Hitachi's Midrange Disk Array as Platform for DLCM Solution

Teiko Kezuka
Ikuya Yagisawa
Azuma Kano
Seiki Morita

OVERVIEW: With the Internet now firmly established as a social infrastructure, the ownership and management of diverse types of digital data is generating considerable expense. To manage this ever increasing amount of data, Hitachi proposes "DLCM (data life cycle management) solution," a comprehensive storage technique that optimizes the cost of data ownership as the value of data changes over time. Hitachi enhanced its midrange disk-array subsystems to provide a hardware platform for the DLCM solution. It provides nearline storage by supporting SATA (serial advanced technology attachment) disk drives in existing midrange storage products, and provides robust data storage solutions by supporting anti-data-tampering functions. Furthermore, with the addition of the 9585V upper-midrange model with enhanced performance, Hitachi also supplies online storage at top-class performance in the midrange market. Hitachi provides a complete lineup of products for achieving a multi-tiered storage platform.

INTRODUCTION

IN recent years, a vast array of information (content) has become accessible via the Internet, and as a result, the amount of data handled by information systems is increasing and content is diversifying. At the same time, the long-term retention of various types of data has become a new issue surrounding the information society. In the United States, for example, legislation has been passed prescribing the retention of corporate data, and in Japan, the management of customer data and personal information is becoming a major issue.

For some time now, Hitachi has been meeting the need for integrated storage management through its SAN (storage area network) solutions. As an addition to these solutions, Hitachi proposes "DLCM solution" to meet the need for long-term retention and storage operations that take data characteristics into account. The idea behind DLCM solution is to store data on optimal storage device in accordance with the value of data which changes over time, and to provide

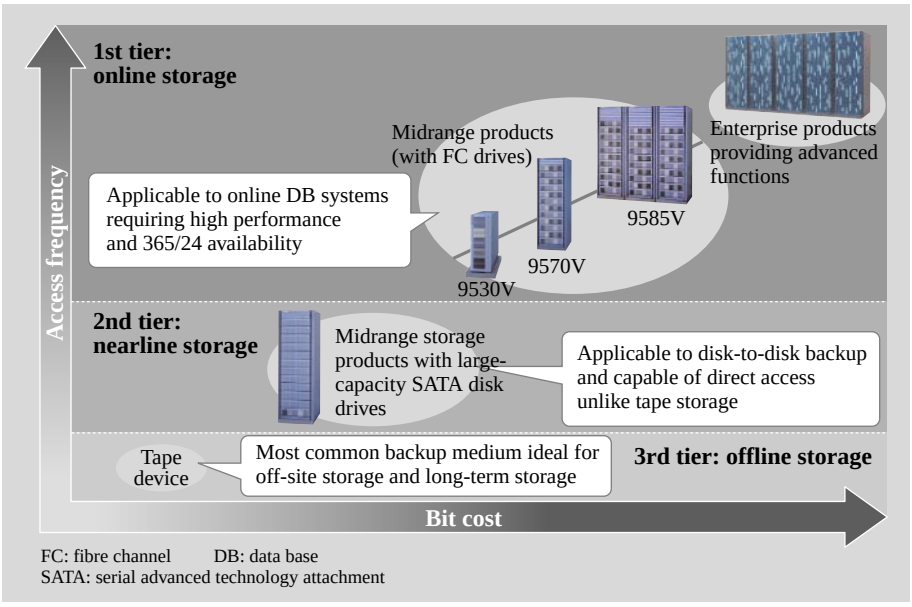

Fig. 1—Hitachi's Multi-tiered Storage Lineup and Midrange Storage Products.
Hitachi's midrange storage products support both online and nearline usages.

appropriate storage measures such as disaster recovery and data-tampering prevention in accordance with the importance of data.

PRODUCT LINEUP FOR MULTI-TIERED STORAGE

Hitachi positions different types of storage devices as "multi-tiered storage" scheme to provide a hardware platform for DLCM.

The concept of multi-tiered storage as a destination for storing data according to DLCM solution is as follows (see Fig. 1). The 1st tier is online storage, whose role is to store primary data such as for an online DB (database) system running 24 hours a day and 365 days a year. The devices making up this tier of storage are the high-availability disk arrays generally utilizing high-performance/high-reliability FC disk drives. The 2nd tier is nearline storage, which is used to store archived data used mostly for reference purposes or backup data that may need immediate access. The devices used for this tier are disk arrays with SATA (serial advanced technology attachment) disk drives as described below. Finally, the 3rd tier is offline storage that makes use of tape devices, the least expensive media. This tier is used for off-site and long-term storage of data in the upper tiers.

