Global Applications for Storage Solutions

Toru Horimoto
Michael C. Hay

OVERVIEW: Cloud computing is the subject of high expectations because of its potential to play a key role in the use of IT to restore the international economy to robust growth and Hitachi's storage solutions have gained a strong reputation among international users for their virtualization technology which is part of the technical underpinning of the cloud. Hitachi's virtualization technology can optimize the total cost of a system with features that include improved storage utilization efficiency and reduced operation and management costs as shown by an example installation at the University HealthSystem Consortium which achieved savings on storage capacity of approximately 26% while also simplifying operation and maintenance work such as volume design and allocation. Hitachi is working to meet the expectations of its international customers through further development of storage virtualization technology to achieve even greater efficiency and operational cost savings.

INTRODUCTION

THE October 15, 2008 edition of the Japan Broadcasting Corporation's (NHK) "Today's Close-up" current affairs program featured an item entitled "The Impact of the 'Cloud'—New Revolution in Information Technology." To recover from an economic crisis that has been described as a once in a hundred years event and restore robust economic growth, business efficiency improvements and innovation that take full advantage of IT (information technology) are essential. Its three key advantages of cost efficiency, business flexibility, and streamlined installation have

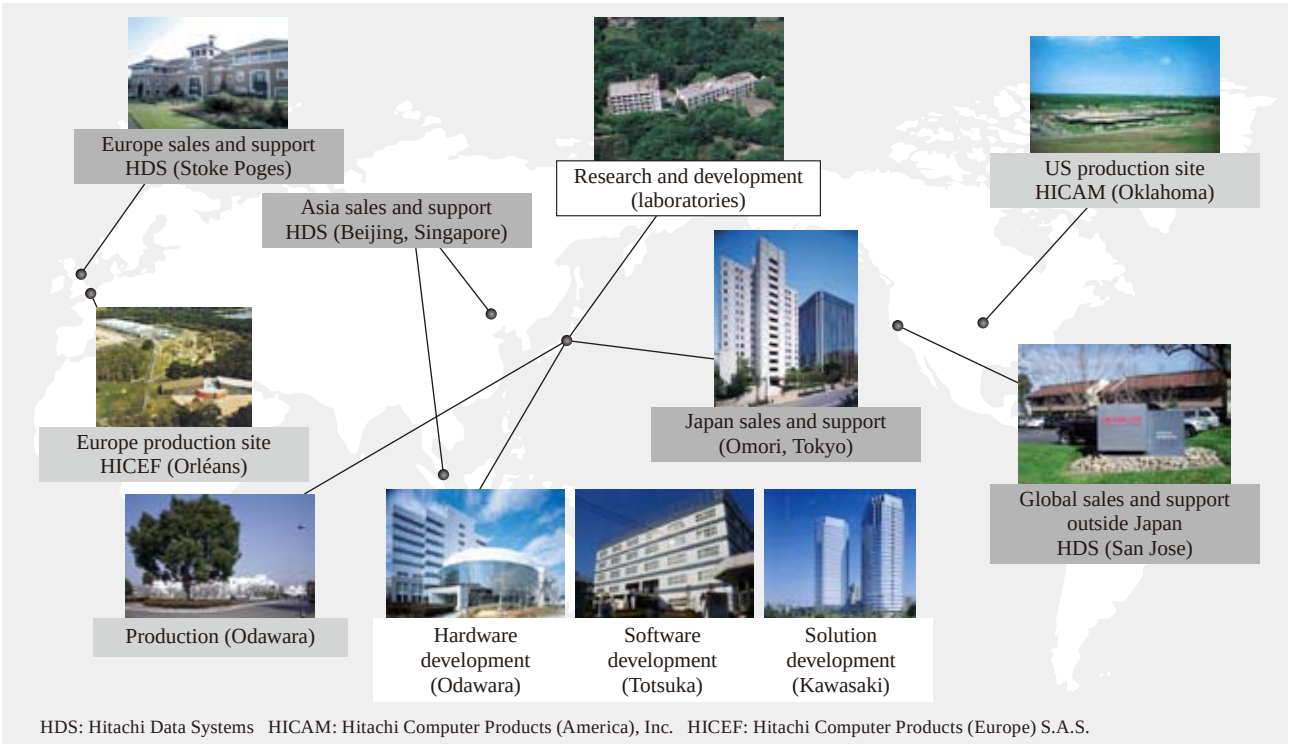

Fig. 1—Global Scope of Hitachi Storage Solutions.
Research and development is centralized in Japan while marketing, production, sales, and support sites are located in the USA, Europe, and Asia to support global customers.

made cloud computing the subject of high expectations around the world because of its potential to play a key role in the utilization of IT and the concept is becoming familiar to people in Japan through reports in the mass media and elsewhere.

