High-Efficiency and High-Functionality Storages for Cloud Systems

ISHIHARA Hirokazu, TAKIYANAGI Masumi, SADANO Shogo

Abstract

Under the conditions of the current economic crisis, enterprises are expecting business innovations and efficiency improvements via IT systems and cloud computing that emphasizes use rather than ownership is enjoying increasing popularity.

The issues accompanying cloud computing, particularly storage at the data center system, include the rapid increase in the data quantity, management complications and the serious effects of operations shutdowns. This paper discusses the "iStorage D Series," which offers a solution to the above issues.

Keywords

storage, disk array system, RAID, cache control, non-disruptive extension
non-disruptive expansion, power consumption function, thin provisioning, backup

1. Introduction

In order to overcome the "once-in-a-century" economic crisis, customers are cutting the costs of existing systems and are expecting much from IT for its possibility of becoming the basis for innovating and improving the efficiency of current business platforms. These expectations have led to the popularization of cloud computing that shifts the emphasis from ownership to usage. It is now a matter of some urgency for enterprise IT systems to meet these big change in the business environment.

The issues with cloud computing, particularly in the case of storage in the data center system are the rapid increases in the data quantity, complications of management and the serious effects of operations shutdowns. In this paper, we will discuss the "iStorage D Series," which is our solution to the above issues.

2. Customers' Issues

The first of the issues troubling customers introducing storages is how to cope with rapid increases in data quantity that are hardly predictable. If a customer introduces a small storage product aiming at a reduction in the initial costs but the data quantity exceeds the maximum capacity of the small storage, the customer should purchase a higher storage product and a system shutdown then becomes necessary to enable data migration. On the other hand, if a customer introduces a large storage product by anticipating a need for headroom, the initial costs would become high.

With regard to the second issue, if individual optimization results in crowding of many systems, increase and diversification of storages increase the administration cost and complicated modifications of storage configurations become necessary when there is a change in the business. For example, an increase in the number of operated storages complicates the management. Storage integration using a large storage for the purpose of TCO reduction necessitates system shutdown to enable data migration as well as meticulous design work.

The third issue is related to the extension of operation time due to the globalization of businesses and linkages of systems between as well as inside enterprises. The storages are always accessed by multiple servers, so if the system should be shut down due to a storage failure, such a shutdown presents a risk of reducing social confidence in the enterprise. The need for non-disruption is particularly noticeable in operating the cloud platform.

3. Storage Solution Proposed by NEC

As a solution to the issues discussed in section 2 above, we at NEC have developed the "iStorage D8/D3-30 Series" of next-generation scalable storage products that feature our ad-

vanced virtualization technology (**Fig. 1**). The product line includes "iStorage D8," a scalable storage that can be extended flexibly to the petabyte class and the "iStorage D1/D3" storages that feature high cost efficiencies to meet all anticipated requirements from the entry class to the high end.

3.1 iStorage D8

iStorage D8 is a SAN (Storage Area Network) compatible storage product. It has the following three significant features.

(1) Scalability

Capacity increases are possible up to the petabyte class without disrupting the business efficiency (**Fig. 2**), as also are linear performance improvements in order to advance a system that began in a small way at introduction. Extension of the unit of storage, called the node, from a minimum of 1 node to a maximum of 4 nodes without disruption. This procedure not only increases the storage capacity but also enables linear improvement of a storage performance that continued to be static even after extensions were adopted.

As the technologies making this feature possible, we adopted the building block structure for the first time in the industry and the distributed cache system (**Fig. 3**). Node and module expansions were made possible using these technologies, thus allowing the system configuration to be modified flexibly. In addition, the previous problem of static performances due to the concentration of accesses to the caches could also be solved and linear performance improvement has become possible because the concurrences of accesses were reduced by distributing them between nodes. Furthermore, the intelligent cache system with which the MPU (Micro Processing Unit) is incorporated in each cache module is adopted in order to execute complex cache data controls inside each cache module and to thereby improve the overall performance.

