



adaptec, inc.

X3T9.2/87-53

Obs - see
revision 1

Memo to: John Lohmeyer, ANSI X3T9.2 Chairman
NCR Corp.
3718 N. Rock Road
Wichita, KS 67226

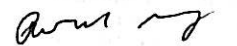
Memo from: Robert N. Snively
Adaptec
580 Cottonwood
Milpitas, CA 95035

Date: March 10, 1987

Subject: Proposed Appendix Describing Power On Protocols

Dear Mr. Lohmeyer:

The following text is a draft proposal for an appendix to define the appropriate system initialization sequence. The intent of this appendix is to provide a guideline for implementers of SCSI driver software that will always assure a successful system initialization process. The initialization should work as well for SCSI devices as for SCSI-2 devices, except that slightly more information may be available from standard SCSI-2 devices.


Robert N. Snively

155

APPENDIX A RECOMMENDED INITIALIZATION PROCEDURE

This appendix defines the normal mechanisms for obtaining the information required for system initialization from SCSI devices. The procedure should operate correctly for all SCSI-2 devices as well as all SCSI devices meeting conformance level 2.

A.1 Information Required for System Initialization

The following list of information is assumed necessary and sufficient for normal system initialization. The procedure will document the steps required to obtain this information and to achieve the desired initial states in the attached devices.

The following information is required for the system to complete its initialization:

- 1) A list of the installed and powered on SCSI device ids for each SCSI port. SCSI devices that are not powered on are treated as not installed, assuming that the terminators are powered from a source other than the un-powered SCSI devices.
- 2) A list of the installed SCSI LU's for each SCSI device. Unpowered or failing LU's may not be completely identifiable.
- 3) The device type for each available LU.
- 4) The manufacturer and model for each available LU.
- 5) The state of readiness of each available LU.
- 6) The critical device type information for each available LU. This information will vary depending on the device type.
- 7) The capability of attached host targets and non-host initiators should be identified by interested initiators and targets.

The following states should be established for each attached LU that has power available and is not failing.

- 1) The LU should be started and established in its initial state if media is installed.
- 2) The LU should be prepared for media installation if media is not installed.
- 3) All error conditions associated with the starting process should be cleared.
- 4) All UNIT ATTENTION conditions should be cleared.
- 5) All data transfer parameters should be established.
- 6) All system tuning parameters should be established where known. Note that these may be modified dynamically to improve the performance characteristics of the system.

n.2 Procedure for Initializing Devices

The system will execute the following steps to perform initialization. Some of the steps will be detailed in subsequent paragraphs. Note that the text represents a primitive pseudo-code that can be converted to the appropriate software object by those who implement device drivers. Since a multiple host environment is assumed, implementation of the Soft Reset option is required. } (2)

n.2.1 Procedure Executed by Host Initiators

Host Initiator Activities:

Power On: It is assumed that each SCSI device as it is powered on performs appropriate internal reset operations and internal test operations. Once powered on, initiators that have target capability must be prepared to respond to a selection within 250 milliseconds.

Soft Reset: At power on time, it is likely that an SCSI device has forced errors on the ongoing activities on the SCSI bus. A soft reset will be generated to notify attached active devices that any activities that may have been occurring should be restarted.

Find Device: Each address other than the testing initiator's address will be tested to determine if any SCSI device is attached. If attached, the INQUIRY command will be executed on LUN 0. The information obtained will indicate the device type, manufacturer and model of the attached

LUN 0. In addition, the standard command set supported by the device will be indicated by the ANSI-Approved Version field.

Find LUN:

Each possible LUN address on the devices found to exist will be tested for existence using an INQUIRY command. Those found to be device type 7Fh (Logical Unit Not Present) will not be included in the list of LU's available. Those described by any other device type indicator will be added to the tables and the device type, manufacturer and model of the attached LUN recorded.

Verify State:

The verify state test (Verify State Test:) will be made to clear any outstanding errors, capture and clear UNIT ATTENTION conditions, and determine the state of readiness of the LUN. The device will be identified as Ready, Not Ready, or Failing at the end of this test.

Device Init:

The device will undergo a device dependent initialization process. This process is described for Direct Access Devices in section n.3.1. The process for Sequential Access Devices will be described in section n.3.2. The process for Processor Devices will be described in section n.3.3. Other devices will not be considered since they tend to be similar to one or another of the already described initialization procedures. The initialization process takes into account the state of the device as identified during the Verify State Test.

