

Date : August 30, 1990

To : John Lohmeyer, Chairman, X3T9.2

From : Gary R. Stephens
IBM Corporation
D67E/060-1
9000 S. Rita Rd.
Tucson, Az 85719
(602) 799-2246

Subject: Logical Model for SCSI-3

The attached document is the third draft of the logical model for SCSI-3. There is a descriptive section which is not considered part of the text for inclusion into the SCSI standard. It gives some background material for the logical operation model.

Some functions which require multiple ports may only be economically possible in the serial physical interface market when cables and connectors are much smaller. Bridge controllers which have back panel space can use the functions described here to achieve controlled multiple port operation. Small devices will not generally have the rear panel space to attach the cables for multiple parallel ports. A combination including a bridge controller with a serial interface on the initiating controller side and parallel interface with information packets on the logical unit side may make a good combination (See X3T9.2/89-130 R02).

The concepts in the model are presented in a form similar to the models developed for the various device classes. This format provides a structure to name the concepts and functions and present them in a logical sequence.

The logical model is proposed for addition to Section 6 of the SCSI standard. Some glossary terms are proposed for addition to Section 3.1.

The command descriptions are not included, but an outline of them is provided at the end of the document. All of the commands are proposed as additions to Section 7 of the SCSI standard.

Gary R. Stephens

DESCRIPTIVE TEXT -- NOT PART OF THE MODEL

The following text provides an introduction to the logical operations model, but it is not part of the logical operations model.

There is no explicit provision in SCSI to prohibit a initiating controller from having more than one port attached to the same SCSI bus. There is no provision to prohibit a initiating controller from having ports on more than one SCSI bus.

There is no explicit provision in SCSI to prohibit a target controller from having more than one port attached to the same SCSI bus. There is no provision to prohibit a target controller from having ports on more than one SCSI bus.

The capabilities provided by such flexible attachment lack a defined control mechanism. How is a target controller to determine that two of its ports and at least one initiator on each port belong to the same initiating controller? As system availability and performance requirements increase, the ability to use an alternate path to communicate with a initiating controller to continue an I/O process becomes increasingly important.

If a primary path fails and an alternate path were available, especially on an alternate SCSI bus to the same initiating controller, a target controller would be able to continue an I/O process on an alternate path. After completion of the I/O process on the alternate path, the original path could be reset with, usually, little recovery processing in the initiating controller. Such an I/O process is an example of a dynamic path process.

The first hints of the need for a defined protocol for handling multiple paths comes in the definition of the soft reset alternative in response to the RESET condition. The hard reset alternative is losing its appeal as sophisticated systems come to the market. The soft reset alternative essentially clears the bus. Note that the RESET condition applies to initiating controllers and target controllers. The intent is to disturb only the current I/O process. A RESET condition occurring when the SCSI bus is in the BUS FREE state should not cause any I/O process to be aborted.

The presence of the reserve and third party reserve functions in SCSI-2 recognizes that a logical device may need controlled access from one or two paths on the same SCSI bus to the exclusion of other paths. Up to two initiators may operate exclusively with a logical unit or some extents within a logical unit. There is no mechanism for enabling a third party reserve on a second SCSI bus. The result of a third party reserve is approximately equivalent to the model description of assignment to two path groups each having one path as described by the logical model.

Dynamic path operations require a change in the management of pointers from SCSI-2. For dynamic path operations, the set of pointers becomes tied to the path group and not the initiator on which the connect is made. This means that with reconnection on a second initiator in the path group, the second initiator retrieves the saved pointers for the I/O process from a common pool of pointer sets, rather than from within the first initiator itself. The pointers are removed from the exclusive domain of a single initiator and are made a resource of the initiating controller.

When multiple initiating controllers share one SCSI bus, there may be times when potentially high priority/performance I/O processes in one system are delayed because another initiating controller is actively using the SCSI bus. The impact can be reduced greatly, if not eliminated, by giving the target controller a choice of paths on which to reconnect to continue an I/O process or giving the initiating controller a choice of paths on which to start an I/O process. (Reference the "don't call me, I'll call you" proposals for initiator and target interaction which make slaves of target controllers to attempt to better control system performance.)

