

Digi-Data Best Practices: Xyratex 4835 (Sumo) Disk Enclosures under Disk Power Management

The Xyratex 4835 disk enclosure provides a rack dense (4U) array of 48 disks. The Digi-Data T-2000 RAID Controller's Disk Power Management application augments the space savings of the 4835 with potential power savings. The convergence of these two technologies makes archive and backup tier storage more viable in the data center environment where it is needed most, but where footprint and energy costs of traditional storage, with low density and unmanaged power, are prohibitive.

Disk Power Management and Scale

The Digi-Data T-2000 Disk Power Management (DPM) design goals are two fold. First, it provides a means to reduce the total power consumption of the storage system. Second, it should not reduce system availability in the attempt to meet its primary goal. While both goals have been met, the massive sizes to which the T-2000 system can scale can interfere with the power savings if DPM is not considered during storage allocation.

Disk Power Management has several rules that govern its behavior:

1. DPM is a system wide setting.
2. Disks are put into power save mode when there is no I/O activity on the RAID Group to which they belong for a configurable period of time. Not when there is a lack of activity on that particular disk.
3. There will never be a case of 100% disk power down. Several disks in the system are required to be fully powered (even if idle) to assure proper system operation. The percentage of disks required for this becomes more negligible the larger the system (based on number of disks and number of disks per shelf).
4. All disks in the system can be active and fully powered simultaneously.

While most other storage platforms can have relatively small numbers of disks in their RAID Groups, the T-2000 is capable of managing 256 disks in a single RAID Group. Using 750GB SATA II drives configured as RAID 5 or 50, the RAID Group could be as large as 180TB. This could then be presented to a host in its entirety or divided into several LUNs. Unfortunately creating such a massive RAID Group and subdividing the required LUNs will rarely be the best design when DPM is also involved. DPM will work best when LUNs with mutually exclusive data sets (i.e. weekly backup sets) are on separate RAID Groups. Fortunately, the model that is derived from a smaller scale implementation can be replicated as the system scales to assure proper operation and cost savings from DPM. The storage can also be aggregated by the host when the initial allocation is no longer sufficient.

Xyratex 4835 and Disk Allocation

Xyratex recommends that the 4835 enclosures have drives allocated in 9 disk, Disk Groups, leaving 3 as spares. These Disk Groups can then be aggregated by the T-2000 into RAID Groups based on storage needs of the application.

Example: 4835 Disk Grouping

| | |
|------------------------------------|--------|
| Qty. 4825 Enclosures | 1 |
| Drive Size | 750 GB |
| RAID Level | 50 |
| Disk Group Usable | 6TB |
| Total Number of Disk Groups | 5 |
| RAID Group Usable | 30TB |

T-2000 users must take into consideration their particular needs and use the system to aggregate the Disk Groups from the 4835 into RAID Groups appropriately (Reference T-2000 Users Manual). The T-2000 Xyratex 4835 combination is thus best thought of in 6TB increments when considering the storage needs of the host servers.

In a backup solution that uses one 4835 unit, 5 weeks of data can be stored. One 6TB LUN can be set aside for each week's data and the weekly backups can be retained for 30 days before being over written. During the data retention period the 6TB set (9 disks) can be powered down. The result is that 4/5th of the system (36 disks) are not drawing full power, but may be accessed in seconds if a restore operation is required.

The T-2000 can scale this solution to any size needed for the backup by aggregating 6TB sets in a performance solution. Using the 9 disk sets across several 4835 shelves allows much greater backup sets to be stored in the power managed environment.

Summary

T-2000 Disk Power Management can be successfully implemented in systems of any scale. It is especially effective in combination with the Xyratex 4835 rack dense disk enclosure. The best practice is to use the recommended 9 disk sets, as atomic units of storage that can be grouped together by the controller and effectively managed.