

Storage Products with Scalable Extendibility and a Flexible Operation Capability

OHTANI Hiroyuki, YOSHITA Hiromitsu, YOSHIKAWA Akihiro, FUJIMORI Hideaki, TAKIYANAGI Masumi, TANAKA Yoji

Abstract

The explosive dissemination of broadband in recent years has increased the exchange of various types of data via the Internet, leading to an annual increase of 160% in the amount of data used in E-mails and movie contents. With regard to storage for use in IT systems, the issues attracting most attention are considered to be sudden increases in the data volumes, complications of management and the impacts of shutdowns. This paper discusses solutions to these issues, including “iStorage D8” featuring scalability, manageability and availability and “iStorage D1/D3” featuring high cost efficiency, easy introduction and space saving design.

Keywords

storage, disk array device, RAID, cache control, extension/expansion without system interruption, energy-saving function, distributed caches, backup

1. Introduction

The explosive dissemination of broadband in recent years has increased the exchange of various types of data via the Internet, leading to an annual increase of 160% in the amount of data used in E-mails and movie contents. It has now become an urgent task that the IT systems of enterprise deal effectively with this new environment. With regard to storage for use in IT systems, the issues attracting most attention are considered to be sudden increases in the data volumes, complications of management and the impacts of shutdowns. This paper discusses the iStorage D Series, which is our solution for the storage issues of enterprises.

2. Customer Issues

The primary issue of customers introducing storage to their systems is how to deal with rapid, unexpected increases in the data volumes. If a small storage product is introduced to reduce the initial costs and the data amount exceeds the maximum capacity of its small storage capability, replacement with an upgraded model would make it necessary to shut down the system and perform data migration. On the contrary, if a large storage product is to be introduced from the outset, a large capital investment cost would be necessary.

Secondly, when systems based on individual optimization are constructed in a haphazard manner, the increase and diversification of storages tends to raise the operations management

costs and complicates the reconfiguration of storages after task modifications. For example, increases in the operated storages make their management more complicated and storage integration aiming at TCO reduction necessitates system shutdowns for data migration as well as for detailed design work.

Thirdly, the globalization of tasks and linkages of inter- and intra-enterprise systems have extended the operating times. The storage is permanently accessed by multiple servers, and a storage problem often makes a system shutdown inevitable. Such a system shutdown may increase the risk of a decline in social confidence in the enterprise.

3. Storage Solutions Proposed by NEC

With regard to NEC’s solutions for the issues described in Section 2, we have developed the iStorage D series of next-generation scalable storages adopting the advanced virtualization technology (**Fig. 1**). From the “iStorage D8” scalable storage with flexible extendibility up to the petabyte level to the “iStorage D1/D3” high cost-efficiency storages for small- to medium-scale systems, the product line is capable of meeting needs from the entry-class to high-end users.

3.1 iStorage D8

iStorage D8 is a SAN (Storage Area Network) compatible storage with the following three attractive features.

(1) Scalability

The system may be small-started initially but the capacity can

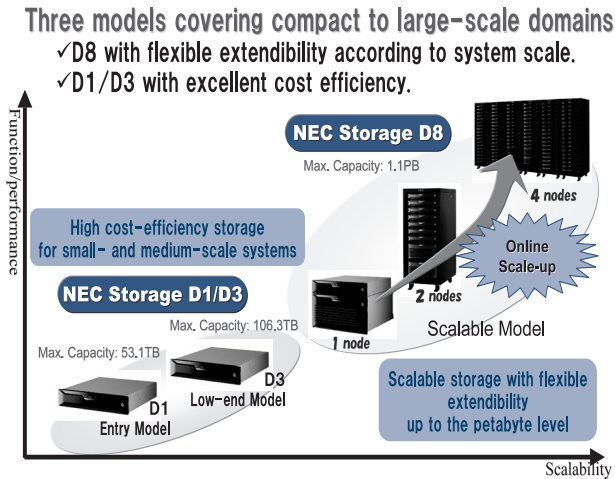


Fig. 1 Product line.

- ◆ Storage extension and flexible reconfiguration without task interruptions.
- ◆ Linear performance improvements according to an increase in task.