The technology that underpins cloud computing is "virtualization." Virtualization enables the integration of IT resources and improves utilization efficiency and operational flexibility. Although server virtualization has gained much of the attention, the maximum benefits are achieved when it is combined with storage virtualization.

This article describes the origins of Hitachi's storage solutions whose virtualization functions are highly regarded around the world as well as details of the technology's latest global applications and its future outlook.

LATEST STORAGE SOLUTIONS

The roots of Hitachi's storage solutions business go back to June 1967. In that year, Hitachi became the first Japanese company to develop and bring to market a magnetic disk drive with a capacity of 7.25 Mbyte. This product was recognized as an "Information Processing Technology Heritage" by the Information Processing Society of Japan in February 2009[1]. Taking the view that storage is a form of social infrastructure, Hitachi has since that time worked hard on research and development based on a philosophy of "taking a rigorous approach with an uncompromising attitude to quality and reliability" and has built up operations with a global scope to meet the expectations of customers from around the world (see Fig. 1).

In the area of research and development, Hitachi works hard to deliver innovative storage solutions that anticipate customer needs based on proactive marketing that brings researchers and designers face to face with leading customers in Europe, America, and Asia.

Core Products for Storage Solutions

Fig. 2 shows the core product range of Hitachi storage solutions.

The product range is extensive including enterprise storage that supports environments which combine mainframes and open systems, high-performance mid-range storage that provides a high degree of scalability in open system environments, file storage that connects directly to the network, and appliances that are designed for archiving large amounts of content.

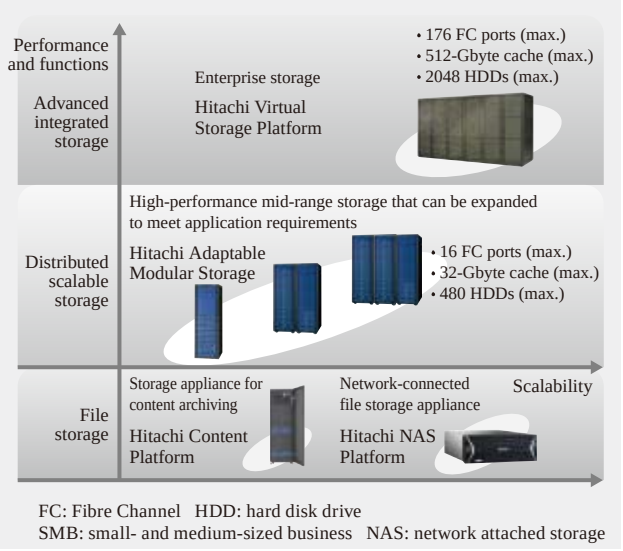

Fig. 2—Storage Solution Range.
Hitachi offers a wide range of products from enterprise to entry-level to meet diverse customer requirements.

Total Cost Optimization Using Virtualization

As the business environment undergoes major changes, the data being handled by information systems is growing rapidly in volume and becoming more diverse. In addition to the secure and reliable storage of data, other demands being placed on storage systems include improving the efficiency of utilization and reducing the workload associated with operation and management. The aim of storage virtualization is to optimize the total cost of storage by improving utilization efficiency and reducing the cost of operation and management, hardware, and power and cooling.

The main storage virtualization functions provided by Hitachi are a storage device virtualization function, volume capacity virtualization function, and storage hierarchy virtualization function. The storage device virtualization function is a way of simplifying storage management by treating multiple storage devices of different types as a single unit that operates like a single large storage device.

Storage device virtualization simplifies the adoption of a tiered storage structure made up of storage units of different types. The Universal Volume Manager function can transfer data within the storage system and this allows data to be moved between tiers with minimal impact on applications. This approach also helps make the most of existing assets because, when additional Hitachi storage units are added, the new units can be connected alongside existing older-model units.

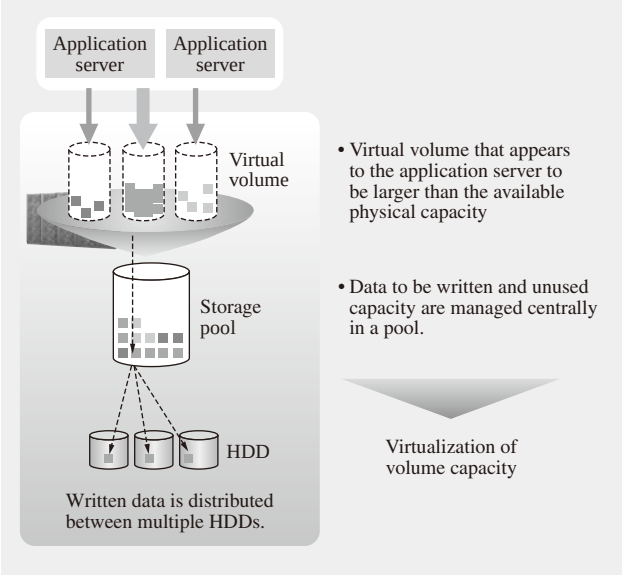

Fig. 3—Virtualization of Volume Capacity.
Multiple HDDs are managed in a single large pool and made available to the application servers as virtual volumes.