Hitachi provides the following products as the devices for realizing multi-tiered storage. As the top tier online storage, the followings are provided; enterprise-class storage products for large-scale centralized- or mission-critical systems, and midrange-storage products with FC disk drives for storage consolidation in a distributed environment and for stand-alone systems. As for the nearline storage, Hitachi provides midrange-storage products with SATA disk drives having advanced functions for data protection and other purposes. Hitachi also provides a suite of integrated-storage-management software for managing the data on these multiple tiers of storage.

The following describes the features and functions of Hitachi's midrange storage products.

MIDRANGE STORAGE PRODUCTS: HARDWARE

Features of Midrange Storage Product Series

Hitachi's midrange storage 9500V series are modular-type products. The components are packed in a high-density, compact enclosure that can be combined with other enclosures, supporting flexible configurations and ease of expansion.

There are three models of midrange products. The

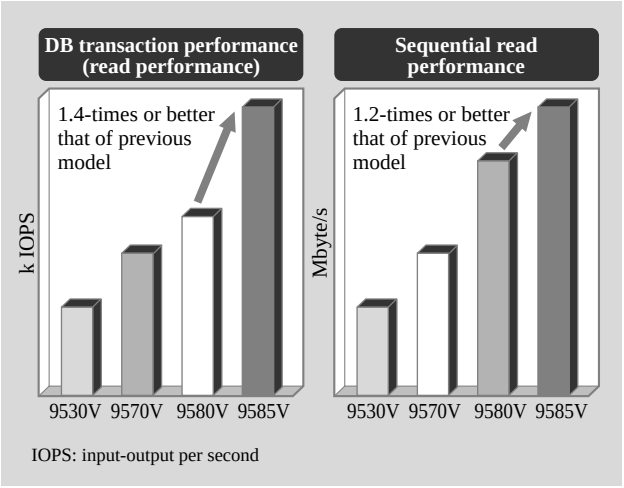

Fig. 2—Achieving High-performance in Midrange Storage. Model 9585V has the highest performance compared to the previous model and other models in the same series.

9585V model has the highest performance in the series. The storage capacity is easily and flexibly increased by adding expansion enclosures with disk drives to its dedicated controller enclosure. Next, the 9570V model features high cost performance. Disks can be placed in the controller enclosure so that the system can be expanded gradually from a small-scale configuration. The 9585V and 9570V models share a common expansion enclosure, and the 9570V can be easily upgraded to the 9585V by adding a controller enclosure. Finally, the 9530V model provides various fixed configurations, making it easy for a customer to implement a storage system by selecting one of the fixed configurations.

The expansion enclosure for the 9570V/9585V models comes in two types: one for FC disk drives and the other for SATA disk drives. The user can select either type depending on the usage and purpose. These two types of enclosures may also be intermixed in the same subsystem, meaning a single disk array can be used to provide both tiers of storage, i.e. online and nearline storage.

9585V High-performance Model

For enhancement of online storage, the controller of the 9580V model has been enhanced to provide the 9585V model with higher performance. As the highest-class model in this series, the 9585V provides a leap in performance compared to the previous model 9580V and lower-class models, and provides top-class performance in the midrange disk array subsystem market (see Fig. 2).

TABLE 1. Comparison of FC/SATA Disk Drives
A comparison of general characteristics of FC and SATA disk is shown.

Feature	FC	SATA
Capacity per drive*1	146/300 Gbyte	250/400 Gbyte
Performance	Sequential performance: high Transaction performance: high	Sequential read performance: medium Sequential write performance: low Transaction performance: low
Expected usage environment	24 h/d R/W access	330 h/month or less R/W access
Guarantee period*2	5 years	3 years
MTBF	$1{,}200 \times 10^3$ h	600×10^3 h

*1 HDD capacities currently supported by leading storage vendors or expected to be supported in the near future
*2 As stated by HDD vendors
R/W: read/write
MTBF: mean time between failures

The 9585V high-performance model features high performance and high scalability. It supports up to 8 host interface ports, 8 Gbyte of cache memory, and 449 disk drives, providing a maximum raw capacity of 107 Tbyte.

