iStorage M5000 Providing a High-Reliability Platform for the Cloud Environment

SADANO Shogo, HAGIWARA Hiroyuki, SOEDA Naoki, HONMA Kazuya, OHTANI Hiroyuki

Abstract

The rapid increase in the amount of data by ICT systems is tending to focus the roles of the storage systems. It is not enough that a storage has high reliability but assuming its use in the cloud environment, it also has to achieve a stable performance in virtualized environments and in performing multiple tasks. iStorage M5000 can provide a high-reliability platform using the sophisticated reliable technologies developed from our experience gained with mainframes and the X4 (by four) architecture. Stable performances are thereby secured even in virtualized environments and when performing multiple tasks.
This paper introduces new features that are developed based on the requirements of NEC Cloud IaaS.

Keywords
iStorage, storage, disk array, cloud, virtualization, backup, high reliability, high availability, high extensibility, stable performance, fault-tolerance, optimal deployment, OpenStack

1. Introduction

The virtualized platform for server integration used in the field of data center for providing cloud services have been required recently to provide both the capability of processing a large amount of data and the high reliability that can avoid processing interruptions.

The issues that are especially important for the storage systems in high-reliability cloud platforms include: the significant impact of task interruption, the need for stable performances, design complexity and development, complications affecting management and backup operations.

In this paper, we propose iStorage M5000 as a solution for the above issues.

2. Storage Issues Posed by High-Reliability Cloud Platforms

One of the most important storage issues in high-reliability cloud platforms is the big impact that results from storage failures. Since virtualized integration runs multiple task services on a single storage, any storage failure tends to impact a large area.

The second issue is how to secure a stable performance against an increase in the data volume. When a new service is added to an existing storage, the performances of existing services are sometimes affected. In addition, a sudden increase in the load of a specific service may also affect the performances of other running services.

The third issue is the fact that designing an optimized storage for each service leads not only to a diversification of configurations but also complicates the configuration changes as well as the design, development and operational management. Furthermore, handling large volumes of data takes time to complete backup and complicates the operational management such as the backup of the job scheduling.

3. Features of iStorage M5000

As NEC's solution to the issues discussed in Section 2, we have developed iStorage M5000 (hereafter the M5000) by fusing the iStorage A series for mainframes and the iStorage M series for open systems and adding an advanced virtualization technology (**Fig. 1**).

With the maximum number of host ports of 64 and the maximum number of drives of 1,536, the M5000 features superior scalability to the lower models (M11e to M710).

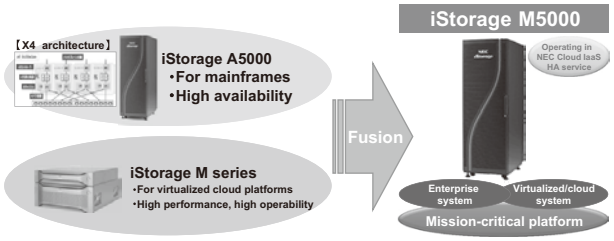

Fig. 1 iStorage M5000.

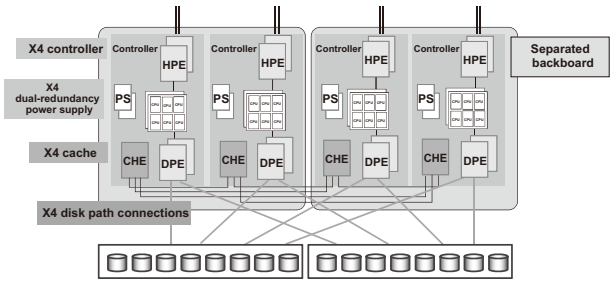

Fig. 2 X4 architecture.

3.1 High Reliability/Availability

(1) X4 (by four) architecture

Developed for use in storages for mainframes in order to avoid interruption of system operations, the NEC-original "X4 architecture" technology (**Fig. 2**) enhances system reliability by using four controllers incorporating caches and eight sets of dual-redundant power supplies and by separating the backboard into two parts.

This design maintains the read/write performance of the storage in the case of a fault because the redundancy of the caches is not lost. Operations may then be continued also in the case of multiple failures or of a backboard failure.

As described above, the excellent reliability and availability of the M5000 makes it ideal for use in important systems such as for mission-critical and virtualized integration platforms.

(2) Nondisruptive maintenance

All of the major components of the M5000 can be exchanged online. The backboard is separated into two parts so that, in the case of a fault, a failure component can be changed while continuing the operations (**Fig. 2**).

When the control software (firmware) of each controller is applied, the redundancy conditions of the connection paths between the servers and storage are checked and the connection paths are switched automatically in linkage with the servers, safely and simply without interrupting operations.

In this way, the M5000 enables nondisruptive maintenance at a high level.

