SGI DMF Storage Tier Virtualization Software

Automated Virtualized Data Management

Key Features

Maximizes storage cost savings while minimizing administrative overhead

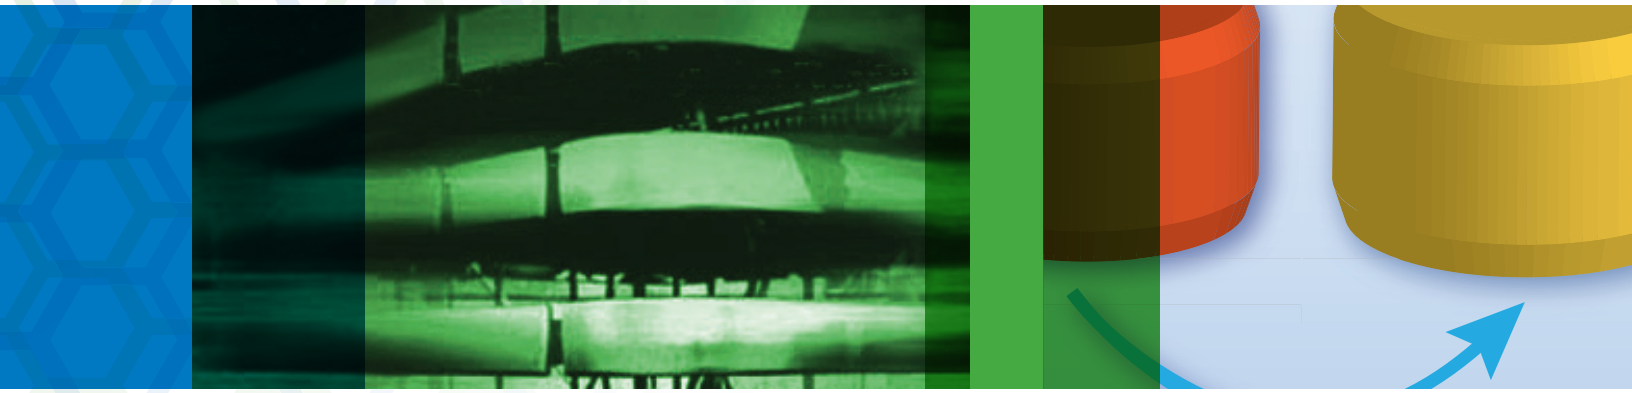

Maximize storage cost savings and minimize administrative overhead.

DMF continuously monitors and automatically moves data among storage devices with different performance and price characteristics. More expensive, faster storage can be augmented with less expensive storage like tape cheaper nearline disk or MAID, lowering costs versus an all-high performance RAID environment. It is not unusual for a DMF environment to save 60 to 90% on storage and storage management costs.

Most storage industry analysts agree that storage under management will increase by a factor of ten over the next five years, and estimate the lifetime cost of storage management at 8 to 10 times the original cost of acquisition of storage products. Because DMF automates data migration and recall, significant cost savings can be realized, breaking the cycle of ever-increasing storage and storage management costs. DMF customers have reported no increase in administrative costs even when data under management increased by a factor of twenty.

A recent study by the Clipper Group shows that tape can cost 23 times less than SATA disk and lower energy costs by a factor of 290. Savings are significantly higher when compared to high performance RAID. By integrating automated tape libraries into a tiered storage strategy, DMF allows datacenters to take advantage of the inherently green nature of tape storage not only for backup and archive, but in everyday online storage operations.

DMF Overview

DMF creates a virtualized storage pool across Fibre Channel and SAS RAID, SATA RAID, and tape presenting a single view of data to applications and users. Data appears online at all times, eliminating any need for modifications to applications or user behavior. Initial data access times can range from milliseconds for disk-based data to as little as 14 seconds for data residing in an automated tape library. Data transfer occurs at rated device speeds. In addition to managing tiered disks, DMF provides special features to integrate low cost, large scale tape libraries into an automated tiered storage management architecture.

DMF Features

DMF provides the most advanced set of features for managing tiered storage architectures including:

Automatic Data Migration

Continuous monitoring and migration of data to lower cost tiers based on administrator defined file-level and volume-level policies. File-level policies include age of file, size of file, user ID, and group ID. Custom policies can be created through user defined tags or using the DMF API.

Volume-level policies allow for “thin provisioning” of disk-based tiers. This feature allows storage capacity utilization to be optimized without heavy administrative overhead. Administrators can define high and low freespace thresholds which trigger data migration to lower tiers and increase free space on selected volumes. Once migration is triggered, file-level policies are used to select the data to be migrated.

Partial file migration allows specified region(s) of data to remain on primary disk storage as the entire file is moved down the hierarchy. Users and applications can leave part(s) of the file on primary disk storage for immediate access.

There is no interruption of service during migration operations.

Automatic Data Recall

Data that has been migrated to lower cost tiers retains the appearance of being online. File location is transparent to users and applications. When accessed, files are automatically retrieved without the need for operator intervention.

DMF provides advanced features to improve application performance in a tiered storage environment.

- Accelerated Access allows users and applications to read data from primary disk storage while the remainder of the file is recalled.
- Partial File Recall allows specified part(s) of a file to be recalled to primary disk storage while the entire file remains on a lower tier. Users and applications can identify content and determine if full recall is required.

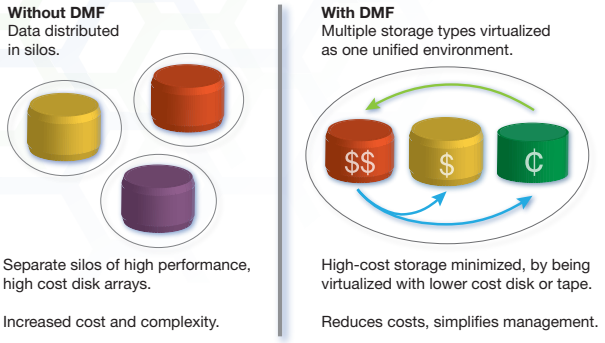

Active Copy

DMF can be configured to automatically make up to 64 copies of a file on disk or tape during the migration process, providing enhanced data protection for critical files. Locations can be specified on a per copy basis and remote copies are supported via FTP over standard IP-based networks.

Web-based Administration

DMF provides a web-based administration tool to simplify management of DMF operations. Administrators can start/stop DMF, view configuration details, view and respond to DMF alerts, and obtain reports and statistics on DMF operations. Many tape library management functions are also available including adding and ejecting cartridges, verifying data integrity, merging data from one cartridge to another, enabling and disabling tape drives and more.

Parallel Data Mover Option

Data Movers are the components of DMF that perform the actual movement of data during migration and recall operations. The basic DMF product incorporates an integrated Data Mover, allowing the DMF system to reside on a single server and minimizing the cost of a DMF implementation.

For users with higher throughput requirements, the Parallel Data Mover Option allows two or more Data Movers to operate in parallel on separate servers, increasing data throughput and enhancing resiliency.

Tape Management Features

DMF provides a rich set of features to automatically manage automated tape libraries including media verification and merging “sparse” tapes.

Availability and Data Integrity

DMF can be configured for high availability via active/passive failover.

Original data copies are maintained until migration completion is verified. For tape resident data, checksums are generated, block position is encoded and written to the tape. Additional copies are automatically used to reassemble a complete file if tape data errors are detected.

Scalability

Total volume of data is only limited by physical media capacity. DMF currently supports production customers with nearly two billion files, and greater than 100 Petabytes under management.

Global Sales and Support: sgi.com/global