FEATURES OF SATA DISK DRIVES ON 9500V SERIES

Performance and Usage of SATA Disk Drives

SATA disk drives are generally used for nearline storage. On comparing the features of SATA disk drives and FC disk drives, the former is superior in capacity but inferior in performance and reliability (see Table 1).

The performance and reliability figures for SATA disk drives implies that SATA is not appropriate in cases FC disk drives are generally used, e.g. online transactions that involve frequent read/write operations to many disparate areas on the disk. It is not recommendable to simply substitute FC disk drives by SATA disk drives. On the other hand, the sequential read performance and capacity characteristics of SATA disk drives indicate that they would be appropriate for the usage with continuous data reads and less writes. This includes the backing up of online DBs on FC disk drives at certain points in time, the storing of data like e-mail after a fixed period of time has elapsed from its creation, and the archiving of graphical and video data to be accessed immediately upon needs.

Measures for Improving Reliability of SATA Disk Drives

The original specifications for SATA disk drives were drawn up for consumer usage as in PCs as opposed to enterprise usage on servers and storage equipment, which is why SATA has inferior characteristics when compared to FC disk drives in terms of reliability and lifetime. To alleviate this problem, Hitachi midrange storage products incorporate the following measures in the disk array controller with the aim of reducing risk in the use of SATA disk drives.

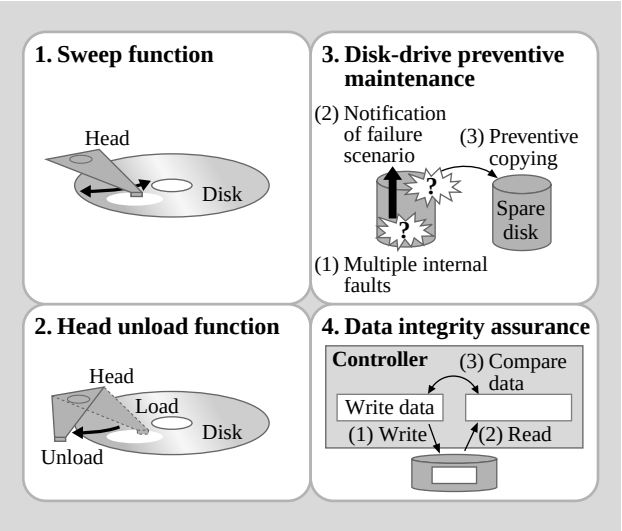

Fig. 3—Measures for Improving Reliability of SATA Disk Drives.
The 9500V series incorporates functions for reducing risk in relation to SATA-disk-drive reliability.

(1) Sweep function to reduce disk-drive failure rate

If the lubricant material coating the surface of a disk has a non-uniform thickness, it increases the risk of contact between the head and media. To counter this problem, a disk-drive has a sweep function whereby the head smoothes out the lubricant material. Hitachi's midrange storage guarantees that this function be performed at fixed intervals to reduce disk-drive failure rate.

(2) Head unload function

Hitachi midrange storage products use an unload

function on disk drives to perform periodic head unloading to achieve reduction in overall head loading time to contribute in stretching the disk-drive lifetime.

(3) Disk-drive preventive maintenance

Hitachi midrange storage supports a preventive-maintenance function that automatically spares out data on a disk where errors are frequently occurring. With SATA disk drives, such preventive maintenance function is coordinated with SMART (self-monitoring analysis and reporting technology) function to periodically collect fault information.

(4) Data integrity assurance

To improve data integrity, Hitachi midrange storage automatically generates and attaches an 8-byte data-integrity code for each 512-byte block of data. Additionally, for improving the reliability of data written to SATA disk drives which are known to have inferior reliability, the system performs a read-after-write function where it reads out written data and compares it with original data in the cache memory.

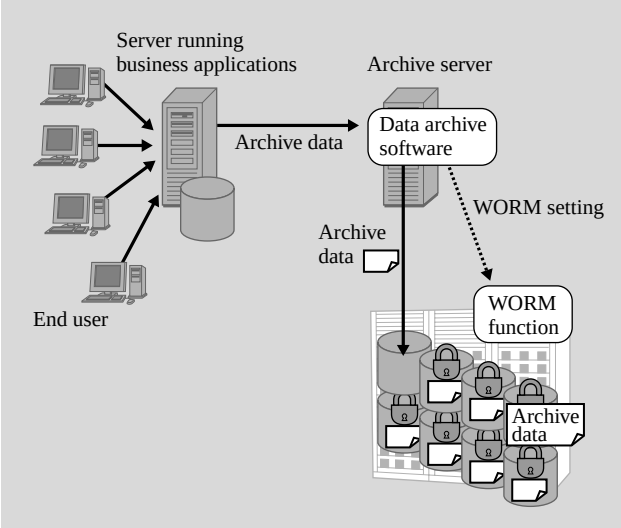

Fig. 4—Example of Data Archiving Solution Using WORM Function.