(3) High-availability service

To maintain normal functioning and optimum operation of the storage system, a proactive maintenance service is performed at the same level as the periodical preventive maintenance generally adopted by open systems (maintenance service detecting the fault location and performing necessary fault recovery work).

In this way, the M5000 makes it possible to provide the high-availability services required for cloud environments stably and continuously.

3.2 Securing Stable Performance

(1) I/O Control

This function controls the service level provided for each tenant by setting the upper and lower limits of the IOPS (Input/Output Per Second) of each logical disk created in the M5000 and handling the I/O flow from the business servers to the storage.

- **Upper Control**

  This function limits the I/O from and to each logical disk so that they do not exceed the specified upper limit. This procedure can reduce the impact on the performances of other logical disks in the same pool (RAID group) even if large numbers of I/O are issued to a specific logical disk (**Fig. 3**).

- **Lower Control**

  This function limits the I/O from and to each logical disk in the pool so that the I/O from/into it does not go below the specified lower limit. An example of its application is to reserve the I/O of critical tasks by reducing the I/O of ordinary tasks (**Fig. 4**).

(2) Cache Partitioning

This function divides each of the M5000 cache memories into several segments (cache partitioning). This procedure

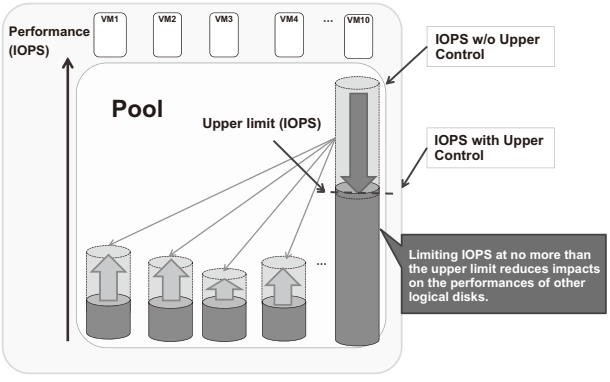

Fig. 3 Upper control.

ensures stable operation of the virtualized environment by restricting the share of the virtualized environment occupied by each tenant (task) to reserve the I/O bandwidth (**Fig. 5**).

Introducing this function brings about the following effects.

- **Securing stable performance of the virtualized environment**
  The impact between tenants in a virtualized environment can be eliminated in order to secure the stable performance of the virtualized environment.
- **Elimination of impacts on the performance of working tasks caused by the addition of tenants**
  The addition of new tenants is possible without affecting the performances of running tasks,
- **Elimination of impacts on performance of working tasks caused by an increase of data**

Even if the amount of data handled by the tasks increases, stability can be maintained by expanding the I/O bandwidth without affecting the performances of other tasks.

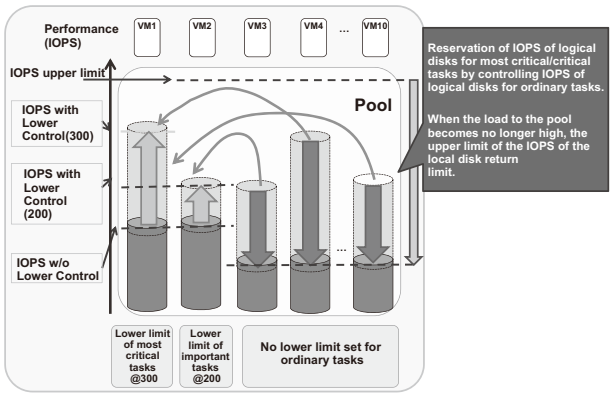

Fig. 4 Lower control.

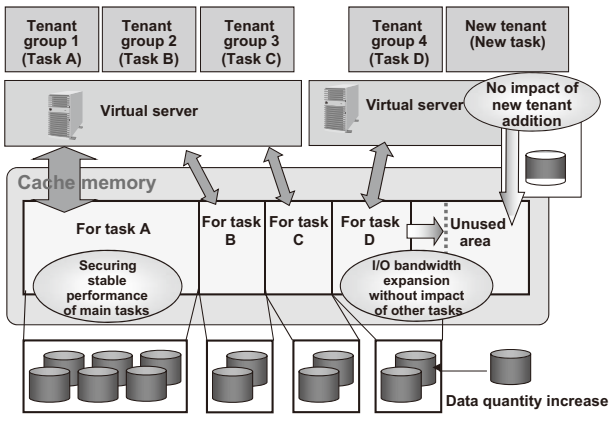

Fig. 5 Cache partitioning.

3.3 Simplification of Storage Design, Development and Operational Management

(1) Performance Optimization

For maximum use of drives (SSD, SAS, NL-SAS) with variable costs and performances, this function autonomously relocates each block of the data stored in the logical disk to optimum drives according to the data access frequency (**Fig. 6**).

It can improve the response of business systems and reduce costs because frequently accessed data may then be placed in high-performance drives and less accessed data can be placed in large-capacity drives.

