



BUILDING A HIGHLY RELIABLE SAN

Focusing on the reliability of individual components often results in greater expense with no improvement in system reliability.

SATA vs. FC/SCSI disk drive reliability differences are immaterial to system reliability when comparing well designed storage systems.

System reliability is a vital component in Storage Area Network (SAN) design that keeps your production environment operating and avoids data loss and downtime. But since SANs are built using both mechanical and electronic parts, component failures due to usage, environmental factors, or manufacturing defects are not uncommon. Even in completely redundant systems, controllers can fail, fans can stop operating, power supplies can burn out, and disk drives can degrade or fail.

To eliminate downtime, IT managers should focus on total system design and how components work together to deliver reliability. A common mistake is to evaluate individual component reliability alone – you can “miss the forest for the trees” and significantly increase expenses with no improvement in overall system reliability.

SYSTEM DESIGN TRUMPS COMPONENT RATINGS

While SAN component reliability ratings are of interest, system reliability is established as a result of all these components working together. For example, a server with one direct-attached disk is only as reliable as that disk; the server’s electronics reliability rating may be very high, but if the disk fails, the system fails. However, configure two internal disks with RAID, and reliability improves dramatically; a disk failure does not impact system operation because of the “designed-in” redundancy.

Some storage vendors want SAN buyers to focus on individual disk drive reliability ratings, but good system design and RAID implementations render these statistics essentially moot. Individual disk drives are commonly assessed using the statistical calculations of Mean Time to Failure (MTTF), measured in hours, and storage arrays are assessed using Mean Time to Data Loss (MTTDL), measured in years. SATA drives currently test at 600,000 to more than 1 M hours MTTF at 100% duty cycle* (and offer a price/performance option that FC or SCSI disks cannot match). Consider a 14-drive RAID system with disk drives rated at 600,000 hours MTTF. For that system, the lowest MTTDL – the storage administrator’s key concern – in a RAID 50 environment is 70,000 years, and in a RAID 10 environment is 360,000 years.

In addition, reliability ratings should be compared in the context of capacity. FC and SCSI disk drives may be rated with higher reliability (800,000 - 1.5 M hours MTTF) than SATA drives (600,000 - 1 M hours MTTF). However, since FC and SCSI drives offer less capacity than SATA drives, it takes more disks to deliver the same capacity, result-