In the SCSI-2 environment, each initiator defines an implicit path to a logical unit by successfully selecting a target controller and completing an I/O process with the logical unit. The target controller is restricted to responding to that I/O process using the target/initiator SCSI ID combination where the connect was made. Since the target controller cannot positively identify the initiating controller behind each initiator, each path is treated as if the initiating controller ID were different. This is an example of a path group containing only one path.

From these examples, several functions required of a multiple pathing definition become clear.

1) A host/path identifying mechanism needs to be provided. Use of the mechanism results in an explicitly defined path which specifies, by the way it is established, that a physical link exists or existed between a initiating controller and a logical unit or target routine.

2) Explicitly named paths must be capable of being formed into groups, called path groups. A path group identifies the set of available logically equivalent paths for communication between an initiating controller and a logical unit or target routine. (A cooperating set of initiators is identified as controlled by a superior element called the initiating controller.)

3) From 1) and 2), a mechanism to cause path groups to be established and disbanded, paths to be added to and removed from a path group, and reports of path group status must be provided. These are singular, non-supervisor functions.

4) A mechanism to temporarily suspend dynamic pathing operation must be provided. This is a supervisor function which may be controlled in the INFORMATION PACKET IDENTIFY message. (See X3T9.2/89-130 R02)

5) A mechanism to assign and unassign logical units, target routines, or extents within logical units to path groups which may span multiple SCSI busses must be provided. These functions require supervisor command status. (See X3T9.2/89-130 R02)

6) A mechanism to permit removing inoperative paths from a path group, or of transferring a logical device to another initiating controller must be provided. This function requires supervisor command status. (See X3T9.2/89-130 R02)

7) A mechanism to control the path on which status resulting from a contingent allegiance is presented, separate from the path(s) used for other segments of processing an I/O process. The path on which status is transferred which results from a contingent allegiance or extended contingent allegiance may be restricted to the path on which the connect occurred. If reporting status resulting from a contingent allegiance is restricted, the term single path status mode is used to define this condition. The alternative mode of operation is dynamic path status mode. All SCSI-2 I/O processes operate in single path status mode.

8) A modification to the definition of contingent allegiance and extended contingent allegiance is required. (See X3T9.2/89-130 R02)

9) Pointer management requires modification for dynamic path operation. (See X3T9.2/89-130 R02)

With properly defined functions, the SCSI-2 restriction of operations to a single path and access enabling functions restricted to a single SCSI bus are removed.

The relationship of several terms are shown in the figure below.

PHYSICAL ELEMENT	LOGICAL ELEMENT	IDENTIFIER
	Initiating Controller	ICID
Port	Initiator	SCSI Address
SCSI ID		
SCSI Bus		
SCSI ID		
Port	Target	SCSI Address Port Number
	Target Controller	
	Logical Unit or Target Routine	LUN TRN

END OF DESCRIPTIVE TEXT

THE LOGICAL OPERATION MODEL TEXT FOLLOWS

3.1. Glossary (Addendum)

ICID. See initiating controller ID.

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

initiating controller ID. An initiating controller ID is an identifier for an initiating controller which is communicated to target controllers over an SCSI bus.

initiator. An SCSI device, operating in initiator mode, through which an initiating controller starts I/O processes. An initiator usually attaches to an initiating controller. An initiator is a port to one SCSI bus.

port. A port is the name for the portion of an SCSI device where it attaches to one SCSI bus. An SCSI device may have more than one port. Each port may attach to a different SCSI bus. Each port has an SCSI ID and an SCSI address unique to the SCSI bus to which it attaches. Ports function as initiators and/or targets.

port number. When a port is in target mode, the port has a unique number within the target controller.

target routine. An addressable function within a target controller which executes I/O processes. A target routine is similar to a logical unit in that it has a name, a target routine number or TRN, and a command set to execute.

target routine number. The name of a target routine used during an I/O process to select a target controller function to execute an I/O process.

6.1. Logical Operation Model

The logical operation model describes device-independent activity using one or more SCSI busses. The next section provides a glossary related to the logical operations model. The glossary is followed by the logical operation model description.

