

CONGRUENT SOFTWARE, INC.
514 48th AVENUE
SAN FRANCISCO, CALIFORNIA 94121

TO: John Lohmeyer
Chairman, X3T9.2

FROM: Peter Johansson

DATE: January 7, 1988

RE: Management of the "Grown" Defect List

This is a resubmission of the issues raised in X3T9.2/87-151, my memorandum of August 28, 1987. The original escaped inclusion on plenary or working group agenda and minutes, so I have reviewed the issues against the current SCSI-2 draft. I remain convinced that the FORMAT UNIT, READ DEFECT DATA and REASSIGN BLOCKS commands for direct-access devices need enhancement and clarification with respect to management of the "Grown" defect list and request the committee's consideration.

143

First, the READ DEFECT DATA command. Just what does it mean to return a list of logical block addresses for defect locations if they are remapped by the target? The third sentence of the paragraph referenced below obliquely addresses the problem, but because defect locations that have been remapped have no logical block addresses, I believe it is not strong enough. I suggest that the first paragraph on page 8-55, Section 8.1.9 of Revision 3 of the working draft be altered as follows:

The Defect List Format field specifies the format of the defect list data returned by the target. Legal values for this field are defined in the FORMAT UNIT command (see section 8.1.1.1, Defect List Format). The logical block and physical sector formats may not include defects in user inaccessible areas. If the logical block format is specified, only those defect locations that have not been remapped by the target shall be returned. If either the bytes from index or physical sector formats are specified, the complete list(s) of defects shall be returned; however, some of the defects may not lie in user accessible areas (depending on the disk geometry used by the most recent FORMAT UNIT command). If the value in this the Defect List Format field does not match the defect list format requested by the initiator in the CDB, the target shall transfer the defect data and return a CHECK CONDITION status with the sense key set to RECOVERED ERROR and an additional sense code of DEFECT LIST NOT AVAILABLE.

This ability to distinguish between defects that are recorded in a target's Primary or "Grown" lists and defects that are both recorded and remapped is useful in conjunction with REASSIGN BLOCKS as described below.

Secondly, the REASSIGN BLOCKS command description makes no mention of the fact that the "Grown" defect list is updated when the command is executed. This fact is buried back in the FORMAT UNIT command description. This command provides the only means of adding to the "Grown" defect list other than at the time the unit is formatted, and because of that should have the added capability to add defect(s) to the "Grown" list without remapping them. Three situations come to mind where this is valuable: 1) A volume in use for some time has exhausted its spare sectors and this fact is revealed by a CHECK CONDITION after a REASSIGN BLOCKS command. The initiator will need to reformat the volume with more spare capacity; rather than go through a tortuous logical block address calculation for the new "virtual" disk geometry and the "physical" logical block addresses (now rather confusedly described under the FORMAT UNIT command) it would be more direct for the initiator to issue a REASSIGN BLOCKS command for the affected blocks (specifying no remapping) followed by a FORMAT UNIT command that retains the "Grown" defect list. 2) The initiator does not remap any defects, performs its own defect management schemes but wishes to utilize the target's capability to record defect lists for future reference. 3) An initiator performs an incoming inspection test of the target with all defect management and remapping disabled and discovers "Grown" defects. These are to be added to the defect list when the target is reformatted with defect remapping enabled; once again, it would be more direct to use the REASSIGN BLOCKS command rather than a calculation of "physical" logical block addresses.

There are at least two ways to modify REASSIGN BLOCKS to add new functionality: a) add a bit in the CDB that specifies that the block(s) are to be recorded in the "Grown" list, only, and not reassigned or b) change the definition of the actions performed after the NO DEFECT SPARE LOCATION AVAILABLE error condition such that the defect location being processed is added to the "Grown" list. The first approach is cleaner and is reflected in the proposed wording for the REASSIGN BLOCKS command attached to this memorandum.

In informal discussions at the Colorado Springs meeting about the REASSIGN BLOCKS command, the objection was raised that all defects recorded in the "Grown" defect list at present are reassigned and that if REASSIGN BLOCKS were to be modified as described above the initiator would not be able to distinguish reassigned from unreassigned defects. The changes to READ DEFECT DATA (which I believe are needed on their own merits), along with the newly added address translation functions of SEND DIAGNOSTIC and RECEIVE DIAGNOSTIC RESULTS answer this objection.

Lastly, there is the management of the "Grown" defect list through the FORMAT UNIT command. Revision 3 of the SCSI-2 draft incorporates some of the changes recommended the August 28 memorandum and is less ambiguous with respect to the logical block format when the CmpLst bit is one. Additional clarification could be provided by revising the first paragraph on page 8-5, Section 8.1.1 as follows.

A Complete List (CmpLst) bit of one indicates that the data supplied by the initiator during the DATA OUT phase of the command (D list) includes the complete list of known defects. Any previous defect map or defect data, except the P list (if any), shall be erased by the target. The result is to purge any previous G list and to build a new G list. If there was a previous P list, it shall be retained. The target may add to this D list as it formats the medium, according to the format options selected in the defect list header (see Table 8-3), by performing a target certification process (creating a new C list). The combination of the D list and the C list (if certification is performed) creates the new G list. When using the block format with a CmpLst bit of one, the defect list refers to the new block length (and not to the previous block length, if it is different) and the defect list refers to physical block addresses (i.e. addresses that the logical blocks would occupy if there were no defects prior to any spare sector or track allocation and prior to any defect reassignment, including that of the P list, not the previous logical block addresses).

Thank you for your consideration of these interrelated problems.

8.1.11. REASSIGN BLOCKS Command

Peripheral Device Type: Direct-Access and Write-Once
Read Multiple
Operation Code Type: Optional

Table 8-36: REASSIGN BLOCKS COMMAND

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (07h)							
1	Logical Unit Number						NoMap	
2	Reserved							
3	Reserved							
4	Reserved							
5	Reserved							
6	Control Byte							

The REASSIGN BLOCKS command (Table 8-36) requests the target to reassign the defective logical blocks to an area on the logical unit reserved for this purpose and to record the defective logical blocks to the "Grown" defect list. More than one physical or logical block may be relocated by each defect descriptor sent by the initiator. This command does not alter the contents or location of the Primary defect list (see section 8.1.1.1, FORMAT UNIT command).

If the suppress reassignment (NoMap) bit is one, the target shall enter the defect(s) in the "Grown" defect list but not attempt to reassign references to the location(s) specified in the defect list.

The initiator transfers...(the remainder of the REASSIGN BLOCKS command description is unchanged).