(2) Manageability

With iStorage D8, iStorage Virtual Storage Partitioning is provided using the virtual storage technology to make possible the free allocation of resources required by businesses. This function pools physical resources in the storages and configures virtual storages to which the resources required for each task type in the pool is allocated. Addition and modification of resources to or in the virtual storages then become dynamically possible, without affecting the servers and applications. A user management function is also provided for setting the management of each of the virtual

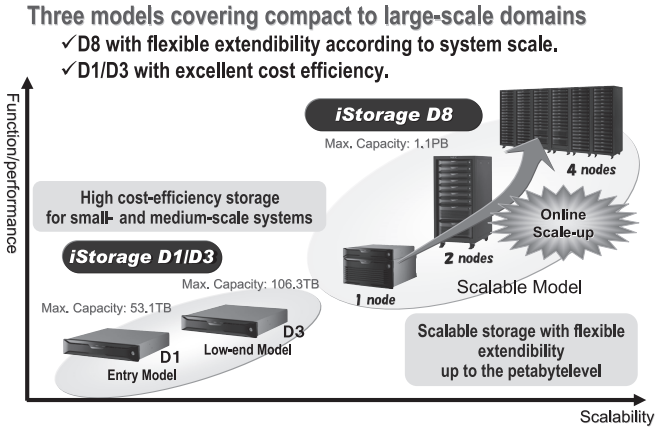

Fig. 1 Product lineup.

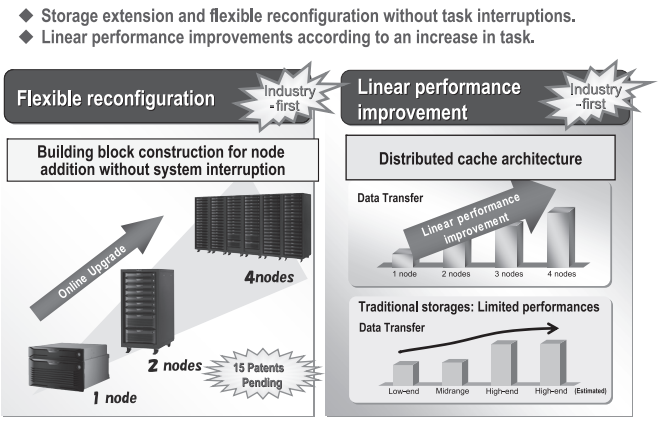

Fig. 2 Feature (1) Scalability.

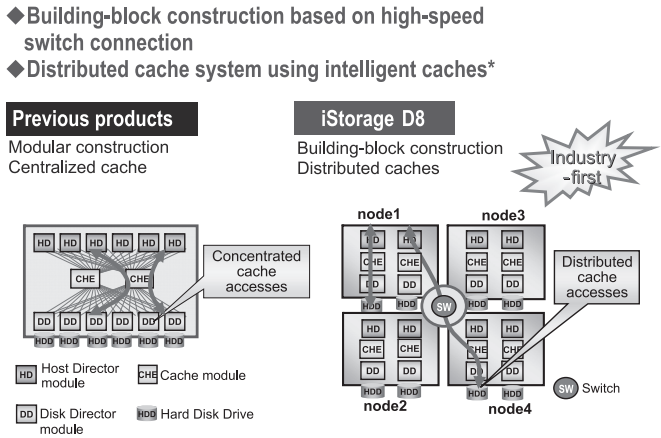

Fig. 3 Scalable technology.

storages. This function limits the access of each user to a virtual storage without the management authority and secures tasks confidentiality by preventing erroneous operation and illegal access (**Fig. 4**).

Resource allocation in the virtual storage is possible either by the performance-emphasized method that allocates exclusive use of resources according to the type of tasks or by the efficiency-emphasized method that allocates shared resources (**Fig. 5**).

1) Performance emphasized partitioning

The host port, caches and disk pools are allocated to a virtual storage for exclusive use. Allocation of exclusive-use resources can eliminate resource competition between tasks and thus minimize the influence on the performances of other tasks.

2) Efficiency emphasized partitioning

The resources of a virtual storage are shared by the tasks so as to optimize the usage efficiency. This method can improve the capacity efficiency of the disk pool and applies meticulous control, for example setting the minimum size allocated per task for the caches exerting important effects on the performance.

In addition, with the iStorage D Series, the capacity virtualization technology (thin provisioning) is added to the above function to allow the customer capacity to be configured flexibly. With thin provisioning, a virtual large-capacity logical disk is given to the operating software at the beginning and the actual capacity is allocated automatically after the fact according to write. In a case when the actual capacity becomes insufficient, capacity from the dynamic pool is added sequentially and a warning requesting HDD addition is generated only when the dynamic pool capacity is about to run out (**Fig. 6**). Therefore, the actual capacity configuration in the initial stage can be reduced and all that is necessary is to add HDDs according to the business expansion needs of the customer. Thin provisioning is naturally applicable both to the performance-emphasized and efficiency-emphasized partitioning.

