacturers Association

ABSTRACT: This standard defines functional requirements for the SCSI-3 Generic Packetized Protocol (GPP). GPP permits computers to attach to each other and to intelligent devices through a variety of physical interfaces.

GPP uses information packets to transfer commands, command parameter data, command response data, logical data, messages, status, and autosense. GPP uses the Fibre Channel as its primary transport layer. The mapping is specified in this standard. GPP uses the SCSI-3 Parallel Interface (SPI) as an alternative parallel transport layer which is specified in the appropriate annexes.

GPP establishes one logical level of interpretation for information packets independent of the physical transport layer. The parallel interface simulates the serial protocol which simplifies SCSI-3 Parallel Interface management compared to the SCSI-3 Interlocked Protocol.

All information transfers in GPP are independent of the physical bus used. GPP defines a method for managing multiple paths between computers and attached devices. See the SCSI-3 Architecture Model standard and the associated device class command set standards.

This is an internal working document of X3T9.2, a Task Group of Accredited Standards Committee X3. As such, this is not a complete standard. The contents are actively being modified by the X3T9.2 Task Group. This document is made available for review and comment only.

POINTS OF CONTACT:

X3T9.2 Chair

John B. Lohmeyer (X3T9.2 Chair)
NCR Corporation
3718 N. Rock Road
Wichita, KS 67226
Tel: (316) 636-8703
Fax: (316) 636-8889

Internet:
john.lohmeyer@wichitaks.ncr.com

X3T9.2 Vice-Chair

I. Dal Allan (X3T9.2 Vice-Chair)
ENDL
14426 Black Walnut Court
Saratoga, CA 95070
Tel: (408) 867-6630
Fax: (408) 867-2115

Internet:
2501752@mcimail.com

//

STATUS of EDITING:

This is the third draft of the proposed standard. This document is the result of a project proposal to update SCSI for serial interface operation.

This document has not been approved by the X3T9.2 committee and has not been reviewed by them as of November 18, 1992.

Technical Editor

Gary R. Stephens
IBM Corporation
9000 S. Rita Rd
Tucson, Az 85744
Tel: (602) 799-2246
Fax: (602) 799-6487

Internet:
gstephen@tucvm2.vnet.ibm.com

Foreword (This foreword is not part of American National Standard X3.xxx-199x.)

The Generic Packetized Protocol provides a general information transfer format between SCSI devices to exchange information using a structure called an information packet. The command sets, data, status, and autosense information are described in the SCSI-3 Architecture Model and one or more SCSI-3 command set standards. These data types are placed in a logical order, relative to an I/O process, in information packets to be sent to a target controller or to be sent to an initiating controller. Certain data types flow only in one direction, relative to the initiator and target SCSI devices, for an I/O process.

The Generic Packetized Protocol is capable of being transported over any physical medium which provides delivery of ordered sequences of bytes. In particular, this standard specifies the appropriate transport layer protocols when sending information packets using the Fibre Channel standard (FC-PH), the SCSI-3 Parallel Interface standard (SPI), and Ethernet.

With any technical document there may arise questions of in interpretation as new products are implemented. The X3 committee has established procedures to issue technical opinions concerning the standards developed by the X3 organization. These procedures may result in SCSI Technical Infomation Bulletins being published by X3.

These bulletins, while reflecting the opinion of the Technical Committee which developed the standard, are intended solely as supplementary information to users of the standard. This standard, ANS X3.xxx-199x, as approved through the publication and voting procedures of the American National Standards Institute, is not altered by these bulletins. Any subsequent revision to this standard may or may not reflect the contents of these Technical Information Bulletins.

Current X3 practice is to make Technical Information Bulletins available through:

Global Engineering documents
2805 McGaw
Irvine, CA 92714
(800) 854-7179
(714) 261-1455

Process Approval Statement

Chair, Vice-Chair, Secretary list

Patent Notice

Organization, Principal/Alternate List X3

Organization, Principal/Alternate List X3T9

Organization, Principal/Alternate List X3T9.2

Editor List

Intentionally Left Blank

Introduction

This American National Standard defines the protocol for communicating between interconnecting computers and peripheral devices using the SCSI-3 Generic Packetized Protocol (GPP). "------ ----- ----- ----- -----" on page --- shows the relationship of this document to other SCSI-3 standards.

```
                     SCSI CAM-2 Common Access Method (CAM)
                                   |
      +-------------+--------+--------+----------------+
      |             |        |        |                |
  SCSI-3 Common  SCSI-3 Block  SCSI-3 Sequential    SCSI-3 Other
   Command Set   Command Sets    Command Sets       Command Sets
     (SCC)          (SBC)           (SSC)                |
      |             |                 |                  |
      +-------------+--------+--------+----------------+
                             |
                 SCSI-3 Architecture Model (SAM)
                             |
      +----------+----------+----------+-------------+
      |          |          |          |             |
 SCSI-3 Interlocked |  SCSI-3 Fiber Channel | SCSI-3 Generic Packetized
     Protocol      |     Phase Protocol    |         Protocol
      (SIP)        |         (FCP)         |           (GPP)
       |           |           |           |             |
 SCSI-3 Parallel   |     Fiber Channel     |       Fiber Channel
    Interface      |       (FC-PH)         |            or
      (SPI)        |                       |      SCSI-3 Parallel
                   |                       |         Interface
        SCSI-3 Storage System     SCSI-3 IEEE 1394      or
        Architecture Protocol         Protocol        Ethernet
              (SSA-1)                  (SBP)             or
                 |                       |              Other
        SCSI-3 Storage System         IEEE 1394       Interfaces
        Architecture Interface
              (SSA-0)
```

Figure 1. SCSI-3 Document Roadmap

The roadmap is intended to show the general applicability of documents to one another, not a hierarchy, protocol stack, or system architecture. For example:

- CAM and SAM and the SCSI-3 command set standards are applicable to all link protocols.
- SIP, SSA, FCP, and SBP are link specific protocols designed to be applied to the physical interfaces directly below them.
- GPP is intended to be used with any physical interface.

GPP features transport layer independence and it uses a logical interpretation of and response to information packets. This standard provides mappings for GPP to:

- Fiber Channel (FC-PH).
- the SCSI-3 Parallel Interface (SPI).
- IEEE 802.2 (ETHERNET)

Mapping to other physical transport layers is outside the scope of this standard.

The SCSI-3 standards provide computers with device type independence within a class of devices. GPP provides physical transport layer independence.

GPP includes references to the necessary standards and specifies options to allow SCSI inter-operability of devices meeting this standard. GPP assures SCSI inter-operability only between two GPP devices having compatible physical transport layer characteristics.

GPP provides a protocol for the movement of commands, data, and responses between devices encapsulated in information packets. The SCSI-3 Architecture Model standard is a prerequisite standard for this standard. GPP permits transfers of information packets on the interfaces identified above. GPP supports buffer-to-buffer transfers of information packets. Logical information is not buried in the transport protocol. That is, all information required to manage SCSI-3 devices is carried in the information packets. Each information packet is self-identifying which permits correct routing of information packets for each I/O process within SCSI-3 devices.

This standard provides a point of incompatibility with the SCSI-3 Interlocked Protocol (SIP) and with serial transfer protocols which emulate the SCSI-3 Interlocked Protocol (SIP). Full duplex interface protocols are, by definition, incompatible with the interlocked nature of the SIP protocol.

Contents

Figures

Intentionally Left Blank

Tables

Intentionally Left Blank

1.0 Scope

The SCSI-3 Generic Packetized Protocol (GPP) provides for connecting multiple initiating controllers and multiple target controllers. GPP also provides control mechanisms to use multiple paths between two devices during an I/O process when there are two or more paths in a logical system.

On the SCSI-3 Parallel Interface (SPI), properly conforming SCSI-3 Interface Protocol (SIP) devices reject all GPP protocol extensions. GPP devices on a parallel interface permit such rejections and allow SIP devices to continue inter-operation with each other without requiring use of the extensions. GPP and SIP devices can share the same parallel interface.

Section 2 provides a reference list of corequisite and reference standards.

Section 3 provides a comprehensive glossary of terms used in GPP.

Section 4 describes the method for encapsulating logical functions in information packets.

Section 5 defines messages and protocols for their use to support the information packet structure.

Section 6 identifies the options in Fiber Channel selected as requirements for GPP on Fiber Channel.

Section 7 specifies Fiber Channel specific exchange protocols needed to support information packet transfer.

Annex A and Annex B describe the SCSI-3 Parallel Interface implementation. Although SCSI-3 Generic Packetized Protocol and SCSI-3 Interlocked Protocol devices may coexist on the same SCSI-3 Parallel interface, these devices cannot operate with each other unless the SCSI-3 Generic Packetized Protocol device changes its definition to SIP operating mode. SCSI-3 Generic Packetized Protocol devices, both initiating controllers and target controllers, attached to parallel busses, have the option of reverting to the SCSI-3 Interlocked Protocol operating mode as necessary.

A change in definition of a GPP device to the SIP operating mode means the SCSI-3 Interlocked Protocol and the SPI standard define all subsequent operations between the two SCSI devices. Refer to the SCSI-3 Interlocked Protocol standard for the SCSI-3 Interlocked Protocol operating mode.

2.0 Normative References

The following corequisite standards contain provisions which, through reference in GPP, constitute provisions of GPP. At the time of publication, the editions indicated were valid. All standsards are subject to revision, and parties to agreements based on this standard are encouraged to investigate the possibility of applying the most recent editions of the following list of standards. Members of IEC and ISO maintain registers of currently valid International Standards. ANSI performs a similar function for American National Standards.

2.1 Corequisite ANSI Standards

ANSI X3.4-1977, American Standard Code for Information Interchange (ASCII),

ANSI X3.xxx-199x, SCSI-3 Architecture Model (SAM).

ANSI X3.xxx-199x, SCSI-3 Parallel Interface (SPI).

ANSI X3.xxx-199x, SCSI-3 Interlocked Protocol (SIP).

ANSI X3.xxx-199x, Fibre Channel - Physical and Signaling Interface (FC-PH).

2.2 Corequisite ISO Standards

ISO 8802.2, ... (Ethernet)

2.3 Reference ANSI Standards

ANSI X3.xxx-199x, SCSI-2 Common Access Medthod (CAM-2).

ANSI X3.xxx-199x, SCSI-3 Common Command Sets (SCC),

ANSI X3.xxx-199x, SCSI-3 Block Command Sets (SBC),

ANSI X3.xxx-199x, SCSI-3 Sequential Command Sets (SSC).

ANSI X3.xxx-199x, SCSI-3 Other Command Sets....

ANSI X3.xxx-199x, Fibre Channel - Fabric Requirements (FC-FG).

ANSI X3.xxx-199x, Fibre Channel - Low Cost Topologies (FC-LT).

ANSI X3.xxx-199x, Fibre Channel - Cross Point Switch Topology (FC-XS).

3.0 Definitions and Conventions

3.1 Definitions

This section contains a glossary of special terms used in this standard. These terms apply to mapping information packet transfers to the Fiber Channel and the terms do not constitute a comprehensive glossary for SCSI-3. See an additional glossary in section "3.2 Definitions of Essential Fiber Channel Terms" on page 6 for selected Fiber Channel terms.

A separate glossary, similar to Section 3.1., is available for terms which apply to mapping information packet transfers on the SCSI-3 Parallel Interface is found in Annex A.

B

3.1.1 bit. See "name" bit.

3.1.2 burst. A portion of an interface logical element which is split to fit into an information packet.

3.1.3 byte. An 8-bit construct.

3.1.4 byte string. A contiguous set of bytes.

C

3.1.5 connect. The initiating controller process which transfers an information packet between an initiator and a target. The result establishes a nexus and begins an I/O process in a target controller. A connect requires a successful information packet transfer associated with the nexus.

F

3.1.6 field. A set of one or more contiguous bits.

G

3.1.7 GPP. SCSI-3 Generic Packetized Protocol.

3.1.8 GPP device. A SCSI device that uses the Generic Packetized Protocol.

3.1.9 GPP address. The unique address for one port on an GPP device relative to the GPP interface to which it is attached. The GPP address is unique on a GPP interface. Multiple ports on the same GPP device connected to the same GPP interface have different GPP addresses. The form of the GPP address varies with the physical interface used.

3.1.10 GPP interface. The set of all GPP devices to which a single GPP port can communicate and the interconnecting transmission media.

3.1.11. A device capable of attaching to one or more GPP interfaces via one or more ports, respectively. The device may allow a port to function in either initiator mode or target mode for any one connection. Each GPP port is capable of operating in either mode.

3.1.12 GPP port. Port.

H

3.1.13 H_C nexus. A nexus which exists between an initiating controller and a target controller. An H_C nexus, or any of its derivatives, may use more than one initiator and one target combination for connections during an active I/O process when using the multiple path option.

3.1.14 H_C_L nexus. A nexus which exists between an initiating controller, a target controller, and a logical unit. This relationship replaces the prior H_C_x nexus.

3.1.15 H_C_L_Q nexus. A nexus between an initiating controller, a target controller, a logical unit, and a queue tag. This relationship replaces the prior H_C_x_y nexus.

3.1.16 H_C_x nexus. A nexus which is an H_C_L nexus. This relationship replaces the prior H_C_x_y nexus.

3.1.17 H_C_x_y nexus. A nexus which is either an H_C_x or H_C_L_Q nexus. This relationship replaces the prior H_C nexus.

3.1.18 H_I_T_C nexus. An H_C nexus, or one of its derivatives, between an initiating controller and a target controller, which operates in single path mode.

