Technical Brief

Data Protection by EonStor® Storage Systems

Abstract

This document describes how EonStor® storage systems protect your data through various kinds of advanced RAID technologies, including RAID6, SMART technologies and Media Scan.

Infortrend® Technology, Inc.

Networked Storage Solution Provider

Revision 2.0

May, 2008

EonStor RAID systems provide the following benefits

Fault Tolerance

RAID technologies enable arrays to sustain drive failure events.

All hardware can fail. A failed component can be replaced. However, a disk drive failure destroys the data within. RAID technologies provide fault tolerance for disk drives by combining multiple disk drives into one or more logical arrays. If a single drive fails, data can still be retrieved from other members of a RAID array. A RAID1/3/5 array can sustain a single drive failure (more with RAID1 if the failed drives still have a mirror); a RAID6 array can withstand simultaneous failures of two members.

Using RAID6 for Advanced Data Protection

RAID6 allows for additional fault tolerance by using a second independent distributed parity scheme.

RAID technology is used to protect data as well as provide more efficient data access and capacity utilization with disk-based subsystems. RAID6 is essentially an extension of RAID5 that allows for additional fault tolerance by using a second independent distributed parity scheme (dual parity).
Why Use RAID6?

1. Growing adoption of less reliable drives

2. High capacity means longer rebuilding time

3. The failure rate significantly rises as the number of hard drives increases

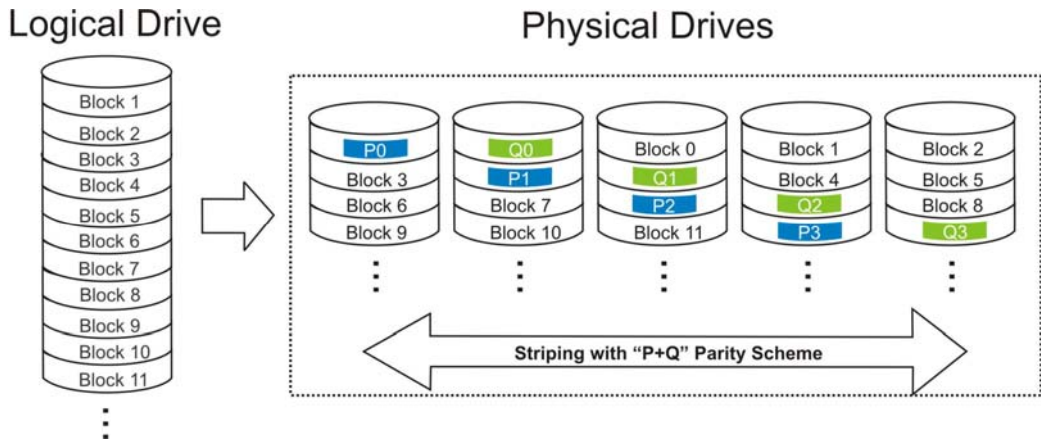

Infortrend's S.M.A.R.T technology (DrvS.M.A.R.T, SysS.M.A.R.T, IOS.M.A.R.T)

S.M.A.R.T is a set of pre-failure detection and corrective actions.

S.M.A.R.T. stands for Self-Monitoring, Analysis and Reporting Technology. It is a set of pre-failure detection and corrective actions, embodying the knowledge base derived from Infortrend's years of experiences working with customers in all aspects of disk array applications. From dealing with individual disk drive behaviors, adapting to I/O characteristics, to coping with changes in operation environments, we provide an abundance of functionality and tunable parameters that best fit customers' requirements.

Disk drives equipped with S.M.A.R.T will make reports to the host when it predicts possible failures according to the status of predetermined attributes.

Disk drives equipped with S.M.A.R.T. will make a status report to the host if it detects degradation with predetermined drive attributes. Not all failures are predictable. S.M.A.R.T. predictability is limited to the attributes that the drive can monitor. Typically, these include: head flying height, data throughput performance, spin-up time, re-allocated sector count, seek error rate, seek time performance, spin try recount and drive calibration retry count.