These functions are particularly optimum for cloud computing that may necessitate large increases in capacity or scale. They will make it possible to select the optimum capacity and optimum performance automatically without shutting down the equipment.

(3) Availability

We eliminate the SPOF (Single Point of Failure) thoroughly to improve the availability that is a factor required for the cloud platform performance. System shutdown due to unexpected faults can be minimized thanks to redundancy and modularization (**Fig. 7**).

In addition to the RAID-6 (Redundant Array of Independent Disks) that enables continuance of business even in the case of dual HDD faults, we implemented RAID-TM (Triple Mirror) that has both the performance of RAID-1 and redundancy of RAID-6 (**Fig. 8**). Complicated RAID computations have previously been done using the MPU in the controller, but these are now done at high speed using the "fast RAID engine," which is an NEC developed LSI.

3.2 iStorage D1/D3

The "iStorage D1/D3" achieves high cost efficiencies, easy introduction and space-saving design at the same time as inheriting the high operability and availability of iStorage D8. A compact enclosure with a 2U size accommodates up to 12 HDDs and up to eight 8 Gbps FC host interface ports, and the number of HDDs can be increased up to 144 by connecting DEs (Disc Enclosures) (iStorage D3). These models meet the high availability requirement of the cloud platform by adopting similar duplication/redundancy to iStorage D8 and employing technologies such as RAID-6 and RAID-TM.

To meet the requirement for introduction cost reduction, iSCSI-dedicated models are available using the GEther interface in place of the storage-dedicated interface (FC) (iStorage D1/D3). In particular, the D3 Series have four 10 Gbps Ether interfaces to offer a performance and affinity that are not inferior to the traditional FC interface, making possible reduction of the cost for the network itself as well as the time taken for construction (**Fig. 9**).

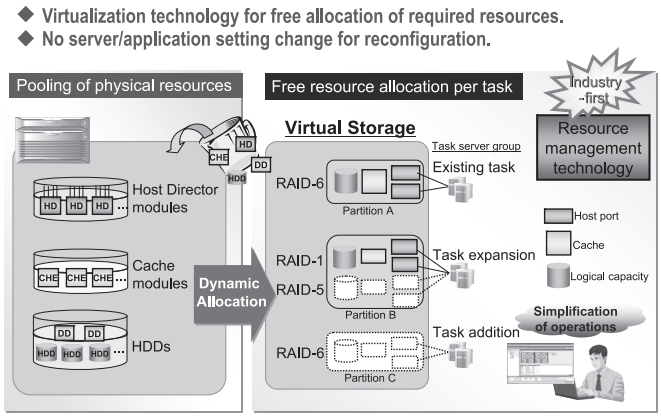

Fig. 4 Feature (2) Manageability.

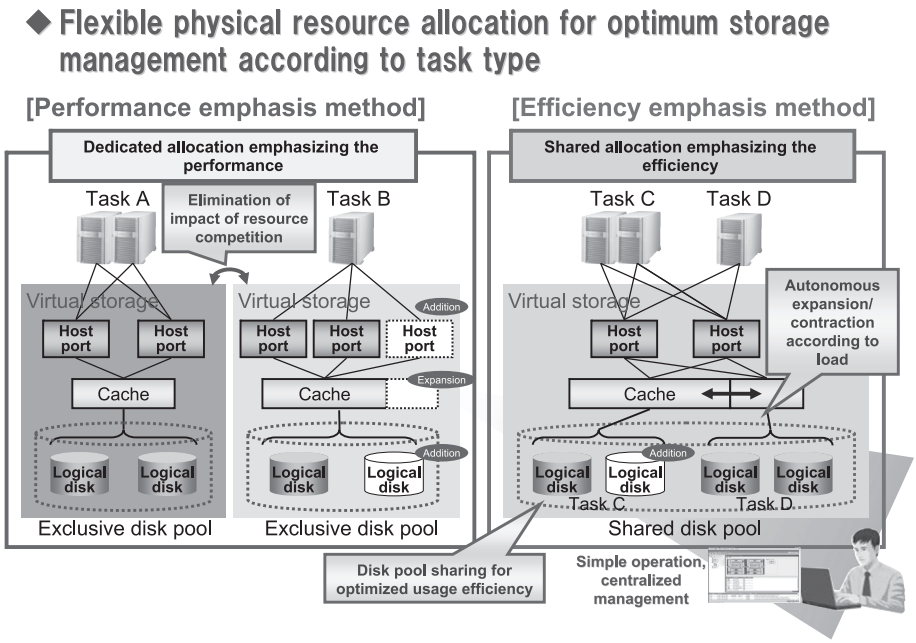

Fig. 5 Resource management technology.

