Technologies of ETERNUS VS900 Storage Virtualization Switch

● Takeshi Ito ● Atsuo Ouchi

(Manuscript received August 10, 2005)

This paper gives an outline of storage virtualization and then describes a new-generation storage virtualization system that is based on Fujitsu's ETERNUS VS900 storage virtualization switch. The VS900 provides a simplified, logical view of a complicated storage system and simplifies storage management. It adopts a network-based, out-of-band virtualization model. Among other advantages, the VS900 provides high performance, reliability, and availability. It is managed by Systemwalker Resource Coordinator, which not only manages a storage system but also manages other system resources such as servers and networks to enable unified system-wide management.

1. Introduction

Due to the relentless growth of user data, the systems that are used to store this data are becoming larger and more complicated. As a result, storage management has become a burden to system administrators, and there is an urgent need to simplify it.

Fujitsu promotes system virtualization and automation under its TRIOLE system infrastructure vision.[1] In TRIOLE, servers, networks, and storages are virtualized and autonomously managed with system management software centered on Systemwalker Resource Coordinator, thereby reducing the burden on administrators and realizing a high level of system management. The virtualization switch ETERNUS VS900 Model 200[2] (VS900 hereafter) provides a virtualization function in the storage domain.

Storage virtualization with VS900 allows flexible assignment of the storage space required to execute applications and middleware. For instance, we can place unused disk space into a large virtual disk and enlarge it whenever it approaches full capacity.

In addition, virtual disks have advanced copy functions similar to those of Fujitsu's ETERNUS RAID (Redundant Array of Inexpensive Disks) systems. By combining VS900 with Softek AdvancedCopy Manager, we can smoothly migrate to a virtualized storage environment using existing management operations without making any changes.

Administrators use Systemwalker Resource Coordinator to manage virtualized storage. It manages the VS900 and server, network, and storage environments as a unified system environment to ease the burden on administrators.

In this paper, we give an outline of storage virtualization. We then introduce the architecture of the VS900 storage virtualization switch and explain its advantages and copy function.

2. What is storage virtualization?

Storage administrators' tasks include:

1) Assigning storage to each application
2) Protecting stored data, for example, by

making backups

3) Migrating data from one storage to another when there is insufficient storage space or a newer RAID system is being installed

Storage systems continue to become more complex as the amount of stored data and storage equipment increases; however, there has been little increase in the number of storage administrators. As a result, the management workload of administrators has increased. Storage virtualization is expected to ease the burden of storage management because it provides simplified, logical views of complicated storage systems.

2.1 Types of storage virtualization

In general, virtualization is a technology that enables system resources such as CPUs, memory, storage, and networks to be used not simply as discrete physical units but also as logical aggregations of resources. Storage virtualization means that servers do not use physical disks directly, but use logical disks created from them (**Figure 1**). A physical disk can be divided into several virtual disks, or physical disks can be concatenated into a single virtual disk.

By assigning virtualized storage to servers, storage managers do not need to be conscious of physical storages. In a virtualized storage environment, physical disks are aggregated into a pool, from which virtual disks are created; therefore, we can improve our storage utilization by gathering unused storage space from each physical disk. In addition, we can reduce the space management workload by using a dynamic volume expansion function, which most storage virtualization solutions provide.

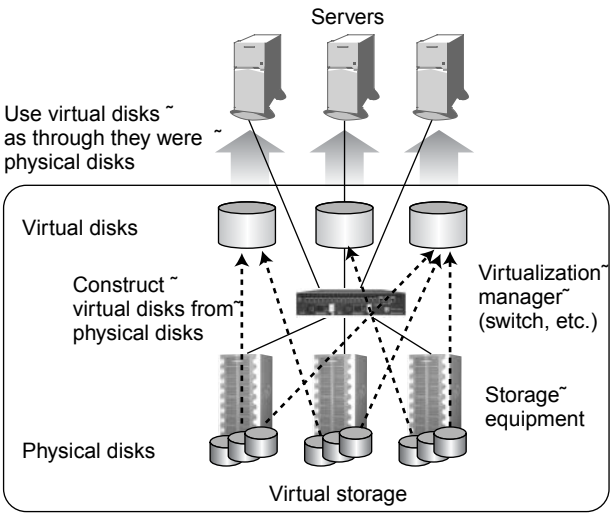

Figure 1
Storage virtualization.