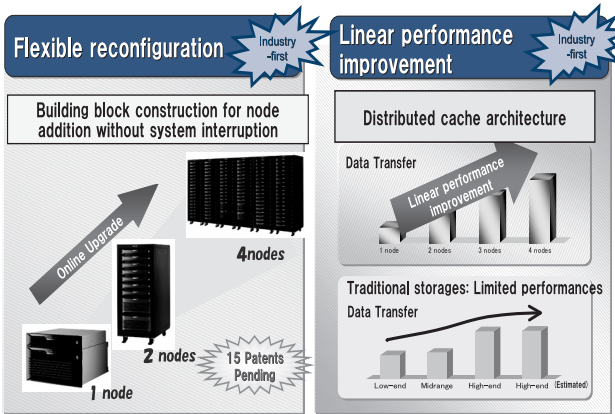


Fig. 2 Feature 1: Scalability.

be expanded to the petabyte scale and the performance can be improved linearly without interrupting tasks (Fig. 2). The storage devices are counted in “node” units and the scale may be extended from a minimum of 1 node to a maximum of 4 nodes without system interruption. This has made it possible to improve the storage performance, which used to be limited previously, linearly at the same time as increasing capacity. The technologies making this feature possible are the industry-first building-block construction and the distributed cache system (Fig. 3). It is these technologies that enabled the extension of nodes and modules with flexible reconfiguration capabilities. While accesses in the previous system were

- ◆ Building-block construction based on high-speed switch connection
- ◆ Distributed cache system using intelligent caches*

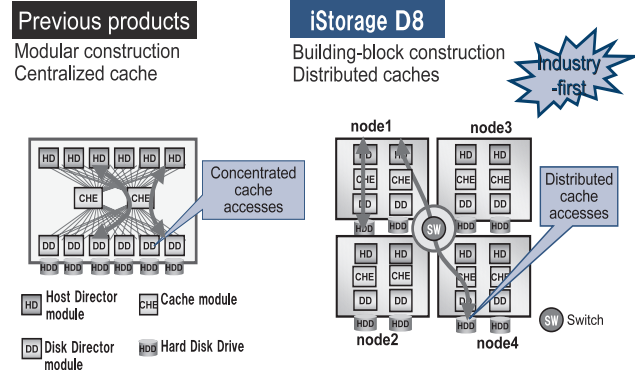


Fig. 3 Scalable technology.

centralized in a cache and the improvement of performance had been limited due this constraint, the new system can distribute accesses to each node in order to reduce competition and enable linear performance improvements. In addition, the intelligent cache system with which PMPY (Micro-Processing Unit) is incorporated in each cache module allows the complicated cache data control to be enclosed within each module, thereby contributing to improved performance.

(2) Manageability

iStorage D8 provides iStorage Virtual Storage Partitioning, which enables free allocation of resources required for tasks by using the virtualization technology. This function works by pooling the physical resources of each storage and forms a virtual storage to which the required resources are selected from the pool and allocated according to the type of task. The resources in the virtual storage can be added or modified dynamically without affecting the server or application. In addition, a user management system for setting up the management of the virtual storage is also provided. This function secures task confidentiality by restricting access that lacks administration authority to the virtual storages thus preventing erroneous operations and illegal accesses (Fig. 4). The resource allocation to the virtual storages is possible in either the performance emphasis mode that allocates exclusive resources according to the task type, or the efficiency emphasis mode that allocates shared resources (Fig. 5).

1) Performance Emphasis Method

Host ports, cache modules and disk pools are allocated to the virtual storages in an exclusive manner. Exclusive allocation of devices can eliminate competition for resources between

Storage Products with Scalable Extendibility and a Flexible Operation Capability

- ◆ Virtualization technology for free allocation of required resources
- ◆ No server/application setting change for reconfiguration.