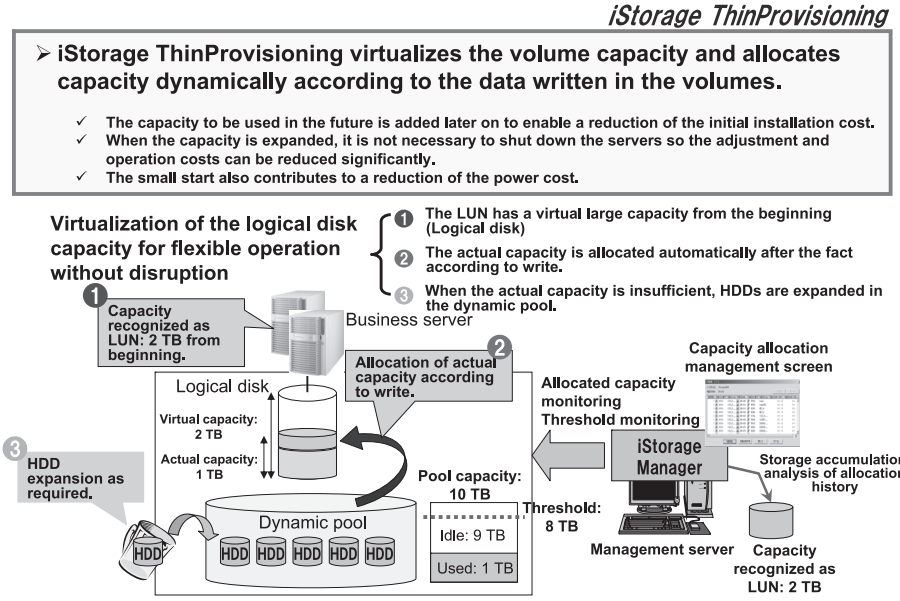

Fig. 6 Thin provisioning.

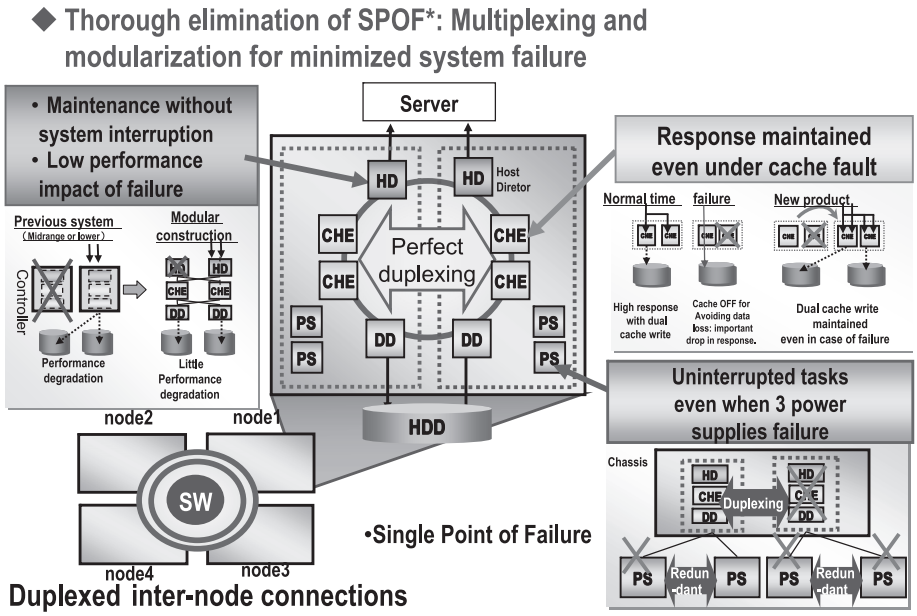

Fig. 7 Feature (3) Availability.