For SCSI drives, an industry standard specification is used as defined in the ANSI-SCSI Informational Exception Control (IEC) document X3T10/94-190.

Normally, SCSI drives with S.M.A.R.T. capability communicate a reliability condition as either good or failing. The specification provides for a sense bit to be flagged if a reliability issue exists. The host may then alert the user.

Some examples for S.M.A.R.T technologies:
(1) Disk Clone for drive relocation and data backup
(2) Smart Bypass to avoid prolonged timeout waiting
(3) Smart Cache to enhance IO performance by performing read-aheads

A few examples of S.M.A.R.T technologies are shown below:

1. Disk Clone: Individual data drives can be cloned for physically relocating disk drives or for safe-keeping purposes when an administrator discovers symptoms of disk drive defects.

2. Smart Bypass: In special applications such as video streaming, bad media on drives can cause observable delays. The S.M.A.R.T bypass quickly determines after a few mini-seconds that drive retries may ensue, and turns to other member drives for data. Instead of waiting for timeout, I/Os are quickly satisfied.

3. Smart Cache: With intelligent algorithms set for the application purposes, read-aheads can be performed, temporarily put in cache, and quickly returned to the host far faster than the actual I/O hits.

There are S.M.A.R.T options deliberately designed for various applications and require understanding of application demands.

Four configurable options relating to S.M.A.R.T function in the firmware utility: Disable, Detect, Perpetual Clone and Clone + Replace

Implementing S.M.A.R.T. Technology

There are four configurable options relating to S.M.A.R.T. function in the firmware utility:

1. **Disable** - S.M.A.R.T. function not activated.

2. **Detect** - S.M.A.R.T. function enabled. Controller will send command to enable all the drives S.M.A.R.T. function. If a drive detects a problem, the controller will send an event log.

3. **Perpetual Clone** - S.M.A.R.T. function enabled. Controller will send command to enable all drives S.M.A.R.T. function. If a drive detects a problem, the controller will send an event log. The controller will clone the drive if there is a hot-spare drive. The faulty drive will not be taken off-line, and the clone drive is still continued to operate as a spare drive. In case the faulty drive stops working, the spare drive will take over immediately. If the faulty drive continues to function and another drive fails instead, the spare drive will become active data rebuild into it.

4. **Clone + Replace** - S.M.A.R.T. function enabled. Controller will send command to enable all drives S.M.A.R.T. function. If a drive detects a problem, controller will send an event log. The controller will clone the drive to the spare drive and take the failed drive off-line.

There should not be any problem with having different brand of drives as long as they are ANSI-SCSI Informational Exception Control (IEC) document X3T10/94-190 compatible.

Media Scan

Media Scan can be used as a preventative measure to examine drives and detect bad blocks. If any bad blocks are encountered during the

scanning process, data from those blocks are automatically retrieved and stored onto undamaged sectors.

Media Scan can be used as a preventative measure to detect bad blocks on drives. During the logical drive rebuild process, even when bad blocks are encountered on yet another drive, Media Scan can ensure that the rebuild goes on and salvage most of the data.

In the worse case when a hard drive fails and the logical drive needs to be rebuilt, Media Scan can ensure the rebuild process keep on going even if bad blocks are encountered on yet another drive. The LBA (Logical Block Address) of those bad blocks will be shown and the rebuild process of the unaffected sectors will continue, salvaging a majority of the stored data. Running Media Scan on a regular basis ensures the early detection of degraded blocks and further minimizes the possibility of data loss.

A scheduler utility can be used to arrange periodical Media Scan tasks.

Infortrend's firmware also comes with a scheduler utility that performs Media Scan by preset schedules. You can arrange the schedules so that Media Scan runs when your application servers are less stressed by daily services.