THE POSSIBILITIES ARE INFINITE FUJITSU

ETERNUS VS900

Virtualization Switch

White Paper
June 2006

Trademarks

Microsoft, Windows, Windows NT, Windows 2000 and Windows Server 2003 are trademarks or registered trademarks of Microsoft Corporation in the United States and other countries.

Sun, Sun Microsystems, the Sun Logo, Solaris and all Solaris based marks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries, and are used under license. This document uses the abbreviation “Solaris OS” for Solaris™ Operating System.

UNIX is a registered trademark of X/Open Company, Ltd. in the United States and in other countries.

The name of systems and products mentioned in this documentation is not necessarily marked with ® or™. The other names of industrial products and companies are trademarks or registered marks.

Table of Contents

1. Fujitsu's Storage Management

Due to the relentless growth of user data, storage systems are becoming ever larger and more complicated. As a result storage management has become a burden on system administrators and is in urgent need of simplification.

Fujitsu's "TRIOLE" system infrastructure vision promotes system virtualization, automation and integration as means to achieve such business infrastructure improvement. Here servers, networks and storage are virtualized and autonomously managed using system management software centered on Fujitsu's Systemwalker Resource Coordinator. This integration of product and function not only reduces the burden on administrators but also realizes a high level of systems management. The ETERNUS VS900 Model 200 (VS900) virtualization switch provides virtualization functions for such storage domains.

Storage virtualization using VS900 enables the flexible assignment of storage space, as required, for the execution of applications and middleware. For instance, VS900 makes it possible to collect all unused disk space and build it into a large virtual disk. This ensures that storage is utilized more efficiently. When, due to increases in stored data, the virtual disk becomes 'full' it is easy to enlarge the available space as it has no "physical" limit. These virtual disks can also use the AdvancedCopy functions in similar ways to those in standard Fujitsu ETERNUS RAID systems. The combination of VS900 with ETERNUS SF AdvancedCopy Manager (previous Softek AdvancedCopy Manager) lets you continue to use existing management operations without change. This lets you smoothly migrated to a virtualized storage environment.

Administrators can also use Systemwalker Resource Coordinator to manage such virtualized storage. Systemwalker Resource Coordinator manages VS900 as part of a unified system environment of servers, network and storage to ease the burden on administrators.

2. What Is Storage Virtualization?

Typically a Storage administrators' job includes:

- assigning storage to each application
- protecting stored data via backups etc.
- migrating data from one storage device to another in response to shortages in storage space or the installation of newer RAID systems

Unfortunately storage systems continue to become more complex as the amount of stored data and the number of storage devices increases, while the number of storage administrators hardly changes. The result is greater management workloads for each administrator. The aim of storage virtualization is to provide a simplified logical view and management structure for complicated storage systems and thereby ease storage management.

2.1. Types of Storage Virtualization

In general, "virtualization" means a technology which enables the use of system resources, such as CPU, memory, storage and networks, in a form different from its purely physical makeup. Storage virtualization means that servers do not use the physical disks directly, but instead use a set of logical disks created from them. (Fig. 1) Such virtualization brings with it greater flexibility as a physical disk may be divided into several virtual disks, or several physical disks can be concatenated into a single virtual disk.

By assigning virtualized storage to servers, storage managers no longer need to be acutely aware of physical storage use. Under a virtualized storage environment, physical disks are aggregated into a pool, from which virtual disks are created. This makes it possible for you to improve your storage utilization by gathering unused storage space from the various physical disks. In addition, you can reduce your space management workload by leveraging the dynamic volume expansion function common to such storage virtualization solutions.

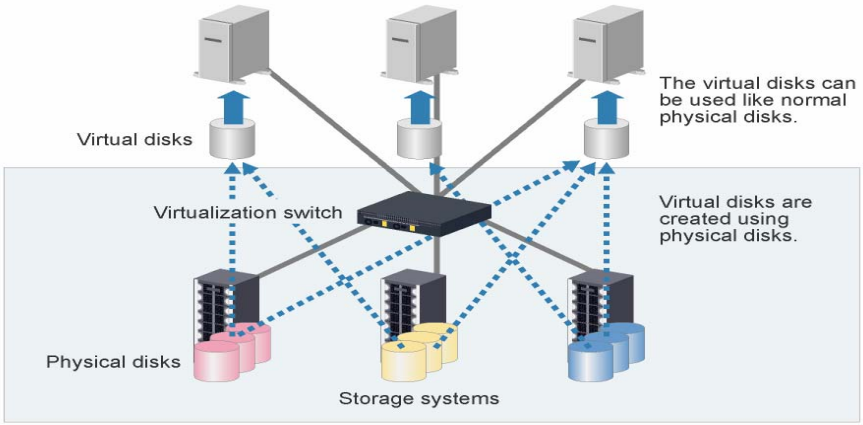

Fig.1. Storage Virtualization

Of the various types of storage virtualization, those embedded inside the storage network equipment, such as VS900, are currently attracting attention. Network equipment-based virtualization has advantages over other virtualization configurations of servers and storage, as all functions are concentrated in the equipment. This means system administrators can introduce VS900 without altering existing configurations of servers and storage to virtualize the storage and centralize the management. In addition, because there is no need to install additional software on the application servers, the administrators' installation workload is minimized. This is especially true in heterogeneous environments where multiple server operating systems, such as Solaris and Linux, are in use.

There are other ways of storage virtualization, such as the installation of virtualization-enabled storage equipment. This virtualizes both the disks inside the equipment and any external disks connected to them. However, introducing this type of equipment to virtualize existing storage regimes means that you pay for additional storage, which you might not need, inside the virtualization equipment. By contrast, you can just introduce VS900, and a new fibre-channel switch to virtualize your existing storage environment.