Among the several types of storage virtualization systems that are used, the ones that have recently been receiving the most attention are systems such as the VS900 that are embedded inside storage network equipment. Network equipment based virtualization has the advantage that the virtualization configuration of all servers and storages is concentrated on the equipment. Therefore, system administrators can introduce the VS900 without altering the existing configuration of servers and storages and can centralize the management of virtual storages. In addition, because we do not need to install additional software on application servers, the administrators' workload during installation is minimized, especially in a heterogeneous environment containing multiple server operation systems such as Solaris and Linux.

There are other methods of storage virtualization, for example, we can use storage equipment that virtualizes the disks inside equipment and also external disks. However, these other methods require the installation of additional storage devices inside the virtualization equipment, and some of the space on these devices might never be used. In contrast, we can simply introduce the VS900, which is a new Fibre Channel switch, to virtualize our storage environment.

Network-based storage virtualization is divided into two subclasses. One is in-band virtualization, which places a virtualization mechanism on the data path, and the other is out-of-band virtualization, which places a virtualization mechanism outside the data path.

In-band virtualization has the disadvantage

that the virtualization device is apt to become a bottleneck because it fetches every communication and process virtualization. On the other hand, out-of-band virtualization is usually performed in cooperation with agents on application servers; therefore, we need to install agents on every server (**Figure 2**).

The VS900 adopts the out-of-band virtualization model. Virtualized storage is managed by a CPU inside the VS900 and by an external management server, while high-throughput virtualization operations are executed by ASICs on each port in the VS900. Moreover, unlike traditional out-of-band switches, the VS900 does not require agents on the application servers, which makes it easier to install (**Figure 3**).

3. Advantages of VS900

In addition to unified system-wide management, the VS900 provides other advantages, for example, high performance, reliability, and availability.

3.1 Virtual disk management

The VS900 has 16 Fibre Channel ports, and we define virtual logical units for each port. Servers recognize these virtual logical units as virtual disks and access them as though they were non-virtualized RAID disks.

The disks of RAIDs connected to the VS900 are aggregated into a logical pool called a virtual storage pool. We can take any amount of storage space from the pool, create a virtual disk, and provide it to servers. Also, we can create multiple storage pools and separate their usage according to the attributes of the physical disks they aggregate, for example, their access speed and RAID level (**Figure 4**). Because there are multiple storage pools with different properties, we can create virtual disks according to the storage properties required by applications.

We can expand the size of a virtual disk after it has been created. Moreover, we can assign a small virtual disk to start a new business and later increase its size as the business expands, eliminating the need to prepare a large volume in the beginning stage.

3.2 Hardware

As described in Section 3.1, each port on the VS900 has a specialized ASIC for virtualization processing. These ASICs execute READ/WRITE operations for virtual disks and provide the same high-speed throughput that can be achieved with

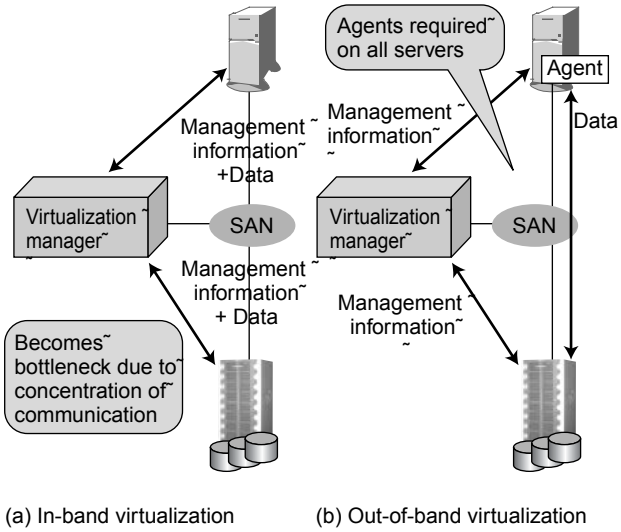

Figure 2
In-band and out-of-band virtualization.

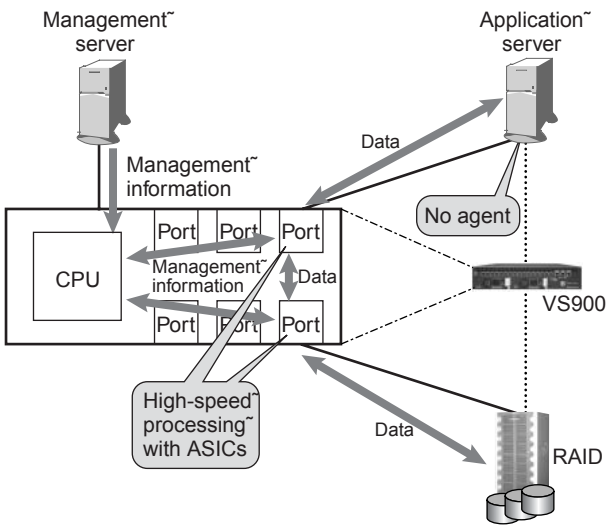

Figure 3
Virtualization with VS900.

non-virtualizing Fibre Channel storage switches. In addition, the VS900 has a CPU for copy functions and operations other than READ/WRITE. This CPU communicates with the management server, receives a virtual disk configuration, and sends it to the port ASICs. The ASICs store the configuration in the memory and execute READ/WRITE operations based on it without further direction from the CPU. This combination of CPU and ASICs provides high-performance, out-of-band virtualization.