Volume capacity virtualization allows virtual volumes of any size to be assigned to the customer's application servers independently of the actual physical capacity (see Fig. 3).

Because physical capacity is allocated to the virtual volumes on an as-required basis, this improves the efficiency of storage utilization and lets users optimize their hardware investment. Also, pre-allocating a large virtual volume to a particular system reduces administration costs because it avoids the changes to server settings and frequency of reboots associated with volume expansion.

Because volume capacity virtualization automatically distributes data written from the servers across multiple volumes, it significantly reduces the workload associated with performance design and other tuning that in the past would need to be considered when configuring the system. Also, because additional disk drives only need to be added as required when the physical capacity becomes insufficient, volume capacity virtualization helps reduce the overall power consumption of the storage system.

Most data access in corporate systems tends to be concentrated on specific data areas. An example would be a system in which 50% of the data access is performed on only 5% of the data area. Accordingly, to optimize cost-performance, it is important to allocate data to an appropriate storage medium based on the frequency with which it is accessed. Storage hierarchy virtualization is a function for automatically allocating data to the optimum storage medium. The function allocates frequently accessed data areas to high-speed SSDs (solid state drives) while less frequently accessed areas are stored on slower and inexpensive storage media. Storage hierarchy virtualization reduces operation and management costs by automating the complex tasks associated with this optimization.

Green IT Initiatives

Hitachi has a project in progress called the Environmentally Conscious Data Center Project that draws on the overall resources of the Hitachi Group to make data centers more energy efficient and it is also working actively to reduce the power consumption of its storage solutions. In addition to the power saving benefits of virtualization described above, other technical innovations include saving power through the use of storage media with low power consumption and by controlling the power supplies of system resources. As examples of the latter, Hitachi supplies MAID (massive array of idle disks) systems that control the power supplies of each drive independently and power supply control mechanisms for individual expansion cabinets that can reduce power consumption by an average of 72%.

Involvement in Standardization

Using equipment from a range of different vendors, data centers in the era of cloud computing need to place even more importance on standardization. Hitachi has been actively involved in the standardization of network storage since the earliest days of the industry. Hitachi Group companies including Hitachi Data Systems (HDS) were heavily involved in the standardization and promulgation of SAN (storage area network) management interfaces through the Storage Networking Industry Association (SNIA)[(2)].

More recently, Hitachi has worked with the SNIA and US Environmental Protection Agency (EPA) on the development of a version of the ENERGY STAR certification program for storage systems with low power consumption. In Japan, Hitachi has cooperated with other industry players to establish an SNIA Japan Forum[(3)] and is taking a leading role in standardization and encouraging the adoption of standards.

EXAMPLE IMPLEMENTATION AT US SITE[(4)]

Customer Issues

The University HealthSystem Consortium (UHC) is a non-profit consortium of university health centers and university hospitals in which approximately 90% of US medical research institutions take part. The

UHC provides various services to consortium members aimed at improving health care and optimizing their administrative and financial operations. To deliver these services, the UHC operates a system that includes storage, databases and 180 servers and which is used to store and analyze large volumes of data collected from member hospitals. In early 2007, this system faced the following issues:

(1) Over the preceding six years, the size of the UHC database had on average doubled every two years. This indicated that, if the volume of data continued to grow at this pace, the organization would run out of storage capacity in 2008. The UHC concluded that it needed a storage system with sufficient scalability to cope with this future increase in data size.

(2) The previous storage system used at UHC required the administrator to perform volume performance design to balance the system load by considering factors such as the spindle positions on the physical disks. Because the volumes were independent of each other and individually assigned to specific server applications, a large amount of capacity was allocated but unused. Also, this design and allocation work took up about one week out of every quarter of the year. They required a solution that could resolve these issues by improving utilization and simplifying administration.

(3) The UHC needs to transfer data from time to time as its importance changes, but on the previous system this had a major impact on its operations because it required that the applications be shut down for 20 to 30 minutes. The UHC also required a way of transferring data that would have minimal impact on its operations.

New Storage Solution

To resolve these problems, the UHC introduced a Hitachi storage solution in January 2008. The USP V (Universal Storage Platform V) system at the core of this solution can be scaled up to a maximum of 247 Pbyte and it was combined with a mid-range AMS500 (Adaptable Modular Storage 500) system to form a three-tier configuration (see Fig. 4).