I

3.1.19 I/O process. An I/O process consists of one initial connection and zero or more reconnections with a minimum of one information packet transfer per connection. The connection(s) pertain to a nexus in which at least one information packet is successfully transferred.

3.1.20 identify function. A process performed by logical elements to establish or reestablish a nexus. The result of the process is a correctly formed set of information packet prefix control fields.

3.1.21 information packet. A set of bytes containing Interface Control Fields and at least one Interface Logical Element.

3.1.22 initiating controller. A logical element which normally starts I/O processes. I/O processes execute using the services of one or more ports in the initiating controller.

3.1.23 initiating controller ID. An initiating controller ID is an identifier which uniquely names an initiating controller. The identifier is communicated to target controllers over the interface as part of the multiple path mode option.

3.1.24 initiator. A port operating in initiator role through which an initiating controller starts and executes I/O processes. An initiator usually attaches to an initiating controller. An initiator attaches to exactly one interface.

3.1.25 initiator role. The operating mode of a port which permits it to initiate I/O processes.

3.1.26 Interface Control Fields. An organized collection of bytes used to encapsulate Interface Logical Elements into information packets. GPP interprets some of these fields to accomplish proper routing of information packets between initiating controllers and target controllers.

3.1.27 Interface Logical Element. An identified structure used to communicate the logical content of SCSI using the Generic Packetized Protocol.

3.1.28 invalid. An illegal, reserved, or unsupported bit, field, code value, or protocol sequence.

L

3.1.29 logical element. An initiating controller or a target controller.

3.1.30 logical system. From the point of view of a single logical element, the set of SCSI devices reachable through at least one port.

N

| 3.1.31 **name bit.** A field containing only one bit. "name" is replaced with a uniquely identifying phrase.

| 3.1.32 **nexus.** A relationship between an initiating controller and a target controller which begins with a connect and ends with the completion of the associated I/O process. The relationship may be restricted to a single logical unit by the identify function.

O

3.1.33 **one.** A true signal value, a single bit field with a value of 1b, or a field value numerically equal to 1b.

P

3.1.34 **path.** A path is a named link between an initiating controller and a target controller. At least one connect must be made from the initiating controller to the target controller before a physical link becomes a path.

3.1.35 **port.** A port is the portion of a GPP device where it attaches to an interface. A GPP device may have more than one port. Each port may attach to the same or a different GPP interface. Each port has a port number relative to its logical element and also one or more addresses related to the implemented interface.

3.1.36 **port number.** A unique number for each port assigned by the controlling logical element.

Q

3.1.37 **queue tag.** The parameter associated with an I/O process which uniquely identifies it from other tagged I/O processes for a logical unit from the same initiating controller. The queue tag is not related to the path or paths used for the I/O process; it is related only to the initiating controller and logical unit.

R

3.1.38 **receive (an Information Packet).** The result of the successful transfer of an information packet between two logical elements.

3.1.39 **reserved.** The term used for fields, code values, and other items set aside for future standardization.

S

3.1.40 **SCSI.** SCSI-3.

3.1.41 **SCSI-3.** Small Computer System Interface - 3.

T

3.1.42 **target.** A GPP port, operating in target role, thorough which I/O processes pass for processing by a target controller. A target usually attaches to a target controller.

3.1.43 **target controller.** A logical element which normally executes I/O processes. I/O processes execute using the services of one or more ports available to a target controller operating in target role.

3.1.44 **target role.** The operating mode of a port which permits the port to receive an I/O process from another port on the same interface.

U

3.1.45 unexpected disconnect. A disconnection which occurs because of a protocol error.

V

3.1.46 vendor-specific. Something (e.g., a bit, field, code value, etc.) this standard identifies, but its use is not defined by this standard and may be used differently in various implementations.

X

3.1.47 xx. Digits 0-9, except those used as section numbers, in the text of this standard which are not immediately followed by lower-case "b" or "h" are decimal values.

3.1.48 xxb. A sequence of one or more of the digits 0 and 1 immediately followed by lower-case "b" are binary values.

3.1.49 xxh. A sequence of one or more of the digits 0-9 and the upper-case letters "A"-"F" immediately followed by lower-case "h" are hexadecimal values.

Z

3.1.50 zero. A field with a value of 0b in each bit position.

3.2 Definitions of Essential Fiber Channel Terms

C

3.2.1 class of service. A general definition of the methods by which communication circuits are maintained between communicating N_ports. Class(es) of service supported by an N_port or the fabric is one determining factor in establishing inter-operability between two N_ports or between N_ports and a fabric.

3.2.2 Class 1 service. An dedicated connection service. Class 1 service does not reflect the peer-to-peer nature of SCSI since the connection between two N_Ports is held until released. That is, the disconnect privilege is not granted in a logical sense to a target controller. With a fabric, the switch connections are set and maintained, whether used or not, until explicitly released. Frames transmitted in a certain sequence at the transmitter arrive in the same sequence at the receiver. Other attributes of Class 1 service are: required login, an establish/break connection protocol is required, end-to-end confirmation of frame transmission, maximum frame size limitations, dual flow control, and busy or reject link continue and link responses are not allowed once the connection is made.

3.2.3 Class 2 service. A frame multiplex service. Class 2 service most nearly reflects the peer-to-peer nature of SCSI. Frames transmitted in a certain sequence at a transmitter may not arrive in the same sequence at the receiver. Other attributes of Class 2 service are: required login, connectionless service, frame multiplexing at the sender or receiver, end-to-end confirmation of frame transmission, maximum frame size limitations, dual flow control, and busy and/or reject link continue and link responses are allowed.

3.2.4 CRC field. A four-byte field at the end of the frame content field. The CRC field contains information which may help determine if transmission errors occur across a link, a Fabric, or passing through any N_port or F_port. This field is normally generated by a hardware element and one or more hardware elements check it at the receivers. The CRC field value is calculated before frame transmission begins. The algorithm for generating the field is specified in the appropriate standard.

D

3.2.5 D_ID. destination identifier.

3.2.6 destination ID. destination identifier.

3.2.7 destination identifier. An identifier unique to each N_port on a GPP interface. The identifier is placed in the frame header field of each frame of a sequence which identifies the receiver of the frame.

E

3.2.8 exchange. A set of one or more sequences defined to perform some action for GPP. Sequence selection and the ordering of sequences is a function of GPP. An exchange operates between the originator and responder N_ports where the connect was made.

3.2.9 exchange protocol. A set of one or more exchanges defined for GPP to use to carry out its functions. GPP provides rules to translate the activities for each nexus into a set of one or more sequences. The exchange protocol is independent of the logical content of the information packets (sequences). That is, two or more logical functions may use the same exchange protocol. A set of exchange protocols is called an FC-4 mapping in the Fiber Channel.

F

3.2.10 F_port. A component of a fabric element which connects to the end of a link. An F_port receives frames from an N_port and directs them into the fabric. An F_port receives frames from the fabric and transmits them to an N_port.

3.2.11 fabric. A Fiber Channel entity with at least one fabric element. A Fiber Channel implementation does not require a fabric (i.e., point-to-point topology).

3.2.12 fabric element. A component of a fabric which, at minimum, performs frame switching or circuit switching. Some fabric elements may perform both functions and may perform the N_port naming function called fabric login. Fabric elements attach to each other via links to provide more sophisticated frame routing and provide a GPP interface to attach more N_ports than a single fabric element is designed to handle.

3.2.13 fabric class of service. The class(es) of service supported by a fabric.

3.2.14 fiber. A single transmission element which transfers frames in a unidirectional manner.

3.2.15 frame. The smallest unit of information transmitted across a link. A frame consists of idle words, a start of frame delimiter, frame content field, an end of frame delimiter, and idle words. A frame, excluding idle words, is from 36 to 2148 characters long. An N_port sends frames in the order established by the exchange protocol selected. Frames may not arrive at a receiving N_port in the order transmitted, depending on the makeup of the fabric, if present, and the class of service.

3.2.16 frame content field. A set of three fields in a frame: the frame header field, frame data field, and CRC field. The frame content field varies from 28 to 2140 characters in multiples of four characters.

3.2.17 frame data field. A variable length field, always a multiple of 4 bytes, which conveys logical information across a link. A frame data field has a zero minimum length and a maximum length of 2112 bytes. The frame header field has a field designated to identify the number of pad bytes, if any, in a frame data field. The concatenation of frame data field extracted from tne frames of a sequence is the principal interface between the physical elements of the Fiber Channel and GPP.

3.2.18 frame header field. A fixed length and fixed format structure, positioned after the start of frame delimiter, used by the fabric and N-ports during transmission of frames. The frame header field provides physical addresses for frame routing on the Fiber Channel. The frame header field has other fields which help determine protocol errors.

I

3.2.19 idle word. An ordered set chosen to represent the idle state of a fiber (i.e., no frame being transmitted). A primitive sequence may substitute for a idle word in some instances.

L

3.2.20 link. A physical interface with two fibers and the associated connectors at each end.

3.2.21 login. A procedure for each N_port to identify itself to, or gain its identity from, the fabric or another N_port. An N_port performs a login sequence with each N_port with which it intends to communicate and successfully exchange service parameters. Each N_port exchanges service parameters with its F_port if the fabric is present.

3.2.22 logout. A procedure for removing the service parameter agreement between two N_ports. When the procedure is complete, login must occur before an N_port can successfully perform GPP operations.

N

3.2.23 N_port. A port on the Fiber Channel which attaches to a link end point. Each N_port attaches, through a link, to exactly one F_port, if the fabric exists. An N-port is equivalent to a GPP port.

3.2.24 N_port class(es) of service. The class(es) of service supported by an N-port.

3.2.25 N_port service parameters. A set of parameters specifying the capabilities of an N_port. N_ports exchange service parameters with an N_port or F_port before attempting to use GPP exchanges.

3.2.26 node. A facility composed of one or more N_ports.

3.2.27 node identifier. An identifier which uniquely identifies a node on a GPP interface. See initiating controller ID in "3.1 Definitions" on page 3.

O

3.2.28 originator. The N_port responsible for initiating an exchange. When two N_ports are communicating either may become an originator and start an exchange. The originator of an exchange may be in either initiator role or target role.

P

3.2.29 point-to-point topology. A Fiber Channel implementation where each N_port communicates with exactly one other N_port.

3.2.30 point-to-point operation. A function of the Fiber Channel which logically connects two N_ports, when a fabric is interposed, to make them appear connected as in a point-to-point topology.

R

3.2.31 responder. The N_port that receives an exchange initiated by an originator. When two N_ports are communicating, either may become a responder to an exchange initiated by the other. The responder of an exchange may be in either initiator role or target role.

S

3.2.32 S_ID. source identifier.

3.2.33 sequence. A protocol in the Fiber Channel specifying how to transmit a set of frames between N_ports. In GPP, a sequence transmits a single information packet. Fiber Channel specifies rules for determining when each sequence ends. Fiber Channel specifies rules whereby the receiver can reassemble an information packet in the correct sequence.

3.2.34 source ID. source identifier.

3.2.35 source identifier. An identifier unique to the N_port on a GPP interface placed in the frame header field of each frame of an exchange that identifies the sender of the frame.

3.2.36 service parameters. A set of field values specifying the capabilities of an N_port or F_port. Communicating ports exchange service parameters before attempting GPP exchanges.

X

3.2.37 X_ID. Exchange identifier. The exchange identifier may be prefixed by the letters "O" or "R" meaning "Originator" and "Responder," respectively.

3.3 Editorial Conventions

Certain words and terms used in this standard have specific meanings beyond the normal English meaning. Glossaries define these words and terms (see "3.1 Definitions" on page 3 & "3.2 Definitions of Essential Fiber Channel Terms" on page 6.) or the definition appears at first use in the text. Names of signals, phases, and messages are in all upper-case letters (e.g., INVALID INFORMATION PACKET).

Words have the normal technical English definition unless the word or phrase is defined in context or in one of the glossaries, then that definition is used. Some words may have definitions unique to the sections where they are used. Every effort has been made to limit the number of such instances.

Numbered items in GPP do not represent any priority. Any priority is explicitly indicated.

The ISO convention of numbering is used (i.e., the thousands and higher multiples are separated by a space and a comma is used as the decimal point.

In all of the text, tables, and figures of this standard, the most significant bit of a binary quantity is shown on the left side and represents the highest algebraic value position in the quantity.

The term "shall" is used to indicate a mandatory rule. If such a rule is not followed, the results are unpredictable unless indicated otherwise. In addition, if such a rule is not followed, the implementation is not in conformance with this standard.

If a field is specified as not meaningful, the entity that receives the field shall not check that field.

If a conflict arises between text, tables, and figures, the order of precedence to resolve conflicts is text, tables, and, lastly, figures. Not all tables and figures are fully described in text. Tables show data formats and values; figures are illustrative of the text and tables. NOTEs and IMPLEMENTATION NOTEs have been kept to a minimum, but when they occur they do not constitute any requirements for implementors.

3.4 Abbreviations, acronyms and symbols

Abbreviations, acronyms and symbols applilcable to GPP are listed. Definitions of several of these items can be found in "3.1 Definitions" on page 3 & "3.2 Definitions of Essential Fiber Channel Terms" on page 6.

3.4.1 Abbreviations and acronyms

FC-PH. Fibre Channel Physical standard

GPP. Generic Packetized Protocol

IC. initiating contgroller

ILE. Interface Logical Element

I/O. input/output

IP. Information Packet

LSB. least significant bit

MSB. most significant bit

SAM. SCSI-3 Architecture Model

SIP. SCSI-3 Interlocked Protocol

SPI. SCSI-3 Parallel Interface

TC. target controller

3.4.2 Symbols

||. Concatenation. Merging two items to form one new item.