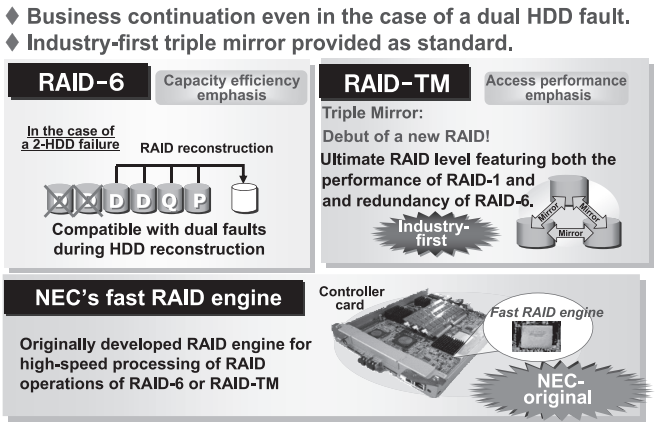

Fig. 8 RAID-6 and RAID-TM.

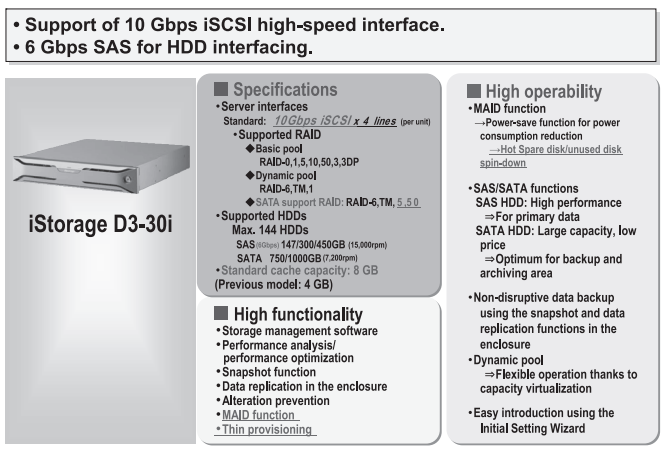

Fig. 9 iSCSI Storage.

3.3 Energy-saving Function

When the scale of a storage system increases the power consumption is also increased as well as the running cost. The iStorage D Series employs the iStorage StoragePowerConserver that applies the MAID (Massive Arrays of Inactive Disks) technology to enable environmentally-friendly, energy-saving operations for the system.

The iStorage StoragePowerConserver controls the stop/start of the operation of the HDDs in the pool according to the usage status of the volumes allocated to the pool. This procedure can reduce the power consumption by up to or more than 30% because the HDDs that are not used permanently, such as the backup HDDs and spare HDDs, are run only when necessary (**Fig. 10**). The iStorage StoragePowerConserver has acquired NEC's Eco Symbol as an environmentally-friendly storage system product.

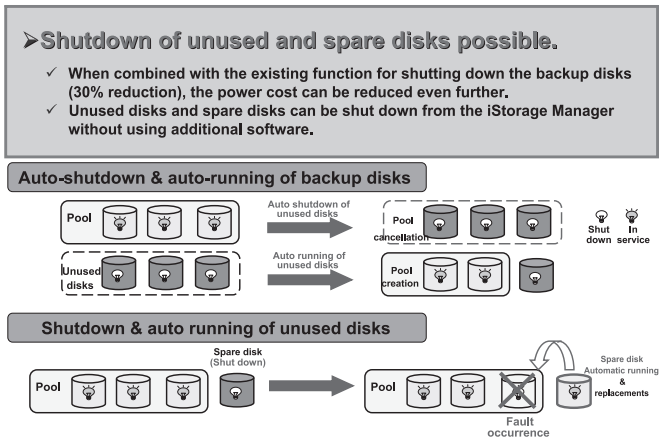

Fig. 10 Progress of MAID technology.

4. Conclusion

In the above, we discussed the "iStorage D Series" solution for the storage issues of customers by focusing on its features. In the future we intend to enhance these products in a timely manner in accordance with changes in customer needs and market trends. We will thus be able to provide our customers with storage products that will continue to offer full satisfaction.

Authors' Profiles

ISHIHARA Hirokazu
Chief Manager
IT Hardware Operations Unit

TAKIYANAGI Masumi
Manager
IT Hardware Operations Unit

SADANO Shogo
Assistant Manager
IT Hardware Operations Unit