The VS900 supports a redundant configuration of two units working in active-active configuration to prevent the VS900 from becoming a single point of failure. With this redundant configuration, a multipath driver installed in servers manages data-path redundancy and load balancing to realize high availability and throughput. In addition, the management LAN can have a dual-path configuration, making the VS900 highly available in every respect.

3.3 Software

The VS900 virtualization system uses virtualization software on the VS900 CPU and management software on the management server (**Figure 5**). The virtualization software communicates with the management software and receives information such as the virtual disk configuration and direction of the copy function. The management software, which is executed on the management server outside the VS900, receives user operation commands through a GUI and then translates them and sends them to the VS900.

By splitting the software into two parts — one of which assumes the minimum virtualization functionality and the other manages complicated functions such as the configuration — future enhancements can be introduced without degrading the performance. Also, once the virtual disks are configured, the VS900 CPU does not need to communicate with the management software, so it can continue operating even if the management server fails.

3.4 User interface

The VS900 management software works as part of Systemwalker Resource Coordinator, which is Fujitsu's total resource management software. By using Systemwalker Resource Coordinator's resource relationship management functionality, we can monitor the connectivity of servers, switches, and RAIDs and easily create

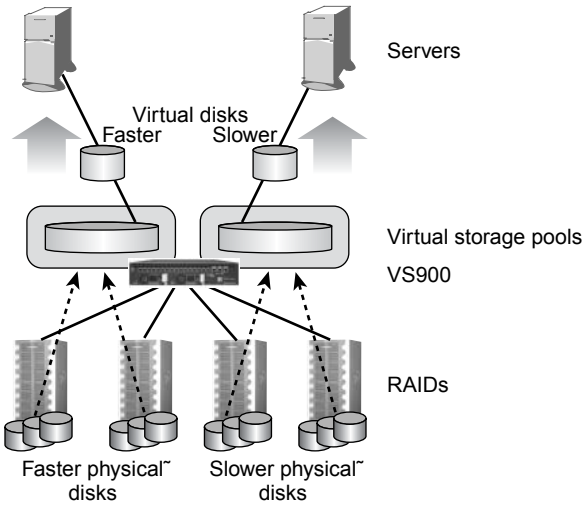

Figure 4
Creating virtual disks.

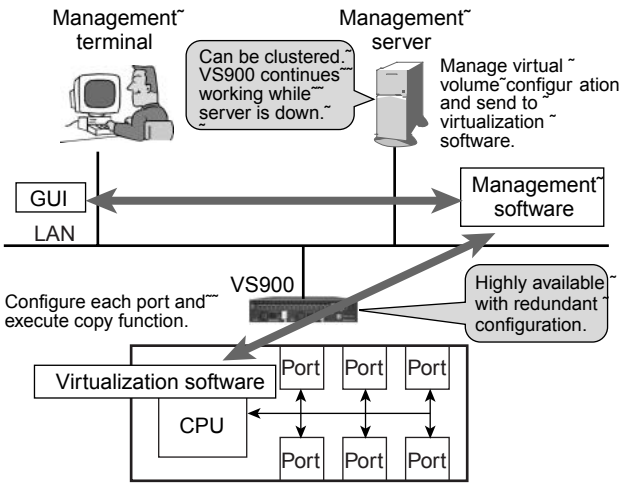

Figure 5
Software components.

virtual disks from physical disks (**Figure 6**).

In addition, Systemwalker Resource Coordinator manages not only a storage system but also other computer system resources such as servers and networks, and we can use it to monitor and control their physical and logical connections.

For instance, under a virtualized storage environment, when a RAID fails, it is difficult to determine which logical disk and hence which applications have been affected. As a result, it takes more time to investigate the situation and resume normal operation. Because Systemwalker Resource Coordinator visualizes the affected resources on a single management console, the administrator can take swift action and quickly resume normal operation.