4.0 GPP Information Packet Structure

This section contains the specfication for information packet construction for Generic Packeticed Protocol.

4.1 Information Packet Structure

An information packet is set of self-describing interface logical elements arranged in prescribed orders. The interface logical elements are surrounded by appropriate interface control fields necessary to send the logical content of the packet on a physical transport layer (Table 1 on page 10).

- Information packets shall be a multiple of 4 bytes in length.
- An information packet shall contain at least one interface logical element and shall not contain more than 255 interface logical elements.

  Interface Logical Elements may have pad bytes added to construct interface logical elements that are multiples of 4 bytes. The value of pad bytes in GPP shall be 00h.

If an information packet is received that is not a multiple of 4 bytes long or there are zero or greater than 255 interface logical elements in the information packet, the information packet shall be rejected and no other action is taken. The target controller returns an information packet with an INVALID INFORMATION PACKET message interface logical element.

Table 1. Information Packet Structure

Bit Byte	7-0
0 to m	Interface Control Prefix Fields
m+1 to n	One to 255 Interface Logical Elements

4.1.1 Interface Control Fields

The interface control fields encapsulate the information content of the logical operations. The identify function for both initiating controllers and target controllers uses the interface control prefix fields to define a nexus.

4.1.1.1 Interface Control Prefix Fields

Table 2. Interface Control Prefix Fields

Bit Byte	7-0
0 - 1	(MSB) Remaining Packet Length (LSB)
2	Packet Type

Table 2 (Page 2). Interface Control Prefix Fields

Bit / Byte	7-0					
3	Reserved					
4	7 MltPath	6 SuspMpth	5 EnbSpvr	4-0 Reserved		
5	7 LUN Valid	6 QNexus	5 HOQ	4 OrdSim	3 ACANexus	2-0 Reserved
6	Initiating Controller Port Number					
7	Reserved					
8	Original Initiator GPP Address					
9	Reserved					
10	Original Target GPP Address					
11	Target Controller Port Number					
12	LUNID					
13	Queue Tag					
14-15	Sender's Sequence Number					

Table 2 on page 10 shows the interface control prefix fields. All reserved fields and bits shall be set to zero.

The remaining packet length field indicates the total number of remaining bytes in an information packet in bytes, including the remaining interface control prefix fields. The actual usable maximum packet length is the result of negotiation and may be further restricted by maximum burst length negotiations. This value shall be of the form ((N*4) - 2) and greater than or equal to 18 ((16 + 4) - 2) and less than or equal to 65530 (65532 - 2).

The packet type field identifies the general content of an information packet.

- A packet type code of 00h indicates an information packet to start a new nexus.
- A packet type code of 01h indicates a target controller initiated packet to terminate an I/O process.
- A packet type code of 02h indicates an intermediate information packet from an initiating controller to continue an I/O process established with an information packet of type 00h.
- A packet type code of 03h indicates an intermediate information packet from a target controller to continue an I/O process established with an information packet of type 00h.
- A packet type code of 04h indicates that a nexus is requested to be established by the initiating controller from a target controller because of asynchronous event notification.
- Packet type code values of 05h through FFh are reserved.

A multiple path (MltPath) field value of one specifies that the initiating controller considers the target controller capable of multiple path operation. When the value is zero, the initiating controller considers the target controller to be in single path mode. The setting has no effect on any other I/O process. This value shall remain unchanged in all information packets for the nexus. The MltPath field indicates the initiating controller state of the path. The target controller shall confirm that its state is the same. If the target controller state is different, it shall not process the information packet and shall indicate the error by sending an information packet with the INVALID INFORMATION PACKET message. Accepting this field set to one is optional.

A suspend multiple path (SuspMpth) field value of one specifies that the initiating controller requires the target controller to stop multiple path operation for the I/O process. The effect is to form an H_I_T_C nexus from the H_C nexus. The SuspMpth field is interpreted for each nexus. The setting has no effect on any other I/O process in the target controller. If the MltPath field is set to 0b, the SuspMpth field is not meaningful and the target controller shall operate in single path mode for the nexus. If the MltPath field is set to 1b and the SuspMpth field is set to 1b, the target controller shall operate in single path mode for the nexus. If the MltPath field is set to 1b and the SuspMpth field is set to 0b, the target controller may operate in single path mode or multiple path mode for the nexus. Accepting this field set to one is optional.

A enable supervisor command (EnbSpvr) field of one specifies that the initiating controller has granted the target controller the privilege of executing any valid supervisor command received during the I/O process identified in which the enable supervisor command field is set to 1b. If the field value is zero and a supervisor command is present in the I/O process, the target controller shall terminate the I/O process with auto-contingent allegiance. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to SUPERVISOR COMMANDS NOT ENABLED. (SAM update required? GRS) The enable supervisor command field setting is interpreted for each nexus. The setting has no effect on any other I/O process. Accepting this field set to one is optional.

A logical unit number valid (LUN Valid) field value of zero specifies that the I/O process is to the target controller and forms an H_C nexus. A LUN Valid field value of one specifies that the information packet is for a logical unit and forms an H_C_x_y nexus. Certain information packets cannot be directed to a specific logical unit. The target controller shall use this field, when 1b, as a valid field indicator for the LUN, and LUNID fields. This value shall remain unchanged in all information packets for the nexus. An I/O process with a LUN Valid field value of zero shall operate in single path mode.

The queue nexus (QNexus) field value when one indicates that the nexus is of the form H_C_L_Q. When the queue nexus field falue is zero, the nexus formed will be of the form H_C_L. The target controller shall use this field value, when 1b, as a valid field indicator for the HOQ, OrdSim, ACANexus, and Queue Tag fields. This value shall remain unchanged in all information packets for the nexus. If tagged queuing is implemented, support of this field set to one shall be implemented.

A HOQ field value of one specifies that the I/O process is to be placed at the head of the I/O process queue for the initiating controller. This value shall remain unchanged in all information packets for the nexus. This field is to be interpreted only when QNexus is 1b. If the QNexus field is supported, this bit shall be supported.

A OrdSim field value of one specifies ordered queue processing in the target controller. A OrdSim field value of zero specifies that a simple tagged I/O process is requested in the target controller. This value shall remain unchanged in all information packets for the nexus. This field is to be interpreted only when QNexus is 1b. If the QNexus field is supported, this bit shall be supported.

A ACANexus field value of one specifies that the I/O process is in response to an established ACA condition. This value shall remain unchanged in all information packets for the nexus. This field is to be interpreted only when QNexus is 1b. If the QNexus field is supported, this bit shall be supported.

The next nine fields identify the path on which the connect was made for the nexus. These fields shall not change in any information packet of any type transferred for the nexus once established as specified below. These fields, along with the LUN Valid, QNexus, HOQ, OrdSim, and ACANexus fields identify the nexus to be reestablished in any later connection on any allowable path for the nexus.

The initiating controller port number field identifies the internal port number assigned by the initiating controller to the port where the connect was made using an information packet packet type code of 00h. This value shall remain unchanged in all information packets for the nexus.

The original initiator GPP address field maps to the external name of the GPP device on the initiating controller where the connect was made using an information packet packet type code of 00h. This value shall remain unchanged in all information packets for the nexus.

The original target GPP address field maps to the external name of the GPP device, from the initiating controller's perspective on the selected GPP bus, where the connect was made using an information packet packet type code of 00h. This value shall remain unchanged in all information packets for the nexus.

The target controller port number field identifies the physical port number assigned by the target controller to the target GPP address. This field shall be FFh when the packet type code is 00h. The target controller supplies its unique port number value in any response information packet. Once the initiating controller receives an information packet for the nexus that contains the target port number, all later information packets sent with packet type code 00h on the same physical path shall have the corresponding target port number placed in this field. The target controller may receive multiple information packets during the connect with a target port number value of FFh. If the target controller receives any information packet with packet type code 00h, other than as specified above, and the target port number does not match the internally assigned port number, the target controller shall not process the information packet and shall indicate the error by sending an information packet with the INVALID INFORMATION PACKET message.

The logical unit number (LUNID) field specifies a logical unit number if the LUN Valid field value is set to one. The response to an invalid value in the LUNID field is described in a later section. This field is not meaningful when the LUN Valid field value is zero. This value shall remain unchanged in all information packets for the nexus.

The queue tag field is meaningful only when the QNexus field value is 1. If the QNexus field value is zero, this field shall be set to 00h. If the target controller detects any value other than 00h when the QNexus field value is zero, it shall indicate the error by sending an information packet with an INVALID INFORMATION PACKET message. This value shall remain unchanged in all information packets for a nexus. If the QNexus field is supported, this bit shall be supported. If tagged queuing is implemented, support of this field set to one shall be implemented.

The sender's sequence number is a value from 0 to 65535 which specifies the sequence of an information packet relative to other information packets for the same nexus. The initiating controller shall set this value to zero when the packet type code is 00h. It then increments the value by one for each subsequent packet it sends for the same nexus. If the value reaches 65535, the next sequence number is 0 (i.e., the sequence counter wraps). The target controller shall set the value of this field to zero on the first information packet sent in response to a nexus. It shall then increment the counter by one for each subsequent information packet. If the value reaches 65535, the next sequence number is 0 (i.e., the sequence counter wraps). The sender's sequence number shall be zero when the packet type code value is 04h.

If any reserved field is not zero, the receiver of the information packet shall not process the information packet further. It shall send an information packet with an INVALID INFORMATION PACKET message.

4.1.2 Interface Logical Elements

The interface logical elements are messages, command descriptor blocks, command parameter data, command response data, logical data, status, and autosense. A self-describing set of fields prefixes each interface logical element to permit the information packet content to be correctly identified, checked, and processed at the receiving GPP port.

A self-describing interface logical element is composed of five fields: the remaining element length, the element type, a pad byte count field, the logical element bytes, and the ILE pad bytes. See Table 3 on page 13.

Table 3. Interface Logical Element

Bit / Byte	7-0	
0 - 1	(MSB) Remaining Element Length (LSB)	
2	Element Type	
3	Reserved	Pad Byte Count
4 - n	Logical Element Bytes	
n + m	Zero to Three ILE Pad Bytes	

The remaining element length is a two-byte binary field. The value range is from 2 (4 - 2) to 65514 ((65532 - 16) - 2). If the element length value is 2, the interface logical element is not meaningful and shall be ignored by the receiver of the information packet. This is not considered an error. In this case, the element type and pad byte count fields are not meaningful. The value is 2 plus the number of logical element bytes and pad bytes.

The element type is a one-byte binary field. When the remaining element length field value is greater than 2, the valid values are defined in Table 4.

Table 4. Element Type Codes

Element Type Value	Interface Logical Element	Source
0	Message	IC/TC
1	Command Descriptor Block	IC
2		
Command Parameter Data	IC	
3	Command Response Data	TC
4	Logical Data	IC/TC
5	Status	TC
6	Autosense	TC
7 - 255	Reserved	N/A

The pad byte count field indicates the number of pad bytes added at the end of an interface logical elements to construct an interface logical element that is a multiple of 4 bytes in length.

The logical element bytes field is a 0 to 65512 byte variable length field composed of the information about the element type declared in the element type field. The maximum length which may be used for this field may vary by the physical interface implemented and the currently negotiated maximum information packet length. The maximum length for any one interface logical element shall be limited by the requirements of the currently negotiated maximum information packet length. The current maximum interface logical element length is calculated as current maximum information packet length - 20.

Each interface logical element shall be rounded up to the nearest multiple of 4 bytes by adding from zero to three pad bytes. The value of each pad byte shall be 00h. If more than 3 pad bytes are present, the receiver shall indicate the error by sending an information packet with the INVALID INFORMATION PACKET message.

If any reserved field is not zero, the receiver of the information packet shall not process the information packet further. It shall send an information packet with an INVALID INFORMATION PACKET message.

4.1.2.1 Message Interface Logical Element

A message interface logical element shall contain one complete message. A multiple byte message shall not be split into bursts.

4.1.2.2 Command Descriptor Block Interface Logical Element

A command descriptor block interface logical element shall contain one command descriptor block as described in the appropriate SCSI-3 standards. Command parameter data is transferred in its own interface logical element. A command descriptor block shall not be split into bursts.

4.1.2.3 Command Parameter Data Interface Logical Element

The command parameter data interface logical element is a burst of data transferred for a command which is sent as an addendum to the command descriptor block. The CDB identifies the total length of the command parameter data in the parameter length fields, etc. Each command parameter data ILE shall contain data for only one command.

4.1.2.4 Command Response Data Interface Logical Element

The command response data interface logical element is a burst of data transferred from a target controller to an initiating controller for a command that is not taken from the logical blocks of the LUN. Sense data and MODE SENSE pages are examples of command response data. The CDB identifies the maximum length of the command response data in the allocation length fields, etc. Each command response data ILE shall contain data for only one command.

4.1.2.5 Logical Data Interface Logical Element

A logical data interface logical element consists of a burst of data for the primary function of a logical unit requested for transfer between an initiating controller and a target controller. Each logical data ILE shall contain data for only one command.

4.1.2.6 Status Interface Logical Element

The status interface logical element contains a status value. Status shall not be split into bursts. Each Status ILE shall contain only one status code value.

4.1.2.7 Autosense Interface Logical Element

The autosense interface logical element is the sense data for an I/O process which terminates with auto-contingent allegiance. Autosense shall not be split into bursts. Each autosense ILE shall contain data to report one event. An auto-contingent allegiance remains in effect until terminated by the initiating controller.

5.0 GPP I/O Process Control

This section defines the message interface logical elements required with the Generic Packetized Protocol. Table 6 identifies the messages and the requirements for implementation by initiating controllers and target controllers.

