



To: INCITS Technical Committee T10
From: Fred Knight, Network Appliance
Email: knight@netapp.com
Date: 19 November 2008
Subject: SPC4

1) *Revision history*

Revision 0 (19 November 2008)

2) *Related documents*

spc4r16 – SCSI Primary Commands – 4
08-149r5 – Thin Provisioning for SBC

3) *Overview*

The XCOPY command has a short set of generic text talking about preparation of the copy source and destination devices (MODE SELECT, etc). This proposal has two possible methods (only 1 should be added) to be more specific about the need to pre-setup the devices to prevent RECOVERED type errors from aborting the XCOPY operation.

Existing text is shown in **BLACK**, new text is shown in **RED**, comments in **BLUE**.

Proposal:

SPC4r16 Changes

6.3 EXTENDED COPY command

6.3.1 EXTENDED COPY command introduction

The EXTENDED COPY command (see table 99) provides a means to copy data from one set of logical units to another set of logical units or to the same set of logical units. The entity within a SCSI device that receives and performs the EXTENDED COPY command is called the copy manager. The copy manager is responsible for copying data from the source devices to the destination devices. The copy source and destination devices are logical units that may reside in different SCSI devices or the same SCSI device. It is possible that the copy source device, copy destination device, and the copy manager are the same logical unit.

Table 100 — EXTENDED COPY command

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE (83h)							
1	Reserved							
9	Reserved							
10	(MSB)	PARAMETER LIST LENGTH						(LSB)
13	Reserved							
14	Reserved							
15	CONTROL							

Before the copy manager is instructed to move data, the application controlling the data movement shall independently take any necessary actions required to prepare the copy source and destination devices for the EXTENDED COPY command (e.g. loading tapes, sending media changer commands, MODE SELECT commands (e.g., to disable reporting of recovered errors by a disk copy target device set the PER bit to zero and to disable reporting of thin provisioning threshold events by a disk copy target device, set the TPNR bit to one in the ERROR RECOVERY MODE PAGE (see SBC)), reservation commands, and/or tape positioning commands). After all preparatory actions have been accomplished; the EXTENDED COPY command should be issued to the copy manager to start the data transfer.

The above text is one of two possible options to be more explicit about RECOVERED ERRORS and specifically Thin Provisioning. A second option is below. One or the other method (eg text vs. note); not both will be recommended based on CAP discussions.

The PARAMETER LIST LENGTH field specifies the length in bytes of the parameter data that shall be contained in the Data-Out Buffer. A parameter list length of zero specifies that copy manager shall not transfer any data or alter any internal state; this shall not be considered an error. If the parameter list length causes truncation of the parameter list in a target descriptor or segment descriptor, then no data shall be transferred and the EXTENDED COPY command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST, and the additional sense code set to PARAMETER LIST LENGTH ERROR.

<...>

6.3.3 Errors detected during processing of segment descriptors

Errors may occur after the copy manager has begun processing segment descriptors. These errors include invalid parameters in segment descriptors, invalid segment descriptors, unavailable targets referenced by target descriptors, inability of the copy manager to continue operating, and errors reported by source or destination copy target devices. If the copy manager receives CHECK CONDITION status from one of the copy target devices, it shall recover the sense data associated with the exception condition and clear any ACA condition associated with the CHECK CONDITION status.

If it is not possible to complete processing of a segment because the copy manager is unable to establish communications with a copy target device, because the copy target device does not respond to INQUIRY, or because the data returned in response to INQUIRY indicates an unsupported logical unit, then the EXTENDED COPY command shall be terminated with CHECK CONDITION status, with the sense key set to COPY ABORTED, and the additional sense code set to COPY TARGET DEVICE NOT REACHABLE.

If it is not possible to complete processing of a segment because the data returned in response to an INQUIRY command indicates a device type that does not match the type in the target descriptor, then the EXTENDED COPY command shall be terminated with CHECK CONDITION status, with the sense key set to COPY ABORTED, and the additional sense code set to INCORRECT COPY TARGET DEVICE TYPE.

If the copy manager has issued a command other than INQUIRY to a copy target device while processing an EXTENDED COPY command and the copy target device either fails to respond with status or responds with status other than BUSY, TASK SET FULL, ACA ACTIVE, or RESERVATION CONFLICT, then the condition shall be considered a copy target device command failure. In response to a copy target device command failure the EXTENDED COPY command shall be terminated with CHECK CONDITION status, with the sense key set to COPY ABORTED, and the additional sense code set to THIRD PARTY DEVICE FAILURE.

If a copy target device completes a command from the copy manager with a status BUSY, TASK SET FULL, ACA ACTIVE, or RESERVATION CONFLICT, **then** the copy manager shall either retry the command or terminate the EXTENDED COPY command as a copy target device command failure.

NOTES

- 23 The copy manager is assumed to employ a vendor specific retry policy that minimizes time consuming and/or fruitless repetition of retries.
- 24 RESERVATION CONFLICT status is listed only to give the copy manager leeway in multi-port cases. The copy manager may have multiple initiator ports that are capable of reaching a copy target device, and a persistent reservation may restrict access to a single I_T nexus. The copy manager may need to try access from multiple initiator ports to find the correct I_T nexus.
- 24.1 **If a copy source device or copy target device responds to an input or output operation with a CHECK CONDITION status with the sense key set to RECOVERED ERROR, or UNIT ATTENTION then the EXTENDED COPY command will be terminated. If terminating EXTENDED COPY commands due to recovered conditions is not desired, the copy manager should disable reporting of recovered errors (e.g., set the PER bit to zero and the TPNER bit to one in the ERROR RECOVERY MODE PAGE (see SBC)) by the copy source device and/or copy target device.**

If a copy target device responds to an input or output operation with a GOOD status but less data than expected is transferred, then the EXTENDED COPY command shall be terminated with CHECK CONDITION status, with the sense key set to COPY ABORTED, and the additional sense code set to COPY TARGET DEVICE DATA UNDERRUN. If an overrun is detected, then the EXTENDED COPY command shall be terminated with CHECK CONDITION status, with the sense key set to COPY ABORTED, and the additional sense code set to COPY TARGET DEVICE DATA OVERRUN.

<...>