6.1.1 Glossary for Logical Operation Model

assigned. An attribute of a logical unit or target routine where it responds to I/O processes only for certain path groups.

disband. A function which breaks up a set of paths established as a path group.

dynamic path status mode. A condition in initiators and targets where any status is transferred on any path in the path group where a connect is made. The status may or may not lead to a contingent allegiance.

establish. A function in target controllers where a non-empty set of paths is treated as equivalent, or grouped, for most I/O process activity.

grouped. The state of a path when it is included in an established path group. The path group must contain at least one path.

logical path. A logical path is the set of all paths which have the same ICID and LUN or TRN in the same target controller which indicates the routes I/O processes may take between an initiating controller and a logical unit or target routine.

password. An identifier used to permit otherwise unauthorized access to a logical unit or target routine.

path. A path is a named physical link between an initiating controller and a logical unit or target routine. At least one connect has been made from the initiating controller to the logical unit or target routine. The name consists of a ICID, an SCSI ID for the initiator, an SCSI ID for the target controller, the port of the target controller where the connect occurs, and the LUN or TRN of the selected logical unit or target routine.

path group. A path group is a logical path formed as a cooperating unit.

single path mode. A condition in a path group for one I/O process where single path status mode is in effect and all connections are made on the path where the connect occurred. The condition is reset when the I/O process completes.

single path status mode. A condition in a path group where status leading to a contingent allegiance is transferred on the path where the connect occurred. An implicit path or an ungrouped path is in single path status mode since it cannot cooperate with any other path in either state.

singular. An attribute of an SCSI command which prohibits its execution if it follows a command in an I/O process which has the Link bit set to 1 or which has the Link bit in its CDB is set to 1.


supervisor command. A command which may not be executed by a target controller unless specifically authorized.

unassigned. An attribute of a logical unit or target routine where it is permitted to respond to I/O processes on any path.

ungrouped. The state of a path when it is not included in an established path group.

Attributes of a Minimum Logical System

Initiating Controller ICID	Target Controller Target Controller Port Number
SCSI Address SCSI ID (parallel intf.)	SCSI Address SCSI ID (parallel intf.) Logical Unit LUN
Initiator - SCSI Device in Initiator Mode	Target - SCSI Device in Target Mode



```

Path = ICID ||
      Initiator SCSI Address ||
      Target SCSI Address ||
      Target Controller Port Number ||
      LUN
Logical path = ICID || LUN
Path implicitly defined
Path in the Ungrouped State
Path Group is Not Established
Path Status is Implicitly Named Path

Single path status mode
LUN is Unassigned
No password exists
    
```

Figure 6.c. MINIMUM LOGICAL SYSTEM ATTRIBUTES

6.1.2. Logical Operation Model Description

A logical system consists seventeen (17) items as follows:

- 1) a minimum of two SCSI ports and one SCSI bus connecting them;
 - 2) a minimum of one SCSI port must be capable of operating in initiator mode (called an initiator);
 - 3) a minimum of one SCSI port must be capable of operating in target mode (called a target);
 - 4) the initiator and target in 2) and 3) above, attach to the same SCSI bus and are active in their respective modes during a connection between them (i.e., not the same port);
 - 5) the logical element attaching an SCSI device which principally starts I/O processes is called an initiating controller.
 - 6) the logical element attaching an SCSI device which principally receives and executes I/O processes is called a target controller;
- NOTE: The names given to the logical elements attaching an SCSI device do not preclude any SCSI device from using all functions of SCSI. The word "principally" in items 5) and 6) imply this. Thus, a copy manager, acting principally as a target controller, may act as an initiating controller and use all defined functions of the commands and the logical system to perform a copy operation. This is the peer capability of SCSI as opposed to a master-slave relationship on some other I/O busses.
- 7) each port has an SCSI address unique to the SCSI bus on which it is attached; the SCSI address translates to the physical SCSI ID on some bus implementations.
 - 8) each port, when acting as a target, has a port number assigned by the target controller. The port number is unique within a target controller;
 - 9) each initiating controller is assigned an initiating controller ID. An initiating controller ID (ICID) must be unique in a logical system to prevent unpredictable results;
 - 10) each target controller has one or more logical units each identified by a unique Logical Unit Number (LUN).
 - 11) each target has zero or more target routines each identified by a target routine number (TRN).
 - 12) the extent of a logical system, from the viewpoint of a target controller, is the set of all initiating controllers having at least one port attached via a SCSI bus to at least one port of the target controller and from which a connect has been made.