Device Online:

The successful completion of the Device Init process will allow the device table entry to be fully enabled and the device may join the system with all key parameters identified and initialized and with the device state known and presented to the system operator.

n.2.2 Procedure Executed by Peripheral Device Initiators

A peripheral device initiator typically performs operations only under the direction of a host initiator. Such operations may include the Read and Write commands associated with management of a Copy command. Other possible applications include the SEND command associated with Asynchronous Event Notification. As such, they need not completely perform the last two steps defined above. Since all commands are simple logical commands whose logical addresses and file structures are managed by a host processor, the Peripheral Device Initiators normally need not recover information about the media density, the sparing algorithms, or other detailed information that may be required

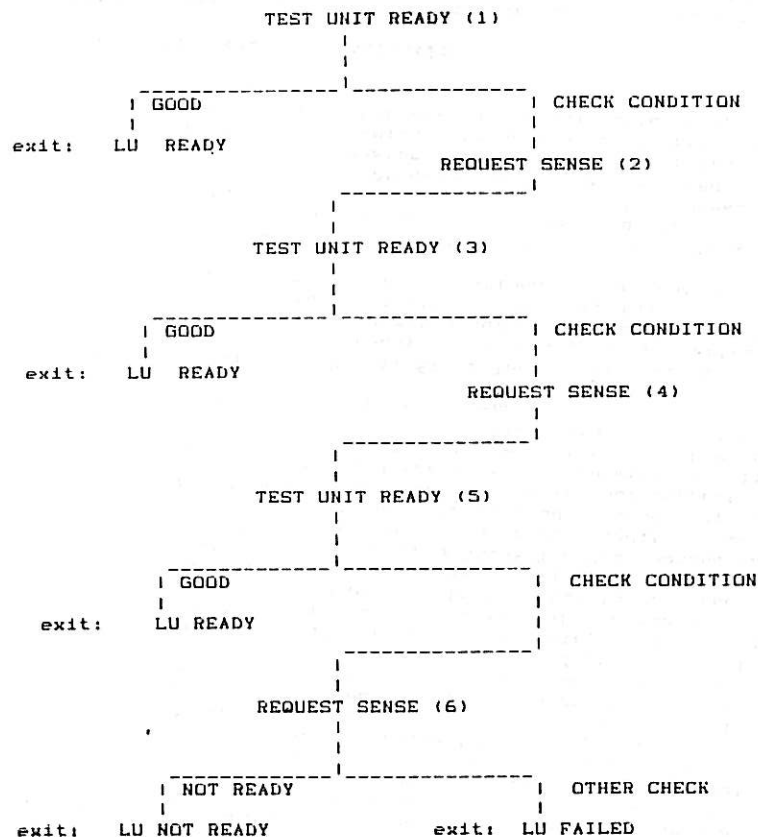
only by a host processor.

n.2.3 Detail of Shared Process Steps and Tests

n.2.3.1 Verify State Test:

The Verify State Test uses the following steps to determine whether a device has any outstanding errors, to clear any outstanding UNIT ATTENTION conditions, and to determine the state of the device. The Verify State Test will be executed against all devices that have returned INQUIRY data with other than a 7Fh device type.

Verify State Test:



TEST UNIT READY (1):

This TEST UNIT READY command is used to determine if any outstanding check condition or UNIT ATTENTION condition exists. If not, the device is indicated to be ready.

REQUEST SENSE (2):

This REQUEST SENSE command is used to clear the outstanding check condition. A normal SCSI-2 LU will probably present UNIT ATTENTION in this sense information.

TEST UNIT READY (3):

This TEST UNIT READY command is used to see if the UNIT ATTENTION or other error was successfully cleared. In some special cases, another error may have been nested with the UNIT ATTENTION and this TEST UNIT READY command may also detect CHECK CONDITION status.

REQUEST SENSE (4):

This REQUEST SENSE command is used to determine the error that was associated with the TEST UNIT READY (3). This may be a NOT READY indication or another unexpected error. The error should be cleared by the REQUEST SENSE.

TEST UNIT READY (5):

This TEST UNIT READY command is used to see if the check conditions have finally been cleared. If they have, the LU is identified as Ready.