The data relocation is performed automatically so the operations management is simplified.

(2) Management software

Integrated management of the configuration of the M5000 is possible based on with the WebSAM SigmaSystemCenter (SSC).

Features according to the SLA (Service Level Agreement) such as “Gold”, “Silver” and “Bronze” can be set to the data stores (LUN). For example, specifying features at the time of creation of a virtual machine can simplify its creation by using optimum data stores and this can facilitate the operations management (**Fig. 7**).

Linkage with the SSC also enables integrated management of the configuration of servers and storage is there-

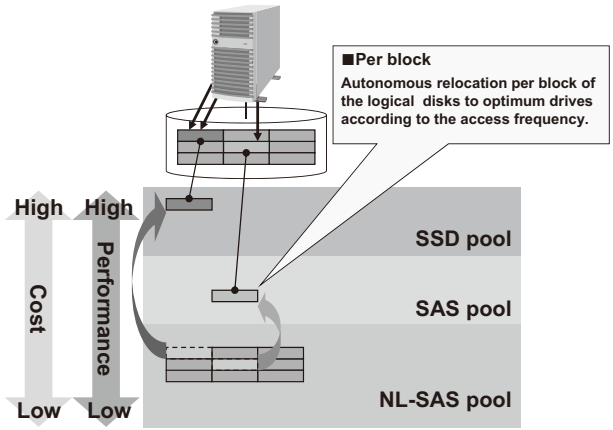

Fig. 6 Performance Optimization.

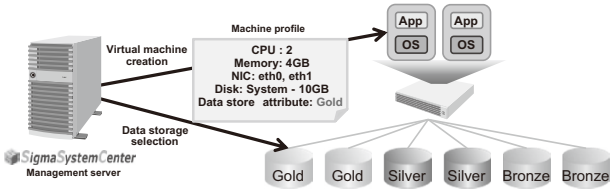

Fig. 7 With WebSAM SigmaSystemCenter.

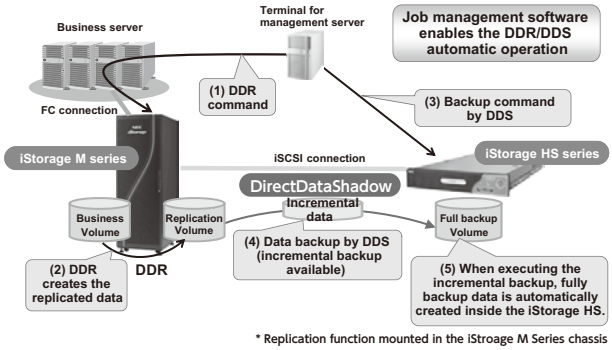

Fig. 8 Illustration of high-speed backup function.

by possible on the web console. For example if there is a task with a slow response, the logical disk used for the task can be determined immediately.

In addition, bottleneck analyses can be performed easily on the SSC by comparing the loads on the server and the storage operation of the task.

(3) Fast backup of a large amount of data

The fast backup of a large amount of data is possible by direct connection between iStorage M5000 and one of iStorage HS series, which is a backup storage featuring a de-duplication function and an excellent data compression effect (**Fig. 8**).This procedure also makes the backup management server and software unnecessary and reduces the initial installation and operating costs of backup.

(4) OpenStack support

The M5000 is compatible with the OpenStack, open source software for building IaaS platform, by providing Cinder driver. The M5000 offers a high reliable and performance stable storage service is open IaaS platform,

4. Conclusion

In the above, the authors introduced functions of the iStorage M5000 packaging the NEC Cloud IaaS. In the future, we will continue to provide storage products to support our customers by the timely enhancement of products according to changes in their needs and in market trends.

Note: The functions described in sections 3.2(1) “I/O control”, 3.3(1) “(1) Performance Optimization” and 3.3(3) “(3) Fast backup of large amounts of data” are scheduled to be supported from FY2015.

* OpenStack is a registered trademark of OpenStack Foundation.

Authors’ Profiles

SADANO Shogo
Assistant Manager
7th IT Platform Department
IT Platform Division

HAGIWARA Hiroyuki
Manager
7th IT Platform Department
IT Platform Division

SOEDA Naoki
Manager
7th IT Platform Department
IT Platform Division

HONMA Kazuya
Manager
7th IT Platform Department
IT Platform Division

OHTANI Hiroyuki
Manager
7th IT Platform Department
IT Platform Division

The details about this paper can be seen at the following.

Related URL:

http://www.nec.com/en/global/prod/storage/

Information about the NEC Technical Journal

Thank you for reading the paper.
If you are interested in the NEC Technical Journal, you can also read other papers on our website.

Link to NEC Technical Journal website

Japanese

English

Vol.9 No.2 Special Issue on Future Cloud Platforms for ICT Systems

Vol.9 No.2
June, 2015

Special Issue TOP

NEC Information