5.1 Message Structure

Messages allow communication of information between an initiating controller and target controller not contained in the other information packet interface logical elements. The logical element bytes of a message interface logical element may be one, two, or multiple bytes in length. Zero or more messages may be transferred along with other interface logical elements in a single information packet.

The initiating controller and target controller controller are required to control the content of information packets when it sends a message interface logical element in an information packet for certain messages as identified in Table 6. The "Single Mesasge ILE in IP" column when "YES" indicates the an information packet containing the message ILE shall be the only ILE in the information packet; no other ILE of any type shall be present in the information packet. The assigned ranges to messages and their lengths is defined in Table 5 on page 16.

All GPP devices acting as initiating controllers shall implement the mandatory messages tabulated in the "Init" column of Table 6 for information packet operations. All GPP devices acting as target controllers shall implement the mandatory messages tabulated in the "Targ" column of Table 6 for information packet operations. This is modified by messages unique to the parallel interface. See the "-" prefix on the message name for these messages.

One-byte, Two-byte, and extended message formats are defined. The first byte of the message ILE logical element bytes determines the format as defined in Table 5 on page 16.

Table 5. Message Element Format Codes

Format Code	Message Element Byte Format
00h	One-Byte Message (I/O PROCESS COMPLETE)
01h	Extended Messages
02h - 1Fh	One-Byte Messages
20h - 2Fh	Two-Byte Messages
30h - 7Fh	Reserved
80h - FFh	Reserved Messages

One-byte messages consist of a single byte where the value of the byte determines which function is to be performed as defined in Table 6.

Two byte messages consist of a message code followed by a parameter byte. The message code determines which function is to be performed using the value in the parameter byte.

A value of 01h in the first byte of a message indicates the beginning of a multiple-byte extended message. The minimum number of bytes in an extended message is three. The maximum number of bytes in an extended message is 258 bytes. The extended message codes and the extended message format are shown in Table 8 on page 18 and Table 7 on page 18, respectively.

Table 6. Message Codes

Code	Reqd IC S/R	Reqd TC S/R	Message Name	Direc-tion	Direc-tion	Single Message ILE in IP
06h	O/N	N/M	ABORT		Out	Yes
0Dh	O/N	N/M	ABORT I/O PROCESS		Out	Yes
16h	M/N	N/M	CLEAR ACA CONDITION		Out	Yes
0Eh	O	M	CLEAR QUEUE		Out	Yes
0Ch	O/N	N/M	DEVICE RESET		Out	Yes
00h	N/M	M/N	I/O PROCESS COMPLETE	In		No
***	M/M	M/M	EXTENDED MESSAGE REJECT	In	Out	Yes
***	M/M	M/M	*INVALID BUS PHASE DETECTED	In	Out	Yes
***	M/M	M/M	INVALID INFORMATION PACKET	In	Out	Yes
0Ah	N/M	M/N	LINKED COMMAND COMPLETE	In		No
0Bh	N/M	M/N	LINKED COMMAND COMPLETE (with Flag)	In		No
***	N/M	M/N	MODIFY DATA POSITION	In		Yes
***	M/M	M/M	*PARITY ERROR	In	Out	Yes
***	M/M	M/M	RESEND PREVIOUS INFORMATION PACKET	In	Out	Yes
***	M/M	M/M	*SYNCHRONOUS DATA TRANSFER REQUEST	In	Out	Yes
***	M/M	M/M	SYNCHRONOUS PACKET TRANSFER REQUEST	In	Out	Yes
11h	O/N	N/M	TERMINATE I/O PROCESS		Out	Yes
12h	M/M	M/M	TRANSFER READY	In	Out	Yes
02h-1Fh			Reserved If Not Listed Above			
30h-FFh			Reserved			

Key: M = Mandatory support, O = Optional support.
TC = Target Controller
IC = Initiating Controller
ILE = Interface Logical Element

IP = Information Packet
In = Target to initiator, Out = Initiator to target.
Yes = IP has Single Message ILE and no other ILEs
No = May be other ILEs in IP
N = Not Sent or Received
* = SPI implementations only.
*** = Extended message Code Value (see Tables Table 8 on page 18)

NOTES:

The extended message length field specifies the length in bytes of the extended message code plus the extended message arguments to follow. The total length of a message is equal to the extended message length plus two. The minimum value is 1; the maximum value is 255;

Table 7. Extended Message Element Format

Byte	Value	Description
0	01h	Extended Message Format Code
1	n	Extended Message length
2	-	Extended Message Code (Table 8 on page 18)
3 - n+1	-	Extended Message Parameters

The extended message codes are listed in Table 8 on page 18. The extended message arguments are specified with each extended message description. For extended messages, the message code is followed by a variable length set of zero or more parameter bytes. The function to be performed is determined by the value in the Extended Message Code field.

Table 8. Extended Message Codes

Code Value	Name
	(Ordered by Extended Message Name)
04h	EXTENDED MESSAGE REJECT
02h	INVALID BUS PHASE DETECTED (SPI Only)
06h	INVALID INFORMATION PACKET
00h	MODIFY DATA POSITION
03h	PARITY ERROR (SPI Only)
05h	RESEND PREVIOUS INFORMATION PACKET
01h	SYNCHRONOUS DATA TRANSFER REQUEST (SPI Only)
07h	SYNCHRONOUS PACKET TRANSFER REQUEST
08h	TERMIONATE I/O PROCESS

Table 8 (Page 2). Extended Message Codes	
Code Value	**Name**
09h	TRANSFER READY
0Ah - FFh	Reserved
	(Ordered by Extended Message Code)
00h	MODIFY DATA POSITION
01h	SYNCHRONOUS DATA TRANSFER REQUEST (SPI Only)
02h	INVALID BUS PHASE DETECTED (SPI Only)
03h	PARITY ERROR (SPI Only)
04h	EXTENDED MESSAGE REJECT
05h	RESEND PREVIOUS INFORMATION PACKET
06h	INVALID INFORMATION PACKET
07h	SYNCHRONOUS PACKET TRANSFER REQUEST
08h	TERMIONATE I/O PROCESS
09h	TRANSFER READY
0Ah - FFh	Reserved

GPP messages are defined below. The requirements for message implementation are given in Table 6 on page 17.

5.2 Message Descriptions

5.2.1 ABORT

The ABORT message is sent from the initiating controller to the target controller to clear the identified I/O process plus any I/O processes for the H_C_L portion of the identified nexus. Pending data, status, and other I/O processes for any other H_C_x_y nexus shall not be cleared as a result of processing this message.

Previously established conditions, including MODE SELECT parameters, reservations, and auto-contingent allegiance shall not be changed by the ABORT message.

The information packet containing an ABORT message shall contain only one interface logical element, a message interface logical element.

It is not an error to issue this message to a logical unit or target routine that does not have any I/O processes for the initiating controller.

If only an H_C nexus has been established, no status or message shall be sent for the current I/O process and no other I/O process shall be affected.

5.2.2 ABORT I/O PROCESS

An ABORT I/O PROCESS message is sent from an initiating controller to a target controller to clear the I/O process identified. If the target controller has already started execution of the I/O process, execution shall be halted. The logical data contents of the logical unit may have been modified before the information packet was processed. In either case, any pending status or data for the I/O process shall be cleared and no additional ILEs shall be sent to the initiating controller.

Pending status, data, and commands for other active or queued I/O processes shall not be affected. Execution of other I/O processes shall not be aborted as a result of processing this message.

Previously established conditions, including MODE SELECT parameters, reservations, and auto-contingent allegiance shall not be changed by the ABORT I/O PROCESS message.

The information packet containing an ABORT I/O PROCESS message shall contain only one interface logical element, a message interface logical element. p It is not an error to issue this message to a logical unit or target routine that does not have any active or queued I/O processes.

5.2.3 CLEAR ACA CONDITION

The CLEAR ACA CONDITION message is sent from an initiating controller to a target controller to terminate an auto-contingent allegiance condition. The auto-contingent allegiance condition ends on successful processing of the CLEAR ACA CONDITION message.

The information packet containing a CLEAR ACA CONDITION message shall contain only one interface logical element, a message interface logical element.

If a CLEAR ACA CONDITION message is received by a target controller when no auto-contingent allegiance condition is active, the message shall not be rejected. The message shall be ignored.

5.2.4 CLEAR QUEUE

A CLEAR QUEUE message is sent from an initiating controller to a target controller to perform an action equivalent to receiving a series of ABORT messages from each initiating controller. All I/O processes, from all initiating controllers, for the specified logical unit shall be cleared. The logical data of the logical unit may have been altered by partially completed I/O processes.

All pending status and data for that logical unit for all initiating controllers shall be cleared. No status or message shall be sent for any of the I/O processes. A unit attention condition shall be generated for all other initiating controllers with I/O processes for that logical unit. When reporting the unit attention condition the additional sense code shall be set to I/O PROCESSES CLEARED BY ANOTHER INITIATOR. Previously established conditions, including MODE SELECT parameters, reservations, and auto-contingent allegiance shall not be changed by the CLEAR QUEUE message.

The information packet containing a CLEAR QUEUE message shall contain only one interface logical element, a message interface logical element.

5.2.5 DEVICE RESET

The DEVICE RESET message is sent from an initiating controller to direct a target controller to clear all I/O processes on that GPP device. This message forces a reset condition in the selected GPP device. The target controller shall create a unit attention condition for all initiating controllers for each logical unit.

The information packet containing a DEVICE RESET message shall contain only one interface logical element, a message interface logical element.

All pending status and data for that logical unit for all initiating controllers shall be cleared. No status or message shall be sent for any of the I/O processes. A unit attention condition shall be generated for all initiating controllers. When reporting the unit attention condition the additional sense code shall be set to ?????. Previously established conditions, including MODE SELECT parameters, reservations, and auto-contingent allegiance shall be reset by the DEVICE RESET message.

5.2.6 I/O PROCESS COMPLETE

The I/O PROCESS COMPLETE message is sent from a target controller to an initiating controller to indicate that the execution of an I/O process has completed and that valid status has been sent to the initiating controller. The I/O process may have completed successfully or unsuccessfully as indicated by the status value.

The I/O PROCESS COMPLETE message when present in an information packet, shall be the last, or only, ILE in the information packet.

5.2.7 EXTENDED MESSAGE REJECT

The EXTENDED MESSAGE REJECT message (Table 9 on page 21) is sent from either the initiating controller or target controller to indicate that a portion of a message in a message interface logical element in the last information packet was inappropriate or has not been implemented.

The content of the extended message identifies the message interface logical element. The message indicates the reason for the rejection. The message also indicates the number of additional information packets cleared for the same I/O process. This provides a logical interlock so that the receiver of the message can determine which message byte is in error.

The information packet containing an EXTENDED MESSAGE REJECT message shall contain only one interface logical element, a message interface logical element.

Table 9. EXTENDED MESSAGE REJECT Message Format

Byte	Value	Description
0	01h	Extended Message Format Code
1	0Ah	Extended Message length
2	04h	Extended Message Code
3	xxh	ILE Position

Table 9 (Page 2). EXTENDED MESSAGE REJECT Message Format		
Byte	**Value**	**Description**
4	xxh	Byte Position in ILE
5	xxh	Reject Reason Code
6-7	xxxxh	No. of Information Packets Cleared
8-9	xxxxh	Sender's Sequence Number of the First IP Cleared
10-11	0000h	Reserved

The ILE Position field value is the ordinal position of the message interface logical element within the information packet containing the error. The ILE Position is a value between 1 to 255.

The Byte Position in ILE is the byte position of the byte in error within the element bytes of the message in error. The value is 0 for the first message byte in the ILE. The range is 0 - 255.

The Reject Reason Code identifies the reason for sending the message. Allowable values for the Reject Reason Code are:

- A reject reason code of 00h indicates that the message code in message is not implemented.
- A reject reason code of 01h indicates that the message is inappropriate.
- A reject reason code of 02h indicates that the extended message length field in an extended message is invalid for the extended message code.
- A reject reason code of 03h indicates that a parameter byte within two-byte or extended messages is invalid.
- Reject reason codes 04h-FFh are reserved.

The Number of Information Packets Cleared field contains a count of additional information packets (not counting the information packet in which the error was detected) for this I/O process which have been cleared as a result of this error.

The Sender's Sequence Number of the First IP Cleared field identifies the information packwet sequence number where the first information packet was cleared.

5.2.8 INVALID BUS PHASE DETECTED (Parallel)

The INVALID BUS PHASE DETECTED message (Table 10 on page 22) is sent from either logical element to the other to inform it that an illegal or reserved bus phase was detected on the SPI.

The information packet containing an INVALID BUS PHASE DETECTED message shall contain only one interface logical element, a message interface logical element.

Table 10. INVALID BUS PHASE DETECTED Message Format		
Byte	**Value**	**Description**
0	01h	Extended Message Format Code
1	02h	Extended Message length

Table 10 (Page 2). INVALID BUS PHASE DETECTED Message Format					
Byte	**Value**	**Description**			
2	02h	Extended Message Code			
3	xxh	7-3 00000b	2 C/D	1 I/O	0 Msg

The C/D bit shall contain 1 if the phase where the phase error was detected has the C/D signal set to true. Otherwise, the C/D bit shall be set to zero.

The I/O bit shall contain 1 if the phase where the phase error was detected has the I/O signal set to true. Otherwise, the I/O bit shall be set to zero.

The Msg bit shall contain 1 if the phase where the phase error was detected has the MSG signal set to true. Otherwise, the Msg bit shall be set to zero.

These three fields identify the information transfer phase which was detected in error.