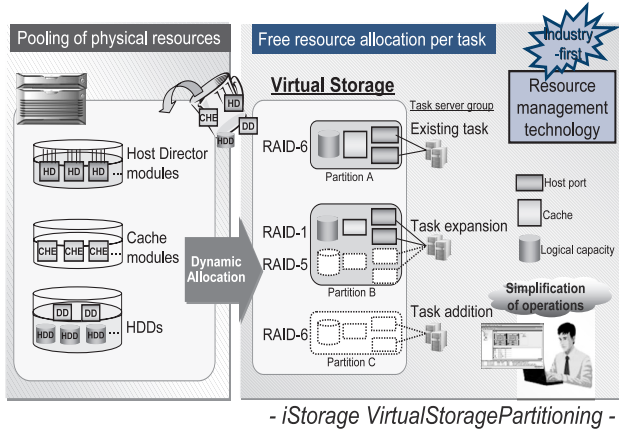


Fig. 4 Feature 2: Manageability.

- ◆ Flexible physical resource allocation for optimum storage management according to task type

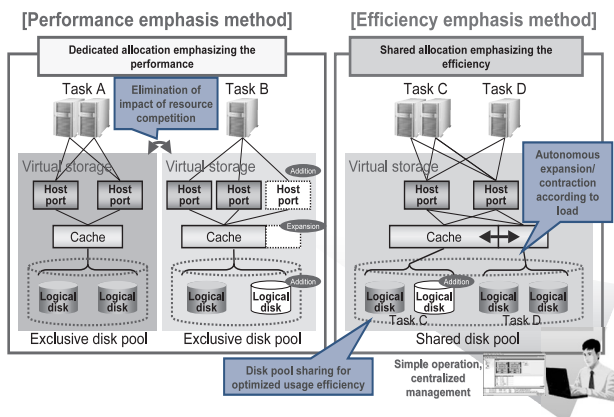


Fig. 5 Resource management technology.

jobs and can thus minimize the effects of other jobs on the performance.

2) Efficiency Emphasis Method

The resources are shared between jobs in order to optimize the usage efficiency of virtual storages. This method can improve the disk pool capacity efficiency and achieves fine control, for example by allowing the minimum size per job to be set for the cache modules that exert an important impact on the performance.

(3) Availability

The availability of iStorage D8 is improved by thorough elimination of SPOF (Single Point of Failure). Multiplexing and

- ◆ Thorough elimination of SPOF*: Multiplexing and modularization for minimized system failure

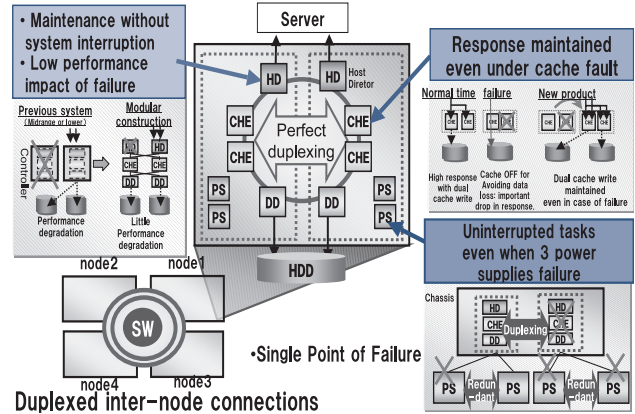


Fig. 6 Feature 3: Availability.

and modularization make it possible to minimize system shutdowns in case of a fault (Fig. 6).

In addition, RAID-6 (Redundant Array of Independent Disks) ensures task continuity even in the case of double faults in the HDDs and RAID-TM (Triple Mirror) offers both the performance and the redundancy of RAID-6. Previous products used the MPU in the controller to perform complicated RAID processing but the new product incorporates originally developed LSI “High-speed RAID engine” to enable high-speed processing.

3.2 iStorage D1/D3

iStorage D1/D3 features high cost efficiency, easy introduction and space-saving design while inheriting the operability and availability of iStorage D8. In spite of a 2U compact size, the system can incorporate up to 12 HDDs and can be extended to have up to 12 FC host interface ports. When a DE (Disk Enclosure) is connected, the number of HDDs can be extended to a maximum of 144 (iStorage D3). Both models are multiplexed and redundant in the same way as iStorage D8 and thus achieve high availability thanks to technologies such as RAID-6 and RAID-TM.

