

Date: January 5, 1994 X3T10/94-026 Rev 0
 To: X3T10 Committee (SCSI)
 From: Doug Hagerman (Digital)
 Subject: Comments on SDA States (93-191)

This proposal is an expansion of the SDA States described by George Penokie. Given the complex nature of RAID devices or any general SCSI subsystem, it seems to me that the initiator used to configure and control the subsystem must have considerable visibility into the current state of the subsystem. The following example traces one possible sequence of events in the use of a RAID subsystem, and shows why the number of states needs to be expanded over George's proposal.

Event	New State
-----	-----
Power on to subsystem	Self-test in progress
Self-test completes	Ready to handle INQUIRY command
Self-configuration completes	Ready to handle any command, P-addressing valid (only)
Define first extent	At least one extent available
Define first redundancy group	At least one redundancy group available, R-addressing valid
Define first volume set	At least one volume set available, V-addressing valid
Enable RAID operation	Normal quiescent operating state
Disk failure	Degraded redundancy
Initiate rebuilding on spare	Degraded redundancy, rebuild in progress
Rebuild complete	Normal redundancy, degraded performance (because of use of spare)
Failed disk replaced	Normal redundancy, copyback in progress
Copyback complete	Normal redundancy, normal performance, erasing spare
Spare reconfiguration complete	Normal quiescent operating state

Many of the above states are volume set states, but P-extent, P-LUI, redundancy group, spare, and overall SDA states are intermixed. Considering that a RAID subsystem of moderate size might have several volume sets, dozens of redundancy groups, and hundreds of P-LUIs, it is clear that the state description needs to be more comprehensive. Perhaps we need a separate table for each of these, as proposed below.

x.0 SCSI-3 Disk Array States

The following sections describe the possible states for each of the addressable components of a RAID subsystem.

x.1 DACL States

An SDA consists of one or more addressable DACL devices at one or more SCSI device addresses. Each DACL can report the following states.

State	Condition
-----	-----
Self-test in progress	Power is applied to system, but DACL is not yet ready to respond to any command.
Ready to handle INQUIRY command	Self-test has completed, DACL is ready to respond to INQUIRY command.
Ready to handle any command	Any self-configuration operations (possibly requiring media access) are

complete, DACL is ready to respond to any command.

Hung

DACL is not responding to commands.

x.2 Volume Set States

A DACL implements one or more volume sets at one or more LUN addresses. Each volume set can report the following states.

State	Condition
-----	-----
Not available	The volume set has been defined, but is disabled.
Active	The volume set is available for normal use.
Degraded performance	Within the volume set some component is in a state that causes degraded performance. All data is still fully protected.
Exposed	Within the volume set some component is in a state that causes the loss of redundancy. All data is still valid.
Failed	Within the volume set some component is in a state that prevents a guarantee that data will not be lost.

x.3 Redundancy Group States

A DACL implements one or more redundancy groups at one or more LUN addresses. Each redundancy group can report the following states.

State	Condition
-----	-----
Not available	The redundancy group has been defined, but is disabled.
Active	The redundancy group is available for normal use.
Degraded redundancy	Something has failed, causing the protection of the data to be lost, and the system does not have a resource needed to provide protection. No data has been lost.
Rebuild in progress	Protection is being rebuilt onto a spare or repaired disk. Protection has been lost, but the system is working to rebuild it.
Degraded performance	A component that affects performance has failed, or a lower-performance spare disk unit is in use. In this state all data is still fully protected.
Copyback in progress	Data is being copied back to a repaired disk. All data fully protected. The system is working to provide full performance.

x.4 P-LUI States

A P-LUI is an underlying device controlled by the DACL. P-LUIs include the disks used to store the data, plus any other devices used in the SDA that can be addressed with the physical addressing mechanism.

Since each P-LUI is an addressable physical device in the system, its state is simply the state of the physical device. SCSI devices do not report "state" as such, but the state can be inferred from the

data returned in response to commands such as INQUIRY, TEST UNIT READY, and REQUEST SENSE.

[Not sure about this. Is the DACL capable of reporting how it is using each P-LUI, e.g. "P-LUI is in use in R-LUI n"?]

Each redundancy group can report the following states.

State	Condition
-----	-----
Not available	The redundancy group has been defined, but is disabled.
Active	The redundancy group is available for normal use.

x.5 Spare States

A spare is an addressable [something] that can be used as a spare.

Each spare can report the following states. Note that there is no provision for reporting the history of prior use of a spare. It is assumed that a spare is available for temporary use but that when the original device is replaced, the data will be copied back to the original device and the spare will be erased and made available again.

State	Condition
-----	-----
Not available	The spare has been defined, but is disabled.
Available	The spare is available for use.
In use	The spare has been allocated for use as a spare.