From the point of view of the initiating controller, the extent of a logical system is the set of all logical unit and target routines to which a connect has been made. The set of paths available to any one logical unit or target routine is determined from the results of the INQUIRY command response data and the REPORT PATH STATUS command response data.

NOTE: A balanced logical system may consist of two SCSI busses with two ports for the initiating controller and two ports for the target controller. One port of each logical element attaches to one SCSI bus. This configuration provides a redundant path from the initiating controller to each logical unit in the target controller.

- 13) An identifier consisting of a ICID, an initiator SCSI address, a target controller SCSI address, a target controller port number, and a LUN or TRN, defines a path when the relationship is established

as the result of a connect started by an initiating controller to a target controller. The LUN or TRN must be valid for the target controller. The logical unit need not be ready or installed (e.g., unpowered but cabled or uncabled). No path exists between a LUN or TRN and an initiating controller unless the LUN or TRN is explicitly the object of a connect started by that initiating controller to the LUN or TRN.

NOTE: A connect using Asynchronous Event Notification does not define a path from the initiating controller to the logical unit or target routine. The logical unit or target routine must be selected as the receiver of an I/O process by the initiating controller. AEN uses a physical path between the target controller, with one port in initiator mode, and a initiating controller, with one port in target mode. From the target controller's point of view, an implicit path is defined as a result of the successful connect from the target controller to the initiating controller. Such a path may be developed in accordance with the Logical Operations Model as needed by the target controller based on the implemented function of the initiating controller when operating in target mode.

An implicitly named path exists when the initiating controller has not transferred its ICID to the target controller, but the initiating controller has made at least one connect to the LUN or TRN.

An explicitly named path exists when the initiating controller successfully completes an I/O process to transfer the ICID of the initiating controller to the LUN or TRN.

NOTE: All actions and functions below which refer to implicitly named paths have equivalent functions in SCSI-2. Any function in the logical operations model referring to an explicitly named path does not exist in SCSI-2.

14) An identifier, consisting of an ICID and either a LUN or TRN in the same target controller, represents one logical path. A logical path consists of a set of one or more paths.

An initiating controller connects with and transfers its ICID to each logical unit and target routine. The initiating controller is not required to transfer its ICID on all available physical paths between the initiating controller and the LUN or TRN, but only on those paths it intends to use for dynamic path operations or for which it intends to perform assignment operations other than for the one path.

NOTE: In the balanced logical system above, each initiating controller must connect with each LUN on two paths and transfer its ICID on each path to establish the set of explicitly named paths.

Explicitly naming a set of paths to a LUN or TRN from the same initiating controller does not establish a path group for conducting dynamic path operations. An explicitly named path is initially in the ungrouped state relative to other paths in the logical path.

An implicitly named path stays in an ungrouped state since the target cannot identify the initiating controller. Each initiator connecting with each target controller port must be considered as attached to a unique initiating controller until the transfer of an ICID from the initiating controller occurs.

15) A set of ungrouped explicitly named paths in a logical path is established as a path group through a command from the initiating controller using any one of the paths in the logical path. This set of paths or logical path, when established as a group, is called a path group. Establishing a path group is a singular operation.

A path may be added to an established group at a later time by a command from the initiating controller on that path.

A path may be removed from a path group by a command from the initiating controller on that path.

The inverse of establishing a path group is to disband a path group. A path group may be explicitly disbanded by a command from the initiating system along any one path in the path group. A path

group is implicitly disbanded when the last path is removed from an established path group using a remove path function rather than a disband command. The logical path remains as it did before the path group was established. Disbanding a path group is a singular operation.

Once a path group is established with two or more paths, the pointers for an active I/O process must be in a position to be shared between the initiators in the initiating controller servicing the paths in the path group.

NOTE: If no path group is established which contains two or more paths, SCSI pointer management in each initiator remains the same as in SCSI-2, since the target controller is restricted to operations only on the path where the connect occurred.

Once a path group is established, the extended functions of assignment and dynamic pathing operations may be used. The initiating controller determines whether either or both functions are used. The target controller by accepting the enabling commands has indicated that it is capable of participating in these functions.

Including a path in a path group does not restrict access to the logical unit or target routine from any other path. Establishing a path group is a naming process. Access is limited by the assignment function described in item 16 below.