REQUEST SENSE (6):

This REQUEST SENSE command is used to determine the final check condition that will be associated with the LU under test. If the check condition indicates NOT READY, the LU is identified as Not Ready. If the check condition indicates some other failure, the LU is identified as Failing and the error condition is logged in the appropriate area.

n.3.1 Direct Access Device Initialization Procedure

The device dependent initialization process for a Direct Access Device may be divided into three independent activities. The first activity enables the minimum logical function required for execution of the boot device READ commands. The second activity is executed against all devices, including the boot device, and is executed to establish all required initial parameters and operating conditions. The third activity is performed against devices that have never been formatted or initialized. This activity is normally performed by an initialization utility.

n.3.1.1 Boot Device Initialization

It is assumed that the boot program and boot device have been prepared in such a manner that proper block lengths, data file

contents, and logical addresses have been implemented by both the boot device and the boot program. The boot program must then prepare the boot device for operation in the following manner.

Verify Ready: The state of the device as determined by the Verify State operation is examined. If the step indicates that the required drive is Failed, the Boot Device Initialization is not performed and appropriate error indications are presented.

Start Device: A START/STOP command is executed with the start bit set to cause the drive to spin up if it is not already spinning. The Immediate bit should be set to 0 in order to guarantee that the returned status reflects the completion of the device start operation. A disconnect operation is likely, since the power on process may take a considerable period of time. If system controlled power sequencing of the drives is required, it is done by managing the timing relationship of the START/STOP commands to different LU's. If GOOD status is returned, the next step is started. If CHECK CONDITION status is presented, a REQUEST SENSE command is executed to determine what error condition was detected. If an ILLEGAL REQUEST sense key is found, the START/STOP command was not supported by the device and the next step is started. If any other error is detected, the Boot Device Initialization is terminated and appropriate error indications are presented.

Verify Ready / Spinning: A Verify State Test is performed. If the device is Ready, the next step is started. If the device is Not Ready or Failing, the Boot Device Initialization is terminated and appropriate error indications are presented.

Boot: The boot READ operations can now be started against the drive. It is assumed that the information read in to the host processor includes the programs that are required to continue the system initialization and bring-up process, including the necessary programs to perform the other system initialization steps.

n.3.1.2 General System Initialization

After booting has been performed, the host is assumed to have the necessary device drivers to perform other initialization functions. A general direct access device initialization procedure is defined below. The initialization procedure must be executed for each attached LU that has been identified as a

Direct Access Device during execution of the previous INQUIRY command. Execution of these steps may be overlapped from one LU to another. If the host is supporting only a limited range of devices, many of these steps and operations may be skipped or simplified.

Verify Ready: The state of the device as determined by the Verify State operation is examined. If the step indicates that the required drive is Failed, the General System Initialization is not performed and appropriate error indications are presented.

Start Device: A START/STOP command is executed with the start bit set to cause the drive to spin up if it is not already spinning. If GOOD status is returned, the next step is started. If CHECK CONDITION status is presented, a REQUEST SENSE command is executed to determine what error condition was detected. If an ILLEGAL REQUEST sense key is found, the START/STOP command was not supported by the device and the next step is started. If any other error is detected, the General System Initialization is terminated and appropriate error indications are presented.

Verify Ready / Spinning: A Verify State Test is performed. If the device is Ready, the next step is started. If the device is Not Ready or Failing, the General System Initialization is terminated and appropriate error indications are presented.

Determine Parameters: If the ANSI-Approved Version field of the previously executed INQUIRY command was 0 or 1, the Mode Sense information is vendor unique and this function should be skipped unless required by the vendor unique initialization protocols.

If the ANSI-Approved Version field is 2, Mode Sense information is partially standardized by the SCSI-2 document. In this case, a MODE SENSE command is executed with the page control field set to request Current Values and the page code set to request all pages. A record is made of the Current Values.

A second MODE SENSE command is executed with the page control field set to request Changeable Values for all pages. A record is made of the Changeable Values.

Any errors that occur during the execution of the two MODE SENSE commands are recorded and the initialization for the failing LUN is terminated.

Set Parameters: If the ANSI-Approved Version field is 0 or 1, the initialization operation is vendor unique and must be executed according to the rules for the particular installed device. The system must have some other source of information concerning these requirements.