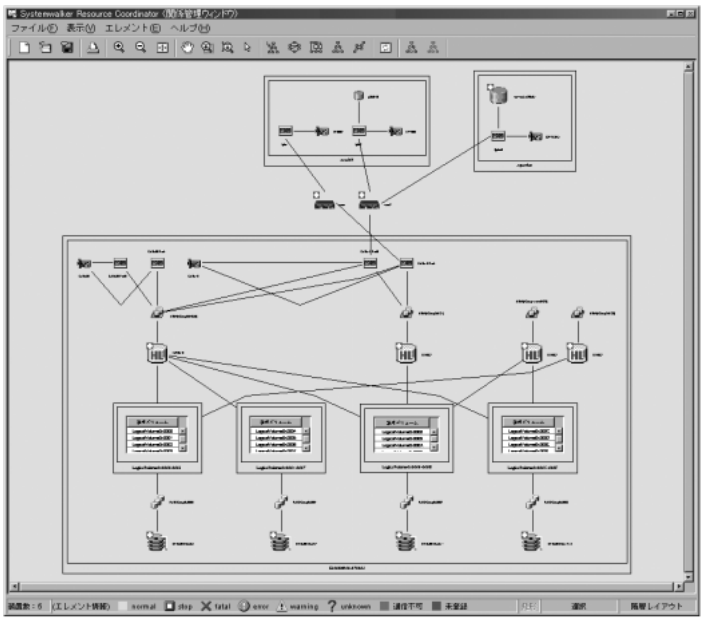

(a) Display of relationships between devices

(b) Display of virtual disk configuration˜

Figure 6
VS900 management GUI.

4. Copy function

Virtual disks are required to have the same additional functions as RAID disks, and by using those functions we can manage a virtualized storage system as well as we can manage raw RAID disks. For example, the VS900 provides the same copy functions as Fujitsu's ETERNUS series of RAID systems. The VS900 also provides migration and replication functions, which are processed by the CPU on the switch independently from application servers and therefore do not affect the application servers' workloads. Furthermore, the VS900 controls mutual exclusion between the copying and access operations of application servers, so business applications can continue during copying.

4.1 Migration

Migration means transferring data between storages. Normally, when we transfer data from one physical disk to another, we need to stop the applications that are using the disks, copy the data, alter the configuration of the server, and then resume operation. However, with the VS900, because applications can access the transferred data transparently during migration, we do not need to stop applications (**Figure 7**).

We can leverage our storage more efficiently with migration. For instance, if the frequency of access to data in a high-speed virtual disk decreases, the data can be migrated to a slower disk so space in the source disk can be made available for more frequently accessed data.

In addition, when replacing an old RAID with a new one, business applications can continue while the VS900 migrates the data to the new RAID.

4.2 Replication

Replication means creating a copy of a virtual disk. The copy is created in the background by specifying the source and destination virtual disks. After the copy is created, the source and copy are synchronized until we decide to split them.