5.2.9 INVALID INFORMATION PACKET

The INVALID INFORMATION PACKET message (Table 11 on page 23) is sent from either the initiating controller or target controller in an information packet to indicate that all or a portion of an information packet transferred was invalid. The format of the information packet does not conform to the rules for information packet construction. Errors in the content of the element bytes of interface logical elements is not pointed to using this message; see sense data.

The information packet containing an INVALID INFORMATION PACKET message shall contain only one interface logical element, a message interface logical element.

Table 11. INVALID INFORMATION PACKET Message Format		
Byte	**Value**	**Description**
0	01h	Extended Message Format Code
1	0Ah	Extended Message length
2	06h	Extended Message Code
3	xxh	ILE Position
4	xxh	Reason Code
5 - 6	xxxxh	(MSB) Relative Position (LSB)
7 - 8	xxxxh	(MSB) No. Information Packets Cleared (LSB)
9-10	xxxxh	Sender's Sequence Number of the First IP Cleared
11	00h	Reserved

The ILE Position field is the ordinal position of the interface logical element group within the information packet containing the error. The ILE Position is a value between 1 to 255. There may be more than one message interface logical element in each information packet. A value of 0 in ILE Position indicates that the information control prefix or suffix fields are in error.

The Reason Code field identifies the reason for sending the message. Valid Reason Code values are:

- A reason code of 00h indicates that the interface logical element type in a descriptor is invalid.
- A reason code of 01h indicates that the interface logical element length field is invalid.
- Reason codes 02h-FFh are reserved.

The Relative Position field identifies the byte position of the field or bit in error. The exact bit position is not provided. The value gives the relative position of the byte to the beginning of the ILE when the ILE Number value is 1 to 255. The value gives the relative position of the byte to the beginning of the information packet when the ILE Number value is zero.

The Number of Information Packets Cleared field contains a count of additional information packets (not counting the information packet in which the error was detected) for this I/O process which have been cleared as a result of this error.

The Sender's Sequence Number of the First IP Cleared field identifies the information packwet sequence number where the first information packet was cleared.

5.2.10 LINKED COMMAND COMPLETE

The LINKED COMMAND COMPLETE message is sent from a target controller to an initiating controller to indicate that the execution of a linked command has completed and that status has been sent. The next command for the I/O process may come from an information packet already transferred or it may come from the initiating controller in a new information packet.

5.2.11 LINKED COMMAND COMPLETE (WITH FLAG)

The LINKED COMMAND COMPLETE (WITH FLAG) message is sent from a target controller to an initiating controller to indicate that the execution of a linked command (with the flag bit set to one in the CDB) has completed and that status has been sent. The next command for the I/O process may come from an information packet already transferred or it may come from the initiating controller in a new information packet.

5.2.12 MODIFY DATA POSITION

Table 12. MODIFY DATA POSITION Message Format

Byte	Value	Description
0	01h	Extended Message Format Code
1	0Ah	Extended Message length
2	00h	Extended Message Code
3 - 6	xxxxxxxxh	(MSB) Change in Relative Position (LSB)
7 - 8	xxxxh	(MSB) No. of Information Packets Cleared (LSB)

Table 12 (Page 2). MODIFY DATA POSITION Message Format		
Byte	**Value**	**Description**
9-10	xxxxh	Sender's Sequence Number of the First IP Cleared
11	00h	Reserved

The MODIFY DATA POSITION message (Table 12 on page 24) requests that the signed argument be added (two's complement) to the value of the current logical block position in the initiating controller. This message shall control or adjust the logical block data transfer position for an I/O process.

The information packet containing a MODIFY DATA POSITION message shall contain only one interface logical element, a message interface logical element.

If the change in relative position field value, in twos complement format, added to the current initiating controller results in a value less than zero or greater than the total length for the command, the initiating controller shall transfer an information packet containing an ABORT I/O PROCESS message and shall not process subsequent information packets for the I/O process.

The Number of Information Packets Cleared contains a count of additional information packets (not counting the information packet in which the error was detected) for this I/O process which have been cleared as a result of this error.

The Sender's Sequence Number of the First IP Cleared field identifies the information packwet sequence number where the first information packet was cleared.

5.2.13 PARITY ERROR (Parallel)

The PARITY ERROR message (Table 13 on page 25) is sent from either logical element to the other to inform it that a parity error was detected during an information packet transfer on a parallel interface.

The information packet containing a PARITY ERROR message shall contain only one interface logical element, a message interface logical element.

Table 13. PARITY ERROR Message Format					
Byte	**Value**	**Description**			
0	01h	Extended Message Format Code			
1	02h	Estended Message length			
2	03h	Extended Message Code			
3	xxh	7-3 00000b	2 C/D	1 I/O	0 Msg

The C/D bit shall contain 1 if the invalid phase detected has the C/D signal set to true. Otherwise, the C/D bit shall be set to zero.

The I/O bit shall contain 1 if the invalid phase detected has the I/O signal set to true. Otherwise, the I/O bit shall be set to zero.

The Msg bit shall contain 1 if the invalid phase detected has the MSG signal set to true. Otherwise, the Msg bit shall be set to zero.

These three fields identify the information transfer phase where the parity error was detected.

If the error is detected during transfer of the information packet prefix fields, the receiving logical element shall assert the ATTENTION signal during the same phase to interlock the error with the current I/P process. The receiving logical element shall form an information packet using the information from the information packet prefix control fields, as received and include a PARITY ERROR message as the only interface logical element in the information packet.

5.2.14 RESEND PREVIOUS INFORMATION PACKET

The RESEND PREVIOUS INFORMATION PACKET message (Table 14 on page 25) is sent by an initiating controller or a target controller to indicate that one or more previously transferred information packets must be resent. The synchronous packet transfer request negotiation has established an agreement for the maximum number of packets to be retained by both the initiating controller and the target controller.

Table 14. RESEND PREVIOUS INFORMATION PACKET Message Format

Byte	Value	Description
0	01h	Extended Message Format Code
1	03h	Extended Message length
2	05h	Extended Message Code
3	00h	Reserved
4	xxh	Previous Packet Number

A value of zero in the previous packet number field causes the device to resend the last information packet transferred. Values of 1 through 255 represent the 2nd to 256th most recent packets transmitted where retransmission is to start. The logical element receiving the RESEND PREVIOUS INFORMATION PACKET message shall resend all information packets, in order, starting with the indicated information packet.

It shall be an error for a GPP device to transmit a value in the previous packet number field which is logically larger than the limit negotiated using the SYNCHRONOUS PACKET TRANSFER REQUEST message. If this error is detected, the logical element processing this message shall send an information packet to the other logical element which contains an EXTENDED MESSAGE REJECT message pointing to the Previous Packet Number field in error.

The information packet containing an RESEND PREVIOUS INFORMATION PACKET message shall contain only one interface logical element, a message interface logical element.

5.2.15 SYNCHRONOUS DATA TRANSFER REQUEST (Parallel)

A SYNCHRONOUS DATA TRANSFER REQUEST (SDTR) message (Table 15 on page 26) exchange shall be initiated by any GPP device whenever a previously-arranged data transfer agreement may have become invalid. The default agreement shall be asynchronous information transfer mode.

The information packet containing an SYNCHRONOUS DATA TRANSFER REQUEST message shall contain only one interface logical element, a message interface logical element.

The SDTR message exchange establishes the permissible transfer periods and the REQ/ACK offsets for all logical units and target routines on the two devices. This agreement only applies to the DATA OUT information transfer phase used by GPP.

Table 15. SYNCHRONOUS DATA TRANSFER REQUEST Message Format

Byte	Value	Description

Table 15 (Page 2). SYNCHRONOUS DATA TRANSFER REQUEST Message Format		
Byte	**Value**	**Description**
0	01h	Extended Message Format Code
1	03h	Extended Message length
2	01h	Extended Message Code
3	xxh	Transfer Period (m times 4 nanoseconds)
4	xxh	REQ/ACK Offset

The transfer period field is the minimum time allowed between leading edges of successive REQ pulses and of successive ACK pulses to meet the device requirements for successful reception of data.

The REQ/ACK offset field is the maximum number of REQ pulses allowed to be outstanding before the leading edge of its corresponding ACK pulse is received at the receiving GPP device. This value is chosen to prevent overflow conditions in the receiving buffer and offset counter. A REQ/ACK offset value of zero shall indicate asynchronous data transfer mode; a value of FFh shall indicate unlimited REQ/ACK offset.

The synchronous data transfer agreement becomes invalid, and reverts to asynchronous data transfer mode, after any condition which may leave the data transfer agreement in an indeterminate state such as:

- after a reset event
- after a DEVICE RESET message and
- after a power cycle.

In addition, an GPP device may initiate an SDTR message exchange whenever it is appropriate to negotiate a new synchronous data transfer agreement.

The originating device (the device that sends the first of the pair of SDTR messages) sets its values according to the rules above to permit it to receive data successfully. If the responding device can also receive data successfully with these values (or smaller transfer periods or larger REQ/ACK offsets or both), it returns the same values in its SDTR message. If it requires a larger transfer period, a smaller REQ/ACK offset, or both to receive data successfully, it substitutes values in its SDTR message as required, returning unchanged any value not required to be changed. Each device when transmitting data shall respect the limits set by the other's SDTR message, but it is permitted to transfer data with larger transfer periods, smaller REQ/ACK offsets, or both than specified in the other's SDTR message. The successful completion of an exchange of SDTR messages implies an agreement, retained by both the initiating controller and target controller. The negotiation shall end when both devices exchange the same values or the agreement reaches a REQ/ACK offset of zero.

5.2.16 SYNCHRONOUS PACKET TRANSFER REQUEST

A SYNCHRONOUS PACKET TRANSFER REQUEST (SPTR) message (Table 16 on page 27) exchange shall be initiated by an GPP device whenever a previously arranged information packet transfer agreement may have become invalid. Each GPP device shall retain at minimum one information packet per I/O process for retransmission (i.e., the last information packet successfully transferred). Each GPP device shall support information packet lengths of 256 bytes or larger.

The information packet containing an SYNCHRONOUS PACKET TRANSFER REQUEST message shall contain only one interface logical element, a message interface logical element.

The SYNCHRONOUS PACKET TRANSFER REQUEST exchange may result in assymetric values between the exchange participants. That is, the initiating controller may establish fewer and shorter information packets than the target controller. This assymetry permits each unit to select the size of information packets it receives and the number of information packets it retains.

In the absence of an explicit packet offset count agreement for a value larger than one, the implied packet offset count agreement shall be 1 for both initiating controllers and target controllers. At the end of an exchange, both the initiating controller and target controller agree to maintain, for retransmitting only, a minimum number of information packets (greater than or equal to one).

In the absence of an explicit maximum information packet length agreement for a value larger than 256, the implied maximum information packet length agreement shall be 256 for both initiating controllers and target controllers. At the end of an exchange of messages, both the initiating controller and target controller agree to maintain information packets of at least the negotiated size. The minimum value for this field shall be 256.

Table 16. SYNCHRONOUS PACKET TRANSFER REQUEST Message Format

Byte	Value	Description
0	01h	Extended Message Format Code
1	0Ah	Extended Message length
2	07h	Extended Message Code
3	00h	Reserved
4	xxh	Originator's Packet Offset Count
5-6	xxxxh	Oringiator's Maximum Information Packet Length
7	00h	Reserved
8	xxh	Responder's Packet Offset Count
9-10	xxxxh	Responder's Maximum Information Packet Length
11	00h	Reserved

The originator's packet offset count represents n+1 information packets that the originator requests be retained by the responding GPP device. A maximum of 256 information packets may be retained per I/O process at the end of an SPTR negotiation (packet offset count value of 255). For initiating controllers which implement the MODIFY DATA POSITION message, there may be an additional requirement to retain additional information packets to support that message..

The originator's maximum information packet length that the originator requests be retained by the responding GPP device.

The responder's packet offset count represents n+1 information packets that the responder requests be retained by the originating GPP device. A maximum of 256 information packets may be retained per I/O process at the end of an SPTR negotiation. This field shall be set to zero in the first message sent by an originator in a SPTR exchange.

The responder's maximum information packet length that the responder requests be retained by the originating GPP device. This field shall be set to zero in the first message sent by an originator in a SPTR exchange.

The synchronous packet transfer agreement becomes invalid after any condition which may leave the packet transfer agreement in an indeterminate state such as:

- after a reset event
- after a DEVICE RESET message and
- after a power cycle.

In addition, an GPP device may initiate an SPTR message exchange whenever it is appropriate to negotiate a new packet transfer agreement.

The SPTR message exchange establishes the minimum number of information packets and the maximum length of each information packet that the originating GPP device and the responding GPP device shall independently maintain to permit retransmitting information packets in case of physical or logical errors in information packets. (See the RESEND PREVIOUS INFORMATION PACKET message.)

The originating device (the device that sends the first of the pair of SPTR messages) sets its values according to its capabilities to retain information packets on a per I/O process basis. The originator shall set the responder's values to zero. If the responding device can also retain information packets at the originator's requested level, the responder returns the same values in its SPTR message in the originator's fields. If it requires a smaller packet offset count or a shorter maximum information packet length, it substitutes an lower value in the appropriate originator field(s) of its SPTR message. Independently, the responder to the original SPTR message fills in the responder's packet offset value and the responder's maximum information packet length value.

The originating device analyzes the response and may initiate another SPTR message with the same or still lower originator's values. If the values in the response message are identical to those transmitted in the previous originator message, the negotiation ends with this agreement. The originator analyzes the responder's packet offset count value and the responder's maximum information packet length value. If the originating device can also retain information packets at the responder's level, it returns the same values in its SPTR message in the responder's fields. If it requires a smaller retention count or a shorter information packet maximum length, it substitutes an lower value in the appropriate responder's field(s) of its SPTR message.