Once a path group is established, it may be important to control the path where status leading to a contingent allegiance is reported. This is especially true if the main recovery mechanism for all I/O processes is located on one path and other paths are treated like data highways by the initiating controller. Therefore, a function of establishing path groups is to identify where status leading to a contingent allegiance is reported. Any status which does not result in contingent allegiance may be sent over any path in the path group. The condition may be altered by disbanding the group and establishing the group with the alternate choice. The initiating controller is given two choices:

- 1) single path status mode is an initiating controller established condition in both the initiating controller and the target controller where status resulting in contingent allegiance is sent only over the path on which the connect was made;

- 2) dynamic path status mode is an initiating controller established condition in both the initiating controller and the target controller where status is sent over the next available path in the established path group.

When an implicitly named path, an ungrouped path, or a path group containing only one path is used to make a connect, dynamic path operation is not permitted. As a result, single path status mode is in effect.

When operating in dynamic path status mode, an initiating controller may temporarily switch to single path mode for an I/O process without affecting the established path group. The function is in effect for all linked commands in an I/O process. In addition to single path status mode, all activity related to the I/O process must occur on the path where the connect was made. The function is called suspend dynamic path operation.

The state of the path may be any one of the following:

- 1) Implicitly named path. No explicit ICID has been received from an initiating controller to any LUN or TRN on this target controller from the initiator SCSI ID/target port combination. An I/O process received on a path in this state is required to perform all operations on this path.

- 2) Path to Other LUNs. An explicitly named path to at least one LUN or TRN on this initiator SCSI ID/target port, other than the selected logical unit or target routine, exists. This is functionally

equivalent to 1) but it imparts additional information to the initiating controller. An I/O process received on a path in this state is required to perform all operations on this path.

3) Ungrouped. An explicitly named path to the selected LUN or TRN exists but it is not currently part of an established path group. An I/O process received on a path in this state is required to perform all operations on this path.

4) Grouped. An explicitly named path to the selected LUN or TRN exists and is currently established in the grouped state. The path group can consist of one or more paths. An I/O process received on this path may respond on any path in this path group unless single path status is in effect or dynamic path reconnection has been temporarily suspended for an I/O process.

Both the initiating controller and the target controller must keep track of the paths, path groups, the state of each path in the group, and the status transfer mode for each LUN or TRN. Both must communicate on the appropriate paths and the target controller must be prepared to report the state of any path to the initiating controller.

All path groups are established and managed by the functions defined above save one exception condition which is described in item 17. Path identification and grouping is not a supervisor mode operation since it does not restrict access to the LUNs in a target controller.

16) Any logical unit, target routine or extent attached to an SCSI bus is initially available to receive I/O processes from any initiating controller attached to that SCSI bus. This use privilege is extended whether the path is explicitly named, implicitly named, and whether for explicitly named paths, the path is grouped or ungrouped. This state of access is logically equivalent to the SCSI-2 bus with no reservations outstanding.

Such unrestricted access may not be appropriate for environments with extensive multi-user access and/or data bases with sensitive information. Therefore, it is appropriate to control access to a LUN logical unit, target routine or an extent within a logical unit at a higher level than the RESERVE/RELEASE functions defined in SCSI-2.

The ability to control use privileges requires path group control when multiple SCSI busses are involved. Two or more explicitly defined path groups may share use privileges to the exclusion of other path groups. A single path group may hold exclusive use privileges.

Use privileges are controlled with two functions called assign and unassign. Because of the implications to system reliability and integrity, these functions are defined as supervisor commands. Their purpose is to act as the logical equivalent of switches or manual cable changes to restrict access to logical units, target routines, or extents within logical units.

Assignment to one path group means that no initiating controller on any path not in the path group can gain access to the logical unit using the functions defined to this point. Certain commands, such as REQUEST SENSE and INQUIRY may be responded to regardless of the source of the command. A logical unit may be assigned to multiple established path groups. An implicitly named path or an ungrouped path can gain assignment for itself, but it cannot add assignment for any other paths or path groups.

Any path holding assignment through an established path group may add assignment of other established path groups. The two functions of assignment are:

- 1) assign this LUN to the path group on which the command was received, or
- 2) add assignment of this LUN for another established path group.