If the ANSI-Approved Version field is 2, the Mode Select format is largely defined by the SCSI-2 specification. The actual set up requirements for the parameters are characteristic of the particular system and must be known to the system. The Current Values and the Changeable Values obtained by the previous step are examined to see if the set up requirements are satisfied and if they can be modified. If all values are correct, the remainder of this function may be skipped. If modifications must be made to the Changeable Values, a MODE SELECT is performed to modify the appropriate pages. This may include modifying error recovery parameters or performance tuning parameters. Most geometry parameters will not be modifiable during General System Initialization.

Determine Capacity: The capacity and block size of the LU are determined by executing a READ CAPACITY command. The information is stored for access by the system device drivers.

The device is now fully initialized and all required device information has been made available to the system. When all available non-failing devices have been initialized, the system is entirely operational and system initialization is considered complete.

159 n.3.1.3 Direct Access Device Media Initialization Procedure

The following initialization procedure is not part of normal power up system initialization. It is assumed to be performed after completion of the General System Initialization process but uses only the INQUIRY information obtained during that process. It is performed to initialize device media and is normally performed only by a system utility.

Determine Format Requirement: The requirement to perform a format operation is normally generated by an operator who has just installed a new device known to require formatting. It may also be

generated by recognition that a device has information that is no longer valid and that must be totally erased. It may also be generated by changes in system requirements, including different block sizes. Finally, reformatting may be required to restructure the defect management.

The General System Initialization has probably identified the device as failing because of the inability of the device to recover Read Capacity parameters. The machine is assumed to have been started during the General System Initialization process. The Verify State Test is now executed again. The device should be Ready according to that test. If the LU is Not Ready or Failing, the Media Initialization process is terminated and appropriate errors are posted.

Determine Format Parameters: If the ANSI-Approved Version field is 0 or 1, the media initialization operation is vendor unique and must be executed according to the rules for the particular installed device. The system must have some other source of information concerning these requirements.

If the ANSI-Approved Version field is 2, a MODE SENSE command is executed with the page control code set to Current Values and the page code set to All Pages. The information obtained should provide all the information available to the LU about its format requirements. The MODE SENSE command is executed again with the page control code set to Changeable Values and the page code set to All Pages. The information indicates what values must be provided by the system to complete the format parameters. If either of these MODE SENSE operations does not complete normally, the media initialization operation is terminated and appropriate error indications recorded.

Set Format Parameters: If the ANSI-Approved Version field is 0 or 1, the format requirements are vendor unique and the appropriate commands must be known to the initialization utility. Those format preparation commands are executed at this time.

If the ANSI-Approved Version field is 2, the LU is prepared for media formatting by execution of a MODE SELECT command. The necessary formatting parameters are selected to meet the system requirements and filled into the changeable value locations. The MODE

SELECT command is then executed. If the command fails, appropriate error indications are recorded and the Media Initialization process is terminated. If the command succeeds, the next step is performed.

Perform Format Operation: The appropriate format parameters are established and the FORMAT command is executed. The FORMAT parameters depend on the system requirements and the device capabilities. These parameters should be made easily variable in the operating system architecture so that modifications can be performed when system or device requirements change. An ILLEGAL REQUEST error may be generated if improper format parameters are selected. If the command fails, appropriate error indications are recorded and the Media Initialization process is terminated. If the command succeeds, the device is fully operational and the next step is performed.

Set Parameters: If the ANSI-Approved Version field is 0 or 1, the initialization operation is vendor unique and must be executed according to the rules for the particular installed device. The system must have some other source of information concerning these requirements.

169
If the ANSI-Approved Version field is 2, the Mode Select format is largely defined by the SCSI-2 specification. The actual set up requirements for the parameters are characteristic of the particular system and must be known to the system. The Current Values and the parameters established by the MODE SELECT / FORMAT sequence are examined to see if the set up requirements are satisfied and if they should be modified. If all values are correct, the remainder of this function may be skipped. If modifications must be made to the Changeable Values, a MODE SELECT is performed to modify the appropriate pages. This may include modifying error recovery parameters or performance tuning parameters. Most geometry parameters were established by the storing of parameters during the MODE SELECT / FORMAT sequence and will not be modifiable.

Determine Capacity: The capacity and block size of the LU are determined by executing a READ CAPACITY command. The information is stored for access by the system device drivers.

The device is now initialized and prepared to partake in system oriented activities. Other system initialization operations may also be required, including the establishment of system volume labels, tables of contents, and others.

n.3.2 Sequential Access Device Initialization Procedure

TO BE DEFINED

n.3.3 Processor Device Initialization Procedure

TO BE DEFINED