The negotiation shall end when both devices exchange the same packet offset value and maximum information packet length value or the packet offset value reaches 0 and the maximum information packet length reaches 256, whichever occurs first, in both sets of fields.

Each device shall retain the number of information packets of less than or equal to the maximum information packet length for retransmission, as agreed upon in the most recent SPTR exchange. If there is any error in the negotiation process, the packet off set count agreement shall be 1 and the maximum information packet length shall be 256 after the I/O process for both pairs of values.

5.2.17 TERMINATE I/O PROCESS

The TERMINATE I/O PROCESS message is sent from the initiating controller to the target controller to advise the target controller to terminate the I/O process without corrupting the logical data of the logical unit. Upon successful receipt of this message the target controller shall terminate the identified I/O process as soon as possible and return I/O PROCESS TERMINATED status. The sense key shall be set to NO SENSE and the additional sense code and the additional sense code qualifier shall be set to I/O PROCESS TERMINATED. The TERMINATE I/O PROCESS message shall not affect pending status, data, and commands for other queued or active I/O processes. However, continued execution and status of other I/O processes queued for the H_C_x_y nexus may be affected by the queue error recovery option specified in the control mode page parameters.

Table 17. TERMINATE I/O PROCESS Message Format

Byte	Value	Description
0	01h	Extended Message Format Code
1	03h	Extended Message length
2	08h	Extended Message Code
3-4	xxxxh	Last IP to Process

The last IP to process field indicates the last IP that the initiating controller requires that target controller process for the terminated I/O process. The target controller shall clear all information packet with a larger value that the value of this field. If processing has progressed beyond this value, the target controller shall immediately terminate all processing for the I/O process and enter the auto-contingent allegiance state for the logical unit.

This message is normally used by the initiating controller to stop a lengthy read, write, or verify operation when a higher priority I/O process is available to be executed.

The information packet containing an TERMINATE I/O PROCESS message shall contain only one interface logical element, a message interface logical element.

If the I/O process that is being terminated has a logical data transfer associated with it, the valid bit in the sense data shall be set to one and the information field shall be set as follows:

1. If the command descriptor block specifies an allocation length or parameter list length in bytes, the information field shall be set to the difference (residue) between the transfer length and the number of bytes successfully transferred.

2. If the command descriptor block specifies a transfer length field, the information field shall be set as defined in the REQUEST SENSE command.

If the I/O process being terminated has no data transfer associated with it the target controller shall set the valid bit in the sense data to zero and terminate the I/O process with I/O PROCESS TERMINATED status. The sense key shall be set to NO SENSE and the additional sense code and the addi;tional sense code qualifier shall be set to I/O PROCESS TERMINATED.

When any error condition is detected while terminating an I/O process the target controller shall ignore the TERMINATE I/O PROCESS message and terminate the I/O process with the appropriate error status and sense data for the error condition.

If the target controller completes all processing for an I/O Process (i.e., all data has been read, written, or processed) and a TERMINATE I/O PROCESS message is received before the I/O process is terminated, the target controller shall ignore the TERMINATE I/O PROCESS message and terminate the I/O process in the normal manner.

If the target controller receives a TERMINATE I/O PROCESS message before any command descriptor block is transferred or the information process identifies an H_C_x nexus that does not have an active or queued I/O process, the target controller shall set the valid bit in the sense data to zero and terminate the I/O process with I/O PROCESS TERMINATED status. The sense key shall be set to NO SENSE and the additional sense code and the additional sense code qualifier shall be set to I/O PROCESS TERMINATED.

If the affected I/O process is in the I/O Process queue (an H_C_x nexus for untagged queuing or an H_C_L_Q nexus for tagged queuing) and has not started execution, the target controller shall record the event with the queued I/O process and wait until the I/O Process is in the active queue (started executing) then terminate the I/O process. The target controller shall terminate the I/O process with I/O PROCESS TERMINATED status. The sense key shall be set to NO SENSE and the additional sense code and additional sense code qualifier shall be set to I/O PROCESS TERMINATED. The valid bit shall be set to zero.

The TERMINATE I/O PROCESS message provides a means for the initiating controller to request the target controller to reduce the transfer length of the current command to an amount that has already been transferred. The initiating controller can use the sense data to determine the actual number of bytes or blocks that have been transferred.

5.2.18 TRANSFER READY

The TRANSFER READY message is sent to indicate that an active I/O process requires additional information packets. The message is used to direct an initiating controller or a target controller to continue sending information packets for the identified I/O process. This message does not make a specific request for any type of interface logical element. The target controller has processed all transmitted information packets and is waiting for more to be transmitted (i.e., an underrun condition).

Table 18. TRANSFER READY Message Format

Byte	Value	Description
0	01h	Extended Message Format Code
1	03h	Extended Message length
2	09h	Extended Message Code
3-4	xxxxh	Sequence Number of the Last IP Processed

The Sequence Number of the Last IP Processed field indicates the last information packet sequence number processed. The I/O process is cannot continue until additional information packets are transferred.

The information packet containing a TRANSFER READY message shall contain only one interface logical element, a message interface logical element.

In response, the initiating controller or target controller picks the path it wants to use based on agreements with the initiating controller for the I/O process and sends the additional information packets.

6.0 GPP FC-PH Usage

This section describes the physical characteristics of Fiber Channel available to SCSI-3 Generic Packetized Protocol (GPP). Fiber Channel specifies the physical requirements for interconnection. For specific details on the Fiber Channel, see the Fiber Channel standards. See Section 2 for the complete titles and document numbers for these standards.

Serial interface operations are incompatible with the parallel interface operations. This section and "7.0 GPP FC-PH Exchange Protocol" on page 36 specifies the use of Fiber Channel as a transport layer for GPP. The parallel interface is described in "8.0 ANNEX A. SPI Physical Characteristics (Normative)" on page 41 and "9.0 ANNEX B. GPP SPI Information Transfer (Normative)" on page 42. A logical system may be constructed using serial transport layers, parallel transport layers, or a combination of the two.

Mandatory requirements for physical operation placed on initiating controllers and target controllers include:

- All GPP ports shall implement an active initiator mode and an active target mode.
- Each initiating controller shall declare itself as a processor device class in response to the INQUIRY command.
- A GPP port which normally acts in initiator role, when selected to receive an information packet to start an I/O process, shall enter target mode.

6.1 Fiber Channel Physical Description

GPP devices attached to Fiber Channel use a 2-fiber connection called a link. The Fiber Channel FC-PH standard defines implementations using different types of conductors.

All Fiber Channel operations are conducted through primitive signals and primitive sequences sent on the fibers or through the transfer of sequences.

The physical port in each GPP device is called an N_Port (for Node Port). A GPP device may have more than one N_Port.

GPP devices may connect to:

- a point-to-point topology (Mandatory)
- a switched point-to-point topology (Mandatory)

  There are two versions:
  - Cross Point Switch Topology (FC-XS)
  - Distributed Fabric Topology (FC-LT).

  Fabric elements may exist which permit inter-operability between different fiber types and different transmission rates.

- arbitrated loop topology (Optional).

  A GPP device capable of being used in an arbitrated loop topology shall also be capable of operating in any of the point-to-point topologies. The arbitrated loop topology requires extensions beyond the basic N_Port.

Each GPP device shall support Fiber Channel classes of service as follows:

- Class 2 service (mandatory)
- Class 1 service (optional)

The physical port in a fabric element is called an F_Port. When a fabric is present, each N_Port is connected by a single link to exactly one F_Port.

The relative flow of information in SCSI is outbound from initiating controllers and inbound from target controllers. Because of the full-duplex nature of Fiber Channel operations, the actual difference in the direction of information flow can only be determined by examining the information packet content. The physical transport layer is symmetrical in all transfer operations.

6.2 FC-PH Encapsulation of an Information Packet

The unit of work between an initiating controller or a target controller is an information packet. The FC-PH construct is called a sequence. The FC-PH sequence delivery protocol forms the building blocks for conducting GPP logical operations on the Fiber Channel.

6.2.1 FC-PH Sequence Management

The Fiber Channel FC-PH interface to an initiating controller or a target controller is an architectural element called an FC-4 level. When an information packet (sequence) is to be transferred, certain mandatory and optional information must be supplied along with the information packet contents. Annex R of the FC-ph standard identifies the information which is supplied in addition to the information packet contents. In general, the information supplied is used to fill in FC-PH frame header fields. The negative displacements shown for the frame header fields keeps all GPP common fields at the same offset for all interfaces. Table 19 on page 32 shows the Fiber Channel FC-PH frame header fields.

Table 19. Frame Header Field (Serial)

Bit Byte	7-0
-24	Routing Bits (in R_CTL) Bits 7-4
-24	Information Categort (in R_CTL) Bits 3-0
-23 to -21	(MSB) Destination ID (D_ID) (LSB)
-20	Reserved
-19 to -17	(MSB) Source ID (S_ID) (LSB)
-16	Data Structure Type (Type)
-15 to -13	(MSB) Frame Control (F_CTL) (LSB)
-12	Sequence ID (SEQ_ID)
-11	Data Field Control (DF_CTL)
-10 to -9	(MSB) Sequence Count (SEQ_CNT) (LSB)
-8 to -7	(MSB) Originator Exchange ID (OX_ID) (LSB)
-6 to -5	(MSB) Responder Exchange ID (RX_ID) (LSB)
-4 to -1	(MSB) Parameter (LSB)

The use made of the individual frame header fields by GPP is specified below for Device_Data sequences. The Fiber Channel standard also specifies the content of Link_Control sequences based on the content of the received Device_Data frames. GPP does not specify any restriction on field use in Link_Control sequences other than that required by the Fiber Channel standard.

The requirements and recommendations below provide high-level Exchange/Sequence/Frame construction guidance for the GPP FC-4 level. To permit interoperability of GPP devices, this level of detail is required to properly construct the GPP FC-4 which operates between the SCSI-3 Logical Element and the FC-PH FC-2 functions in each N_Port.

6.2.1.1 R_CTL (Offset -24, Bits 7-4)

All sequences shall set this field to 0000b indicating Device_Data for all GPP information packets. This specifies a Device_Data frame with uncategorized information as defined by FC-PH. No other value shall be used.

6.2.1.2 R_CTL (Offset -24, Bits 3-0)

All sequences shall set this field to 0010b indicating an unsolicited control type for all GPP information packets. This specifies a Device_Data frame with Unsolicited Control information as defined by FC-PH. No other value shall be used.

6.2.1.3 D_ID (Offset -23 to -21)

This field shall be set to the Destination Identifier value of the GPP device to receive an information packet. This value is acquired through N_Port login with each GPP device in the logical system. N_Port login is a function of Fibre Channel and is not discussed in this standard. See the Fiber Channel standard.

There is a one-to-many mappint from the original Target GPP address to the FC-PH D_ID field values. The FC-4 level shall select one D_ID based on the current state of multiple pathing for the I/O process.

6.2.1.4 Reserved (Offset -20)

All sequences shall set this field to 0000 0000b in all frames containing GPP information packet data.

6.2.1.5 S_ID (Offset -19 to -17)

The FC-4 level shall set this field to the Native Address Identifier value assigned to one of, or the only, N_Port of the sending GPP device N_Port during fabric login (or N_Port login for a point-to-point topology). This value is normally acquired through fabric login of each GPP device in the logical system. Fabric login is a function of Fibre Channel and is not discussed in this standard. See the Fiber Channel FC-PH standard.

6.2.1.6 TYPE (Offset -16)

All sequences shall set this field to 0000 1001b to indicate that the sequence contains a GPP information packet.

6.2.1.7 F_CTL (Offset -15 to -13)

Exchange Context This field shall be set to 0b in all Device_Data frames of each sequence. Two exchanges are used in GPP between each pair of GPP devices so that each GPP device has initiative to transmit to any other communicating GPP device at any time. These exchanges are started after N_Port Login.

Sequence Context This field shall be set to 0b in all Device_Data frames of each sequence.

First_Sequence This field shall be set to 1 in each frame of the first sequence of an exchange. The Exchange Initiator is responsible for setting this field. It shall be set to zero in all subsequent sequences of the same exchange.

Last_Sequence If the Exchange Originator requires termination of the exchange, this field shall be set to one only in the last Device_Data frame of the last sequence.

Since the exchange protocol is open-ended in GPP, the value in this field is normally 0 in all sequences of the exchange.

End_Sequence For each sequence, this field shall be set to 1b on the last frame of the sequence. The field value shall be 0b on all other Device_Data frames of the sequence. The value is assigned by the FC-4 level or by the FC-2 level.

End_Connection (E_C) This field is meaningful only in Class 1 connections and only in the last Device_Data frame of a sequence. Use of this field as defined in the Fiber Channel FC-PH standard has no effect on the GPP protocol.

Chained_Sequence (C_S) This field shall be 0b for all sequences.

Sequence Initiative This field shall be set to 0b in all sequences. The two exchange protocol used with GPP requires that each Exchange Originator retain sequence initiative.

Support for exchange ID management is not a required function of GPP. Setting any of the next three fields to a value or combination of values which requires exchange ID management to be performed is outside the scope of GPP. If exchange ID management is required, it shall be used outside the GPP protocol according to the Fiber Channel standard.

New X_ID Assigned Refer to the Fiber Channel FC-PH standard for use of this field. Its use has no effect on the GPP exchange protocol.

Invalidate X_ID Refer to the Fiber Channel FC-PH standard for use of this field. Its use has no effect on the GPP exchange protocol.