The new Hitachi virtualization solution combined the USP V and AMS500 in such a way that they could be managed as a single virtual storage pool. This improved utilization by reducing the amount of capacity that was allocated but unused, eliminating waste equivalent to approximately 26% of the storage capacity.

Pooling storage simplifies storage administration.

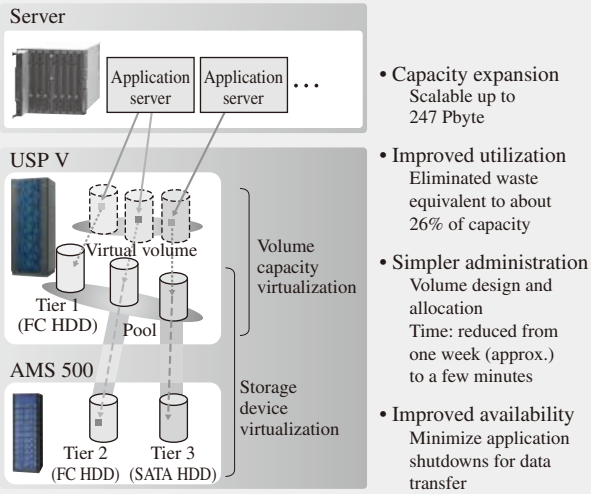

Fig. 4—Example of Hitachi Storage Solution.
Adoption of the Hitachi storage solution provided a high level of scalability, improved utilization, simpler administration, and superior availability.

For example, volume performance design and allocation tasks that previously took one week every quarter of the year were reduced to only a few minutes of routines with introduction of the new storage solution. The load balancing provided by the storage pooling and adoption of the latest hardware have also increased the database response time by between two and eight times.

The adoption of a three-tier storage configuration based on the UHC's business requirements means that data can now be transferred between tiers as required. Because the UVM function allows data to be transferred within the storage system, it has minimized the impact on the applications running on the servers as they no longer need to be shutdown to perform a data transfer.

FUTURE STORAGE SOLUTIONS

Data centers in the coming era of cloud computing will require even higher efficiency and lower operating costs. By further developing storage virtualization, Hitachi is working in the field of IT platforms to meet these needs through two technological innovations. These are automation achieved through interoperation with server virtualization technology and the hiding of the storage medium hierarchy.

Server virtualization technology has made it possible to improve the efficiency of new virtual server allocation dramatically. Even further efficiency improvements can be anticipated by having the

allocation and configuration of the required storage capacity be performed automatically when a new virtual server is created. Similarly, storage utilization can also be improved by automatically returning free space to the pool when it is made available by, for example, a server application deleting a file.

The volume of images, audio, video, and other content produced by companies and individuals will increase dramatically in the future. It is anticipated that this data will be managed centrally in data centers and distributed widely for use. Hitachi uses the term "content cloud" to refer to this type of data use. In addition to providing long-term, safe, and secure storage of large volumes of content, the content cloud will require new forms of storage capable of providing advanced services such as search, sharing, distribution, and hierarchical management. Hitachi will continue to work on the research and development of the underlying storage systems needed to meet the expectations that users from around the world have for the content cloud.

CONCLUSIONS

This article has described the origins of Hitachi's storage solutions whose virtualization functions are highly regarded around the world as well as details of the technology's latest global applications and its future outlook.

The increase in data quantity and advances in technology over the last 40 or more years have been astounding and digital information can be expected to continue growing in quantity and importance into the future. Hitachi will continue to contribute to the information society by collecting and storing long-term records that contain information about the world and its wisdom and knowledge and making it available for access through storage technology that has become part of the infrastructure of modern society.

REFERENCES

(1) News Release, "Recognition as an Information Processing Technology Heritage," http://www.hitachi.co.jp/New/cnews/month/2009/02/0223.html in Japanese.
(2) SNIA, http://www.snia.org
(3) SNIA Japan, http://www.snia-j.org in Japanese.
(4) R. Freedman, "Hitachi Breathes New Life into Critical Online Knowledge Base Applications," http://www.hds.com/assets/pdf/hitachi-breathes-new-life-critical-online-knowledge-base-applications-solution-summary.pdf

ABOUT THE AUTHORS

Toru Horimoto
Joined Hitachi, Ltd. in 1983, and now works at the Product Planning Department, Disk Array Systems Division, Information & Telecommunication Systems Company. He is currently engaged in planning for Hitachi storage solution. Mr. Horimoto is Chair of the SNIA Japan Forum (SNIA-J).

Michael C. Hay
Joined Hitachi Data Systems in 2001, and now works as the Chief Strategist within Strategic Planning, Operations and Pricing (SPOP). He is currently engaged in file storage futures and planning.