Fujitsu includes the VS900 and its

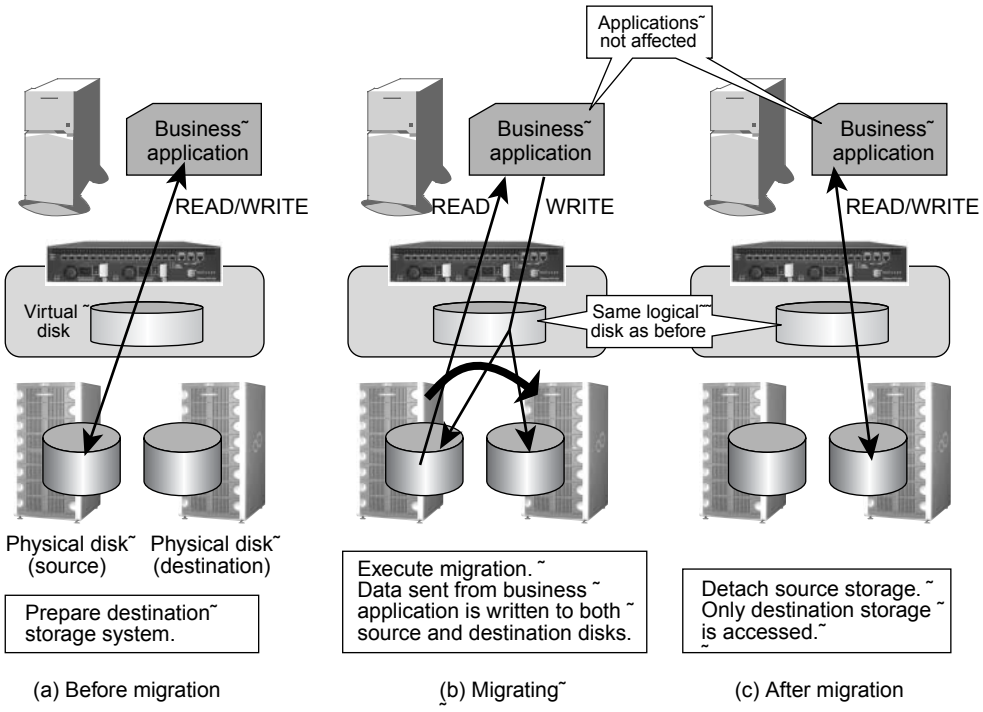

Figure 7
Migration.

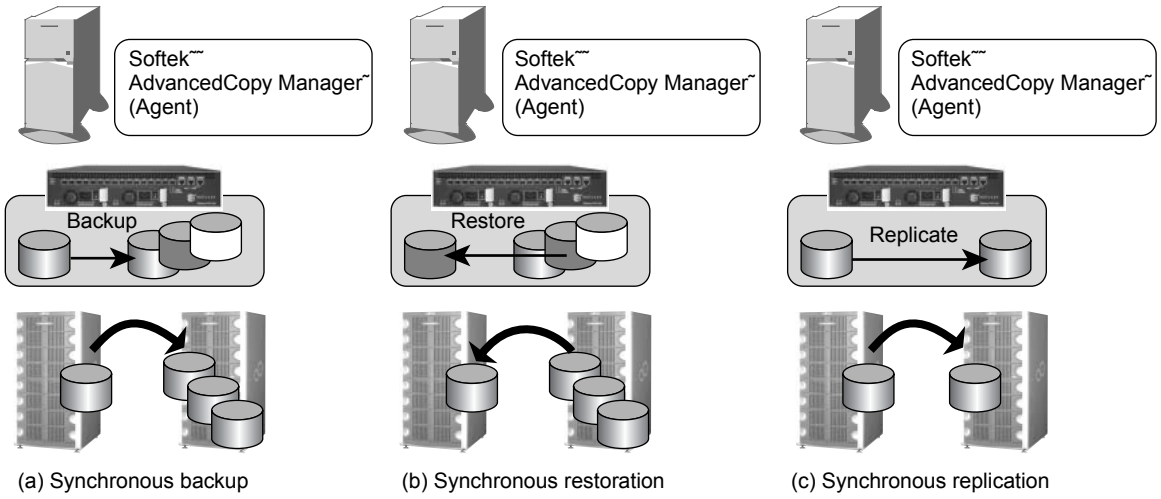

Figure 8
Replication and Softek AdvancedCopy Manager.

replication function, combined with Softek AdvancedCopy Manager, as part of its data back-up solution (**Figure 8**).

5. Conclusion

By using Fujitsu's ETERNUS VS900 storage virtualization switch for storage virtualization, we can mitigate the workload of storage system management. The VS900 is a flexible, high-reliability, high-performance storage environment that provides hardware-processed virtualization, availability-dedicated system configurations, and copy functions that are the same as those of RAID systems. Furthermore, by using Systemwalker Resource Coordinator to perform unified system management, we can flexibly assign the storage resources required for business applications while displaying all the physical and logical connections between servers, networks, and storages.

Under Fujitsu's TRIOLE vision for an autonomous and virtualized system infrastructure, system infrastructures such as servers and storages are virtualized and then flexibly assigned and managed. Our goal is to realize advanced system management and reduce the workload of system administrators through automatic management of those resources. The VS900 is one of the products we offer for realizing the TRIOLE vision through the virtualization of storage resources, and we will continue to enhance its functionality to achieve our goal.

References

1) TRIOLE; Fujitsu strategy for IT optimization. *http://www.computers.us.fujitsu.com/www/solutions.shtml?solutions/triole*
2) Virtualization switch ETERNUS VS900 model 200. *http://www.computers.us.fujitsu.com/www/news.shtml?aboutus/pressreleases/pr_100405*

Takeshi Ito received the B.E. degree in Information Science from Kyoto University, Kyoto, Japan in 1981. He joined Fujitsu Ltd., Numazu, Japan in 1981, where he has been developing data management software. He is also engaged in development of autonomic systems software.

Atsuo Ouchi received the B.E. and M.E. degrees in Information Engineering from University of Tokyo, Tokyo, Japan in 1996 and 1998, respectively. He joined Fujitsu Ltd., Numazu, Japan in 1998, where he has been developing data sharing software. He is also engaged in strategic planning of system management software.