Continue Sequence Condition Refer to the Fiber Channel FC-PH standard for use of this field. Its use has no effect on the GPP exchange protocol.

Abort Sequence Condition This field is meaningful only on the first Device_Data frame of the first sequence of an exchange. The value shall be set to 00b. The nature of GPP makes it impossible to process partial information packet transfers.

The sequence responder, after detecting an error in a sequence shall discard the sequence and set the value of this field to 01b in the ACK frame of the Link_Response frame. No processing shall be performed on the information packet so it can be retransmitted. The imbedded information packet ordering prevents out of order processing by the destination if a later sequences is successfully transmitted before retransmission occurs.

Reserved for Exchange Reassembly This field shall be set to 0b in all Device_Data frames of each sequence.

Fill Data Bytes This field shall be set to 00b in all Device_Data frames of each sequence. All information packets are an even multiple of 4 bytes long. The value is assigned by the FC-4 level or by the FC-2 level.

6.2.1.8 SEQ_ID (Offset -12)

The sequence initiator shall assign a unique value relative to the current exchange. Any sequence ID which is not active may be used as the SEQ_ID value. See FC-PH for specific sequence ID assignment requirements. The value is assigned by the FC-4 level or by the FC-2 level.

6.2.1.9 DF_CTL (Offset -11)

The SCSI-3 GPP protocol does not require a device header field. That field shall be 00b. Setting any of the other non-reserved fields to a value other than 0b is outside the scope of GPP. If they are required, they shall be used outside the GPP protocol according to the Fiber Channel standard. That is, the content of the optional headers is not presented to or supplied by the the GPP logical element.

6.2.1.10 SEQ_CNT (Offset -10 to -9)

This is a sequential count of frames transmitted in a sequence, starting at 0000h and incrementing by one for each subsequent frame of the sequence. The count is reset to 0000h for each information packet. Since the maximum information packet length is 65532, a wrap on SEQ_CNT cannot occur in GPP. The value is assigned by the FC-4 level or by the FC-2 level.

6.2.1.11 OX_ID (Offset -8 to -7)

A two-byte value assigned by the Exchange Originator. The value is assigned by the FC-4 level or by the FC-2 level.

6.2.1.12 RX_ID (Offset -6 to -5)

A two-byte value assigned by the Exchange Responder. The value is assigned by the FC-4 level or by the FC-2 level.

6.2.1.13 Parameter (-4 to -1)

This field shall be set to FFFFFFFFh. Each sequence is reassembled using the SEQ_ID || SEQ_CNT fields.

6.2.2 GPP FC-PH Response Frames and Signals

FC-PH provides well-defined responses to Device_Data sequencess received as part of an exchange. Fiber Channel provides for FC-4-to-FC-4 and/or N_Port-to-N_Port notification of transmission status for Class 1 and Class 2 service. Fiber Channel specifies the exact sequence frames shall be transmitted, the normal response, and the abnormal responses.

6.3 FC-PH Events

FC-PH has two asynchronous events: the unexpected disconnect event and the unsuccessful information packet transfer event. Each event establishs a condition which causes a logical element to perform certain actions.

6.3.1 FC-PH Unexpected Disconnect Event

GPP devices normally expect IDLE primitives to begin on its receive fiber after one of the following occurs:

1. after the transfer of one or more information packets.
2. after one of the FC-PH link protocol sequences.

All other occurrences of a unexpected state cause the GPP device to report an Unexpected Disconnect Event to the sending logical element. The sending GPP device shall handle this as an unsuccessful information packet transfer condition.

6.3.2 FC-PH Unsuccessful Information Packet Transfer Event Processing

For an unsuccessful information packet transfer event, the sending GPP device:

1. may attempt to retransmit the information packet in error up to a limit determined by the sending FC-4.
2. if operating in target mode for the I/O process, the sending GPP device shall terminate the I/O process after the specified number of retries by clearing all pending interface logical elements for the I/O process and preparing sense data.
3. if operating in initiator mode for the I/O process, the sending GPP device shall abort the I/O process. The target controller prepares sense data. It is recommended that the initiator then request sense data from the target

controller if an autosense ILE is not provided by the target controller for the I/O process.

7.0 GPP FC-PH Exchange Protocol

An active serial interface consists of a continuous stream of encoded words transmitted along one fiber between two ends of a link. This stream of transmission words is

- between two N_Ports in a point-to-point topology.
- between an N_Port and its associated F_Port in a switched point-to-point topology.
- between two F_Ports in a switched point-to-point topology between two fabric elements.
- between two N_Ports or between an N_Port and an F_Port in an arbitrated loop topology.

The receiving port is not aware that it is about to receive an information packet. The sending GPP port may not be aware of the present state of the destination GPP port.

The information transfer protocol assumes delivery of complete information packets and recovers on that boundary when that assumption proves incorrect. That is, error free information packet delivery is not guaranteed at the time of transmission in some topologies. The exchange protocol provides for positive responses from the opposite end of the link at some time after transmission. The GPP exchange protocol, as well as other FC-PH mappings, consider a fabric a transparent element in the protocol.

For any point-to-point topology, the source has complete control of when an information packet is transmitted. The destination also has complete control of when an information packet is transmitted on the reverse fiber.

For each information packet which contains a logical error, but which is otherwise received error free, the receiving GPP FC-PH port shall send an acknowledgement response for the frames of a sequence. The receiving GPP logical unit shall not process this information packet. Any logical error is reported using an information packet sent in the opposite direction.

For each information packet received with a physical error in any frame, such as a CRC error, the receiving GPP FC-PH port shall not present any portion of the information packet to a GPP logical element. The N_Port shall not process this sequence further and shall not retain it.

7.1 Introduction to the GPP FC-4 Exchange Protocol

Two GPP devices which require interaction must have a means to communicate:

- The lowest level Fiber Channel protocol involves individual frame transmission and reporting of buffer-to-buffer and end-to-end flow control. This level of operation is transparent to GPP.
- The next level of protocol is sequence construction, transmission, and reassembly. GPP provides the set of bytes in the form of an information packet to be transmitted in a sequence to the FC-4 level, and the receiving FC-4 provides a reassembled information packet to the GPP logical element at the receiver. The content of the sequence payload is not known by, and therefore is not interpreted by, FC-PH or the GPP FC-4 level. The actual sequence transmission details are managed by FC-PH, but the GPP FC-4 level is the lowest level of interaction between GPP and FC-PH.
- The next level protocol in FC-PH is the exchange which consists of the transfer of one or more sequences and the corresponding responses. The actual management details are performed by the FC-4 level and FC-PH.

  An exchange permits transfer of Device_Data frames in only one direction at a time (by the owner of sequence initiative). To be able to effectively use the entire bandwidth between two N_Ports, two or more exchanges are required (and assuming sufficient total workload). GPP specifies the use of two exchanges per pair of N_Ports in a logical system, each with sequence initiative in the opposite direction. Additional pairs of exchanges may be active if required. See Figure 2 on page 37.

  There is no FC-PH tie between these exchanges, and there is no GPP tie between these exchanges. The exchanges are for the owners of exchange sequence initiative to use as needed. Since only one frame can be transmitted at a time, use of multiple sets of exchanges is not required. Sequences may be streamed up to the N_Port login limit for Open Sequences per Exchange.

  The SCSI concept of an I/O process is not related to an exchange. The exchange is the vehicle used to transmit sequences between two N_Ports. Two consecutive sequences in the same exchange between two N_Ports may have no relation to each other, except that the logical unit(s) and initiating controller functions reside inside the respective N_Ports.

```
EXCHANGE w/          GPP Sequences ------------->
Sequence
Initiative     <--------------- FC-PH Responses
+---------+                                         +---------+
|         |                                         |         |
|         |-------- Outbound Information -------->|         |
|INITIATOR|                                         | TARGET  |
| (N_Port)|<------- Inbound Information ----------| (N_Port)|
|         |                                         |         |
+---------+                                         +---------+
               <--------------- GPP Sequences       EXCHANGE w/
                                                    Sequence
                     FC-PH Responses ----------->   Initiative
```

Figure 2. GPP Exchange Protocol

Each initiating controller N_Port may have one or more open exchanges per target controller N_Port in its logical system. Each target controller N_Port may have one or more open exchanges per initiating controller N_Port in its logical system. If it is necessary for one of the N_Ports to terminate some or all of its exchanges, it may do so. Since sequence initiative is not transferred within one exchange, the N_Port must reestablish at least one exchange to continue or start an I/O process.

Also, since there is no affinity between a SCSI GPP I/O process and any FC-PH exchange, multiple N_Port pairs may be used to conduct a single I/O process when multipathing is supported. This permits the use of multiple fabrics for highly available systems.

- The highest level of Fiber Channel protocol which affects GPP is N_Port login. This protocol is specified in FC-PH and executed outside GPP, but it affects GPP. Without successful N_Port login, exchanges for the purpose of transferring GPP information packets cannot be established with another N_Port. See Figure 3 on page 37.

```
N_Port Login               Determines set of Accessible N_Ports

N_Ports in set             Potential SCSI Logical System
that accept GPP
FC-PH Type (00001001b)

Establish Exchange(s)      Desired GPP Devices

Information Packet         With one Other N_Port
Causes Sequence
Transmission

Sequence Transmitted       Assemble, Transmit, and Reassemble
as Individual Frames       an Information Packet
```

Figure 3. Hierarchy of Protocols Affecting GPP

7.2 GPP FC-PH Exchange Protocol

This section specifies how GPP uses Fiber Channel and does reflect the full capabilities of Fiber Channel. The exchange protocol is composed of the following elements:

- Sequences
- Exchanges

7.2.1 FC-PH Sequences

A GPP information packet maps one-to-one with a FC-PH sequence. Once N_Port login is complete and an exchange is established in both directions between two communicating N_Ports, information packets may be transferred between them using FC-PH sequences.

7.2.2 FC-PH Exchanges

An exchange estsablishes a logical connection between two N_Ports. An exchange is established by a sending N_Port after N_Port login. The originator of an exchange identifies its exchange by the concatenation of S_ID || D_ID || OX_ID (of S_ID). The responder assigns a responder exchange identifier (RX_ID) to the exchange.

In the reverse direction, a second exchange is established with the exchange identification D_ID || S_ID || OX_ID (of D_ID). The responder assigns a responder exchange identifier (RX_ID) to this exchange.

In normal practice, the exchanges may be established while attempting to negotiate a synchronous packet transfer agreement using information packets containing the SYNCHRONOUS PACKET TRANSFER REQUEST message (SPTR). This negotiation is normally performed only once per N_Port login. The negotiation is terminated when either N_Port logs out (implicitly or explicitly) with the other.

For example, a GPP initiating controller N_Port is assigned a Native Address Identifier of 777777h and a GPP target controller N_Port is assigned a Native Address Identifier of 111111h during their respective fabric logins. Figure 4 on page 39 shows a portion of an I/O Process which establishes the pair of bi-directional exchanges used with GPP. If the SPTR negotiation is not used, the exchanges can be established as part of any I/O process which requires information packets in both directions. Refer also to Figure 2 on page 37.

INITIATING CONTROLLER NA_ID = 777777h	TARGET CONTROLLER NA_ID = 111111h
1. Develop Information Pkt in Logical Element 2. Select Target Controller 111111h 3. Assign OX_ID = 2222h 4. Transmit Sequence for D_ID = 111111h S_ID = 777777h OX_ID = 2222h RX_ID = 0000h (Unknown)	
	5. Sequence Received Assign Local RX_ID = 4444h 6. Send acknowledgement for sequence D_ID = 777777h S_ID = 111111h OX_ID = 2222h RX_ID = 4444h 7. Deliver Information Pkt to Logical Element 8. Develop Information Pkt Response in Logical Element 9. Use S_ID from Sequence = 777777h 10. Assign OX_ID = 6666h 11. Transmit Sequence for D_ID = 777777h S_ID = 111111h OX_ID = 6666h RX_ID = 0000h (Unknown)
12. Response Sequence Received Assign Local RX_ID = 8888h 13. Send acknowledgement for sequence D_ID = 111111h S_ID = 777777h OX_ID = 6666h RX_ID = 8888h 7. Deliver Information Pkt to Logical Element	
...	...

Figure 4. Example of Establishing Bi-directional Exchanges

Since all GPP information is transferred in information packets, normal operation can now proceed. The initiating controller develops information packets which start I/O processes and the target controller and its associated logical unit(s) act on the I/O processes and prepare the appropriate responses as required by the SCSI-3 Architecture Model and the SCSI-3 command sets.

7.2.3 Transport Layer Independence

Since all information packets are fully self-identifying, relative to the I/O process with which they are associated, there is no connection between these two pairs of exchanges. Additional exchanges may be set up with the same or a different set of N_Ports between the same initiating controller and target controller. Each node may use any available exchange to transfer a sequence to the other.