Hitachi provides DLCM solution with host-based archival software and storage hard ware-based WORM function.

ADVANCED FUNCTIONS OF 9500V SERIES

Backup and Disaster Recovery Functions

In the snapshot function, a replica of a logical unit is created at the specified time within the same array. This replica is comprised of the only differential data and updated data, therefore consuming less physical capacity. Moreover, replicas of multiple generations can be created and managed.

The volume replication function duplicates a logical unit within the same subsystem. This copied volume can be used to perform high-speed online backups in parallel with regular online operations. Because such copied volumes can be managed by a controller not directly involved with online workload, high-speed backups can be achieved with minimal effect on online I/O (input-output) performance.

The remote copy function can duplicate a logical unit to another midrange storage product connected by an FC. This real-time duplication of data enables immediate switching to a remote site in the event of a failure at the main site so that operation can continue with minimal disruption.

Data copying can also be performed between midrange storage and enterprise-class models by utilizing the advanced functions provided by enterprise-storage products.

Since all of the above functions can be performed without host interception, they can be used to construct disaster-recovery and backup systems with minimum effect on the servers and the network.

Support Function for DLCM Solution

In conjunction with SATA disk drives, Hitachi storage provides a WORM (write once, read many) function closely linked with host function to support DLCM solution. The WORM function can set a disk area allocated to the server OS (operating system) to attributes such as "read-only" or "access-denied" to prevent data tampering or unauthorized access. The period for such data protection can also be set.

This WORM function is a hardware-based function controlled by the array controller. Such function that can be used to prevent the accidental leaking or loss of data due, for example, to an erroneous operation, are provided on the hardware and independent of applications.

Fig. 4 shows an example of DLCM solution that combines the above WORM function with data archive software. In this way, the system can provide a solution for long-term and secure data storage.

CONCLUSIONS

Among the disk-array subsystems offered by Hitachi, this article described the features and functions of midrange storage products to satisfy the needs of small- and medium-scale online storage and nearline storage. For the future, Hitachi plans to further pursue its study of the problems encountered by customers in data storage and to promote the development of new and innovative storage products.

REFERENCES

(1) T. Kezuka et al., "Modular Type Storage System 'SUNRISE9500V Series'," *Hitachi Hyoron*, **85**, pp. 253-258 (Mar. 2003) in Japanese.

(2) Product introduction site (Japanese only): http://www.hitachi.co.jp/Prod/comp/storage/diskarray/index.html

ABOUT THE AUTHORS

Teiko Kezuka

Joined Hitachi, Ltd. in 2000, and now works at the Products Planning Department, the Disk Array Systems Division, the Information & Telecommunication Systems. She is currently engaged in the product planning of midrange disk array subsystems. Ms. Kezuka can be reached by e-mail at te-kezu@itg.hitachi.co.jp.

Ikuya Yagisawa

Joined Hitachi, Ltd. in 1993, and now works at the Disk Array Software Development Department III, the Disk Array Systems Division, the Information & Telecommunication Systems. He is currently engaged in the development of midrange disk array subsystem microcode. Mr. Yagisawa is a member of The Information Processing Society of Japan (IPSJ), and can be reached by e-mail at i-yagisawa@itg.hitachi.co.jp.

Azuma Kano

Joined Hitachi, Ltd. in 1994, and now works at the Disk Array Hardware Development Department, the Disk Array Systems Division, the Information & Telecommunication Systems. He is currently engaged in the development of midrange disk array subsystem hardware. Mr. Kano can be reached by e-mail at a-kano@itg.hitachi.co.jp.

Seiki Morita

Joined Hitachi, Ltd. in 1997, and now works at the Disk Array Software Development Department III, the Disk Array Systems Division, the Information & Telecommunication Systems. He is currently engaged in the development of midrange disk array subsystem microcode. Mr. Morita can be reached by e-mail at se-morita@itg.hitachi.co.jp.