The inverse of assign is unassign. Use privileges may be unassigned from any path group to which assignment currently exists from any path for which assignment currently exists. Assignment may be transferred from one initiating controller to another without passing through a state where no assignment exists.

Initiating controllers, through their use of the path naming functions and grouping functions, may be assigned use privileges or they may be unassigned. The mechanism by which an initiating controller obtains the path group name required for adding or removing assignment of additional path groups is not established by or a concern of the logical system.

17) The provision for assignment permits multiple initiating controllers to control use privileges. The last function in the logical operations model concerns breaking an assignment if some error or failure in an initiating controller which currently has assignment. The break may be temporary or permanent, but it must be controlled, as are other functions which can lead to reliability, availability, and data integrity problems. This function, above all the rest, requires supervisor mode. This is not a singular function.

Assignment permits use privileges through assigned path groups. Controlled access is a function which permits access outside the bounds of defined assignment functions. The mechanism to prevent deliberate or accidental loss of assignment protection is the control access function, enabled by a password, and checked by the affected target controller.

A password is established by an initiating controller having assignment. The password is not reported by a target controller on any path. The target controller checks its established password, if any, against password supplied by the control access command function from an initiating controller not having assignment.

If the target controller has a password established and it matches the password with the command, the control access command and any commands linked to it are executed, if possible. The mechanism by which the unassigned initiating controller acquires the correct password is not established by or a concern of the logical system.

The control access command has three functions:

1) establish a password in a target controller. The command must be received on a path currently holding assignment.

2) general unassign. The control access command is received from a path that does have assignment. If no password has been set or the password supplied matches the password in the target, the target controller removes assignment for all paths in any path group having assignment when the command was processed. The result is that the logical unit, target routine, or extent has no assignment protection.

3) request temporary unassignment. The control access command is received from a path that does not have assignment. If the password matches the password in the target controller, commands linked to the control access command are executed to the extent possible. A status which would lead to contingent allegiance on the unassigned path is not permitted, since that would grant the unassigned path permission to continue operations with a REQUEST SENSE command to retrieve the sense data. The contingent allegiance is made with an assigned path whether functional or not.

When a request for temporary assignment is granted, the issuing initiating controller may link to an assign command which will grant assignment to the once unassigned initiating controller. The initiating controller can then use normal I/O processes.

An I/O process containing a valid control access command requesting temporary assignment, a control access command performing a general unassign, and an assign command for the new initiating controller breaks all old assignments and transfers assignment of the LUN to the new initiating controller. This operation permits continued operation without loss of the function provided by the logical unit. Some I/O processes may be aborted.

7.x.z. SCSI-3 Path Control Commands and SCSI-2 Command Changes

The path control commands required to carry out the logical operation model are:

- 1) SET ICID. Explicitly name the initiating controller using this path on which a connect occurred, establish and disband groups.
- 2) REPORT PATH STATUS. Report the status of the path relative to path groups. This is a singular command
- 3) ASSIGN. Define the path groups on which a logical unit may operate. This is a supervisor command.
- 4) UNASSIGN. Remove a path group from the set of authorized path groups granted access to a logical unit. This is a supervisor command.
- 5) The SCSI-2 functions in the RESERVE and RELEASE commands must be included in the ASSIGN and UNASSIGN commands.
- 6) CONTROL ACCESS. Transfer password on an assigned path to the target controller for a path group; permit general unassign from an assigned path group; and, allow an unassigned path, having the correct password, access to a logical unit for one I/O process. This is a supervisor command.
- 7) SUSPEND DYNAMIC PATH OPERATIONS. Suspend for the length of a single I/O process the use of dynamic path operation. All reconnects occur on the path where the connect occurred. This is a message level control as opposed to the command level control in item 8) which causes a general suspension of dynamic pathing operations within the path group.
- 8) SUSPEND DYNAMIC PATH OPERATIONS. Suspend the use of dynamic path reconnection even though the path group is correctly formed. All reconnects occur on the path where the connect occurred.
- 9) Considerations for when and how path groups are terminated and when and how assignment is terminated must be clearly stated.

The exact formats and processing rules are TBD. However, they obey the rules of the logical operations model.

END OF DOCUMENT X3T9.2/89-133 R02