This lack of dependence on the physical transport details permits the multiple port option of GPP to be used with any type of I/O process within a single fabric or spanning two or more fabrics. (See Figure 5 on page 40.)

```
NODE A                                                 NODE B
with 2 N_Ports                                         with 2 N_Ports
_________________                                     ____________________

EXCHANGE w/         GPP Sequences ------------->
Sequence
Initiative    <--------------- FC-PH Responses
(OX_ID/RX_ID)
(2222h/4444h)
+---------+                                            +---------+
|         |                                            |         |
|         |-------- Outbound Information ------->|         |
|INITIATOR|                                            | TARGET  |
| (N_Port)|<------- Inbound Information ---------| (N_Port)|
|         |                                            |         |
+---------+                                            +---------+
                                                       (6666h/8888h)
                                                       (OX_ID/RX_ID)
              <--------------- GPP Sequences           EXCHANGE w/
                                                       Sequence
                    FC-PH Responses ----------->       Initiative


EXCHANGE w/         GPP Sequences ------------->
Sequence
Initiative    <--------------- FC-PH Responses
(OX_ID/RX_ID)
(2020h/4040h)
+---------+                                            +---------+
|         |                                            |         |
|         |-------- Outbound Information ------->|         |
|INITIATOR|                                            | TARGET  |
| (N_Port)|<------- Inbound Information ---------| (N_Port)|
|         |                                            |         |
+---------+                                            +---------+
                                                       (6060h/8080h)
                                                       (OX_ID/RX_ID)
              <--------------- GPP Sequences           EXCHANGE w/
                                                       Sequence
                    FC-PH Responses ----------->       Initiative
```

Figure 5. GPP Logical System with Multiple Ports

8.0 ANNEX A. SPI Physical Characteristics (Normative)

This section specifies the parallel physical transport layer available to GPP. Each transport layer standard defines the physical requirements for interconnection. For specific details on the SCSI-3 parallel bus, see the SCSI-3 Parallel Interface (SPI) standard. See Section 2 for the complete titles and document numbers for these standards.

Electrical Type	Connector Type			
	Shielded		Non-Shielded	
Single Ended (6 M)	A1	A2	A1	A2
8-Bit (A Cable)				
Termination A1	-	-	-	-
Termination A2	X	-	X	-
16-Bit (P Cable)				
Termination A1	-	-	-	-
Termination A2	X	-	X	-
32-Bit (P+Q Cable)				
Termination A1	-	-	-	-
Termination A2	X	-	X	-
Bus Width Intermix				
Termination A1	-	-	-	-
Termination A2	-	-	-	-
Differential (25 M)				
8-Bit (A Cable)				
5 MHZ	-	-	-	-
Fast Synchronous	X	-	X	-
16-Bit (P Cable)				
5 MHZ	-	-	-	-
Fast Synchronous	X	-	X	-
32-Bit (P+Q Cable)				
5 MHZ	-	-	-	-
Fast Synchronous	X	-	X	-
Bus Width Intermix				
5 MHZ	-	-	-	-
Fast Synchronous	-	-	-	-

Figure 6. Parallel Transport Layer Options

8.1 Parallel Interface Physical Description

The implementer has a selection of busses, shown in Figure 6 on page 41, from the SPI standard. The "X" identifies valid GPP options.

Single ended and differential implementations are mutually exclusive. GPP does not support intermixing 8-bit, 16-bit, and 32-bit interfaces. All GPP devices which are principally initiating controllers shall supply terminator power as required for termination option A2.

9.0 ANNEX B. GPP SPI Information Transfer (Normative)

For GPP, the DATA OUT phase is used for all information packet transfers across a SCSI_3 Parallel Interface. Interpretation of the data in the information packet shall not be permitted which halts the transfer of the information packet on the interface. That is, GPP permits no byte-by-byte processing.

For SPI, the initial information packet transfers are in asynchronous data transfer mode. GPP devices shall not have an interlock mechanism to identify individual bytes for protocol or parity errors as they are transferred. Only the actual transfer time may be slower than in synchronous data transfer mode. Parity errors are handled by retransmitting the information packet.

The information transfer phases defined for the SCSI-3 Interlocked Protocol (SIP) may be used as an alternate mode of operation for GPP devices operating a parallel bus. After detecting a SIP device on parallel bus, a GPP device may change its definition to the SIP operating mode for that device. If such a change occurs, the SIP standard shall control the protocol between the SIP operating mode level devices.

9.1 GPP Implemented SPI Bus Phases

GPP uses the following SIP phases with the SPI:

- BUS FREE phase
- ARBITRATION phase
- SELECTION phase
- SETUP phase
- DATA OUT phase

All GPP devices shall shall implement the phases named above. The default mode of information transfer shall be asynchronous transfer mode, at maximum DATA BUS width for the parallel interface. Full width transfers result from disallowing intermix of GPP devices with different DATA BUS widths. (See Annex A.) A GPP device shall implement synchronous transfer mode.

For a single connection between an initiator and a target, information packets for various I/O processes can be transferred unidirectionally. All information packet transfers are from the selecting device to the selected device.

Each GPP device establishes a physical path between an initiator and a target using the SELECTION phase. Information packet transfers are transferred, uninterpreted, to the logical element for later interpretation.

9.1.1 GPP versus SIP Operating Mode

The GPP operating mode and the SIP operating mode for information transfer are mutually exclusive. It shall be an error to attempt a phase for another operating mode within the connect or a later connection for an I/O process.

If the first information transfer phase of a connect is a DATA OUT information transfer phase, the target is stating its intent to operate in GPP mode. If the first information transfer phase of a connect is not a DATA OUT information transfer phase, then the target is stating its intent to operate in SIP mode.

Since ports in target role control the selection of information transfer phases, a GPP target shall initiate a procedure, at power-on or after a reset event, which declares its intent to operate in GPP mode. The target selects a GPP address. The first attempt at information transfer, after power-on reset or after processing the reset condition by a target, shall contain a request to negotiate synchronous data transfer by expecting to use a DATA OUT information transfer phase.

The selected GPP or SIP device shall respond with one of the following:

1. If it does not implement target mode, it does not respond to selection, and the target discovers that it is a SIP device or it is powered off.

2. If it implements target mode, selection without attention, followed by the DATA OUT information transfer phase, the GPP devices transfer the information packet.

3. If it is a SCSI device that only operates in SIP operating mode and implements target mode, the SCSI device responds to selection without attention, the selecting GPP device shall assert ATN on the first assertion of REQ for the non-DATA OUT information phase and shall not assert the ACK signal.

If ATN is asserted on the first assertion of REQ, the target controller disconnects if it implements only GPP mode. With disconnection, the selecting SIP device shall release its ATN signal. The target controller has detected that it is possible to operate with that SCSI device address using SIP operating mode. The target controller then may decide to change its definition to SIP mode for that SCSI device to permit operation.

If the GPP device decides to switch to SIP mode it shall respond to subsequent selections by that SCSI device address. If the target controller implements only GPP mode, the target controller shall not attempt to use GPP mode with that SCSI device address until after a power-on reset, or a reset event has been processed. The selected SIP device detects an unexpected disconnect and reacts appropriately according to the SIP standard.

9.1.2 SIP BUS FREE Phase

The BUS FREE phase indicates that no SCSI device is actively using the bus and that it is available. Sometimes, a target reverts to BUS FREE phase to indicate an error condition that it has no other way to report. This is called as an unexpected disconnect event. GPP devices shall detect the BUS FREE phase according to the requirements specified in SIP for the BUS FREE phase.

9.1.3 SIP ARBITRATION Phase

The ARBITRATION phase allows one SCSI device to gain control of the SCSI bus so that it can initiate or resume an I/O process.

The procedure for an GPP device to obtain control of the bus is the same as specified in SIP.

9.1.4 SIP SELECTION Phase

The SELECTION phase allows:

1) an initiating controller to select a target controller to initiate some target function (e.g., READ or WRITE command).

2) an initiating controller to reconnect with a target controller after a disconnect.

3) a target controller to reconnect with an initiating controller after a disconnect.

The procedure to perform SELECTION is the same as in SIP except as follows:

- Selection shall be performed with ATN not asserted since the only information transfer phase is the DATA OUT.

  The GPP target controller begins a DATA OUT information transfer phase.

  (**EDITORS NOTE** To SIP add text, now missing, about what to do when selection without ATN is attempted. That is, the target controller shall go to the BUS FREE phase and ignore the selection.)

- the single bit selection option shall not be implemented.

9.1.5 SIP SELECTION Time-out Procedure

The SELECTION time-out procedure is the same as required by SIP.

9.1.5.1 SIP Asynchronous Information Transfer

Asynchronous information transfer shall be implemented as described in SIP. It shall be used only to perform a synchronous data transfer negotiation using the appropriate infomation packets.

9.1.5.2 SIP Synchronous Data Transfer

Synchronous data transfer shall be implemented for use during the DATA OUT information transfer phase, except as noted above to negotiate the synchronous mode. It shall be used in this phase if a synchronous data transfer agreement has been established. The agreement specifies the REQ/ACK offset and the minimum transfer period. The default agreement is asynchronous data transfer. Synchronous data transfer shall be implemented as specified in SIP.

9.1.5.3 SIP Wide Data Transfer

Wide data transfer is mandatory if the DATA BUS width exceeds 8 bits, excluding parity. For wide data busses, the ports shall use a full-width data transfer for each REQ/ACK handshake during a DATA OUT information transfer phase. No wide data transfer agreement is required since connection of data busses of unequal widths shall not be supported. The default agreement is the maximum width of the implemented data bus and shall not be changed. Wide data transfers shall implemented as specified in SIP.

9.1.6 SIP DATA OUT Information Transfer Phase

The DATA OUT phase is the only SIP information transfer phase used with GPP. The phase is used to transfer information packets on the DATA BUS.

The C/D, I/O, and MSG signals used to identify the physical information transfer phase are identical to the DATA OUT phase defined in SIP. The port in target role controls these three signals and therefore controls all changes from one phase to another. The port in initiator role can request a DATA OUT phase by asserting the ATN signal at certain times. The target can disconnect by releasing the C/D, I/O, MSG, and BSY signals.

The DATA OUT information transfer phase uses one or more REQ/ACK handshakes to control the information transfer. Each REQ/ACK handshake allows the transfer of one full width bus of information.

9.1.7 SIP SETUP Phase

During one DATA OUT information transfer phase, the port in target role shall continuously envelope the REQ/ACK handshake(s) with unchanging C/D, I/O, and MSG signals. An information transfer phase ends when the C/D, I/O, or MSG signals change after the negation of the ACK signal. This is called the SETUP phase in SIP. Assertion of REQ is not required to signal the end of a DATA OUT information transfer phase. It shall be sufficient that when one of the MSG, C/D, or I/O signals changes the previous information transfer phase is ended.

After entering the SETUP phase, the target may prepare for a new phase by asserting or negating the C/D, I/O, and MSG signals. These signals may be changed together or individually. They may be changed in any order and may be changed more than once. A new information transfer phase does not begin until the REQ signal is asserted for the first transfer of the new phase.

A port in initiator role is allowed to anticipate a new DATA OUT information transfer phase, based on early information provided by changes in the C/D, I/O, and MSG signals while in the SETUP phase. However, the actual phase is not valid until the REQ signal is asserted.

9.2 SPI Events

SPI defines three asynchronous events:

1. the unexpected disconnect event;
2. the attention event;
3. and the reset event.

These events cause the GPP device to perform certain actions and can alter the phase sequence.

9.2.1 SPI Unexpected Disconnect Event

If an initiator detects a disconnect by the target at any other time than for the conditions below, the target is indicating the unexpected disconnect event to the initiator. The target may perform this disconnect independent of the state of the ATN signal. The initiator shall manage this event as an unsuccessful information packet transfer condition.

GPP initiators normally expect the BUS FREE phase to begin after one of the following occurs:

1. after a reset event is detected.
2. after the transfer of one or more information packets in the DATA OUT phase.
3. after failure of a target following selection to transfer more than 16 bytes of an information packet using the DATA OUT phase. The port in target role is indicating a busy condition. The initiator shall attempt the transfer at a later time.
4. after an unsuccessful selection.
5. after a selection of a SIP device (i.e., selection without ATN asserted). The initiating controller may then decide to operate with the target controller in SIP mode. If so, all subsequent communication follows the SIP standard.

All other occurrences of a BUS FREE state cause the initiator to report an Unexpected Disconnect Event to the Initiating Controller.

9.2.2 SPI Attention Event

The attention event allows a port in initiator role to inform a port in target role during a connection that it has an information packet to send. The port in target role gets this information packet by performing a DATA OUT phase.

The port in initiator role generates the attention event by asserting ATN as follows:

- not during the BUS FREE phase.
- before the ACK for the first byte of any SIP information transfer phase other than DATA OUT to indicate that the initiating controller does not operate in SIP mode. The initiating controller may then perform one of the following:
    - not assert ACK which shall cause the target controller to time out on the phase, abort the I/O process and go to the BUS FREE phase.
    - respond with ACK, continue to respond with ACK for additional assertions of REQ with 00H byte values and good parity during the same phase, and when the target controller switches to the MESSAGE OUT phase, send a SIP ABORT message to the target controller. This provides for compatability with SIP devices which still support SCSI-1 selection modes.
    - during a DATA OUT information transfer phase to indicate that it has an additional information packet to send during the same connection.

The port in initiator role shall assert the ATN signal at least two deskew delays before negating the ACK signal for the last byte transferred in a bus phase for the attention condition to be honored before the SETUP/BUS FREE sequence.

A port in target role shall respond to the attention event during a DATA OUT phase as follows:

1. If the ATN signal becomes true during a DATA OUT phase, the target shall complete the current information packet transfer, enter the SETUP phase, and enter a new DATA OUT phase.
2. the target shall not disconnect, except for an unexpected disconnect, before responding to the attention event.

9.2.3 SPI Reset Event

The reset event is used to clear immediately all GPP devices from the bus. This event shall take precedence over all other phases and events. Any GPP or SIP device may create the reset event by asserting the RST signal for a minimum of a reset hold time. During the reset event, the state of all bus signals other than the RST signal is not defined.

All GPP devices shall release all bus signals (except the RST signal) within a bus clear delay of the transition of the RST signal to true. The BUS FREE phase always follows the reset event. The reset event has the same effect as defined in SIP.

GPP devices shall implement the soft reset option in response to the reset event.

END OF DOCUMENT X3T9.2/91-013 R03