Furthermore, in order to satisfy requests to reduce introduction costs, the introduction process has been simplified to make installation easy, even without expert knowledge of the storage. The initial setup wizard of WebSAM iStorage Manager navigates the series of storage introduction procedures from physical installation to software installation, storage construction and disk allocation to the servers. The storage also has the

- Navigation from storage installation to disk allocation to servers -

- ◆ Thorough reduction of labor and time from installation to construction
- ◆ Easy introduction without expert knowledge on storage

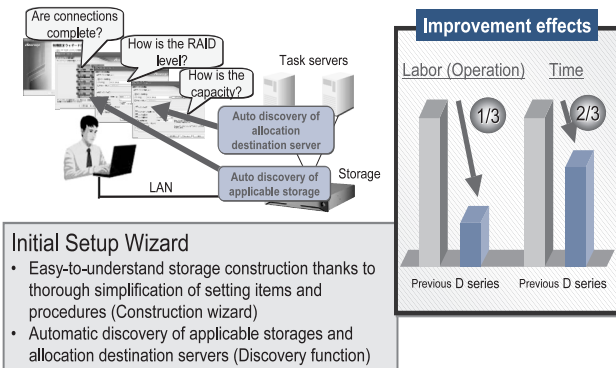


Fig. 7 Easy storage introduction.

discovery function for automatic discovery of the server to which the introduced storage and disk are to be allocated.

These functions thoroughly simplify the labor and time required for storage introduction, reducing the labor to 1/3 and the time to 2/3 compared to the introduction of previous storages (Fig. 7).

3.3 Energy-Save Function

An increase in the scale of storage systems means increased power consumption as well as running costs. The iStorage D series includes the iStorage Storage Power Conserver that utilizes the MAID (Massive Arrays of Inactive Disks) technology to save the energy consumption and make the system friendly to the environment. It controls the shutdown/startup of each HDD in the pool to which the volume belongs according to the volume usage status. It achieves power consumption reduction of up to 30% by running the HDDs that are not used permanently, for example those for backup purpose, only in the required period (Fig. 8). The iStorage Storage Power Conserver has acquired NEC's Eco Symbol as an environmentally friendly storage system product.

4. Conclusion

In the above, we have discussed the iStorage D series that solves the storage-related issues of customers. In the future, we will continue to enhance our products in a timely manner in order to match the changes in customer needs and market

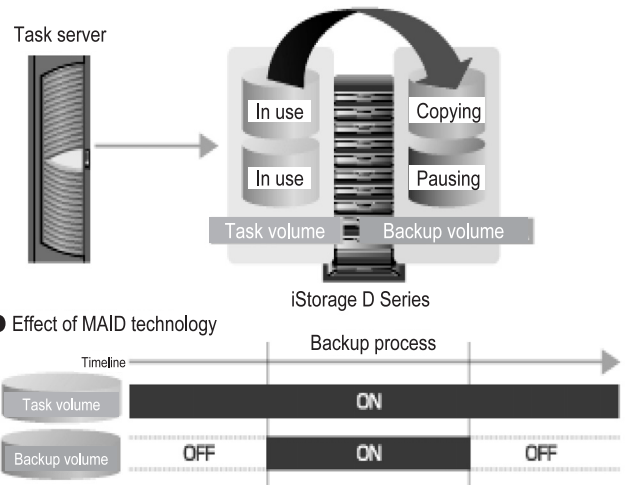


Fig. 8 Energy-save operation using MAID technology.

trends so that we may provide customers with those storage products that can satisfy their requirements.

Authors' Profiles

OHTANI Hiroyuki
Expert,
1st Computers Software Division,
Computers Software Operations Unit,
NEC Corporation

YOSHITA Hiromitsu
Assistant Manager,
2nd System Solutions Division,
NEC Engineering, Ltd.

YOSHIKAWA Akihiro
Researcher,
Software Development Division,
NEC Software Tohoku, Ltd.

FUJIMORI Hideaki,
Senior Manager,
System Storage Products Division,
1st Computers Operations Unit,
NEC Corporation

TAKIYANAGI Masumi
Technology Expert,
System Storage Products Division,
1st Computers Operations Unit,
NEC Corporation

TANAKA Yoji
Staff,
System Storage Products Division,
1st Computers Operations Unit,
NEC Corporation