Typically network-based storage virtualization is divided into two sub-classes. One is in-band virtualization which has its virtualization mechanism on the data path, and the other is out-of-band virtualization which has its virtualization mechanism out of the data path.

In-band virtualization has the disadvantage that the virtualization device is apt to be a bottleneck because it fetches every communication and process virtualization. Out-of-band virtualization, on the other hand is performed usually in cooperation with agents on the application servers; therefore you need to install agents on every server. (Fig. 2)

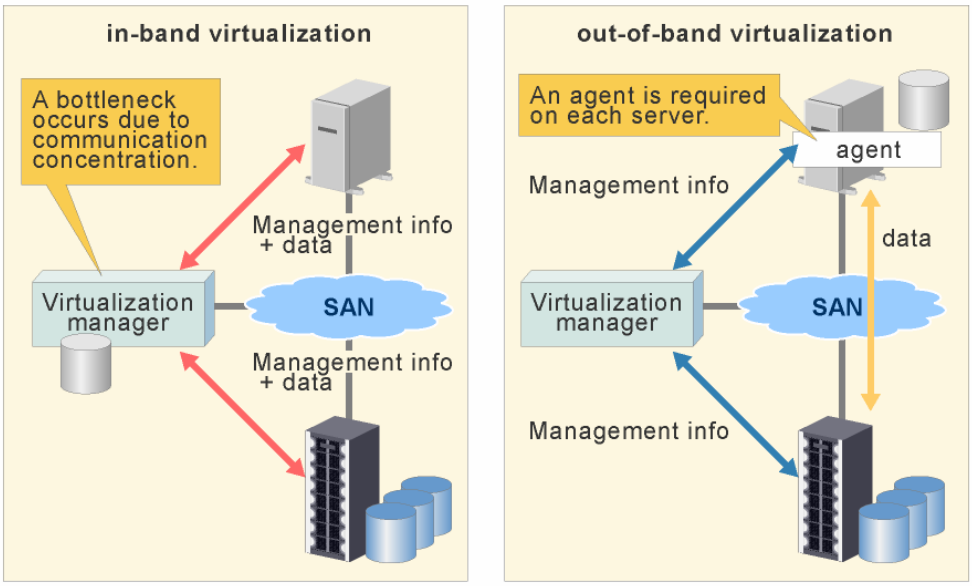

Fig.2. In-band and out-of-band virtualization

VS900 makes use of an enhanced out-of-band virtualization model. Storage virtualization is managed by the CPU inside the VS900 and an external management server. The virtualization operation is executed using Application Specific Integrated Circuits (ASICs) on every port on the VS900, to provide high throughput. Importantly, unlike a traditional out-of-band model installation of VS900 is easier as there are no agents to install on the servers. (Fig. 3)

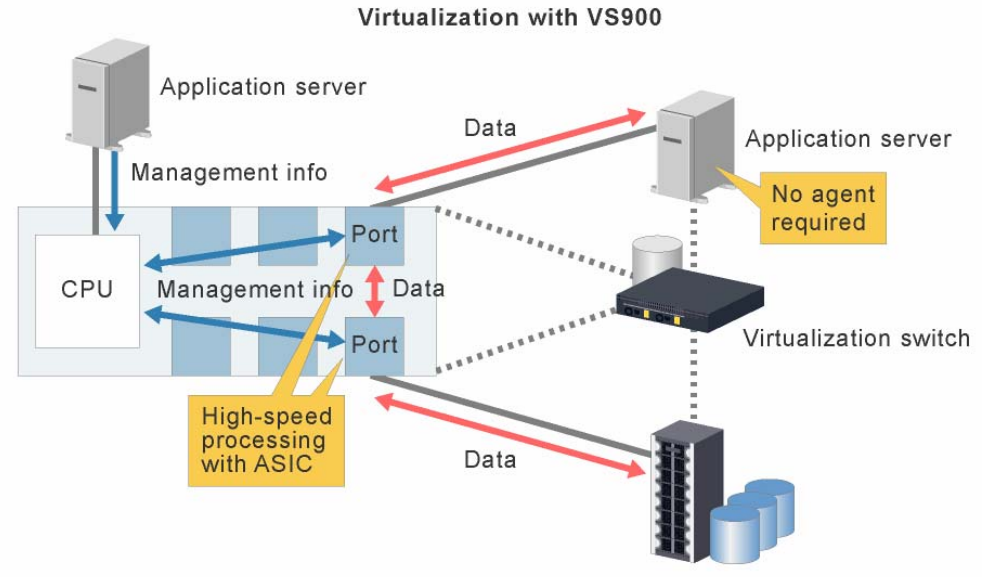

Fig.3. VS900 adopts a new type of out-of-band virtualization

3. Advantages of VS900

VS900 provides several important advantages. These include high performance, reliability and availability, in addition to unified system-wide management.

3.1. Virtual Disk Management

VS900 has 16 fiber-channel ports that let you define the virtual logical units for each port. Servers recognize these virtual logical units as virtual disks and access them.

Disks provided by RAID devices connected to VS900 are aggregated into a logical pool called the virtual storage pool. This allows you to take any amount of storage space from the pool, create a virtual disk and provide it as a disk system to servers. You can create multiple storage pools and separate their usage according to the attributes of the physical disks aggregated into them, such as access speed and RAID level. (Fig. 4) Having multiple storage pools with different properties lets you create virtual disks in accordance with the storage properties required by specific applications.

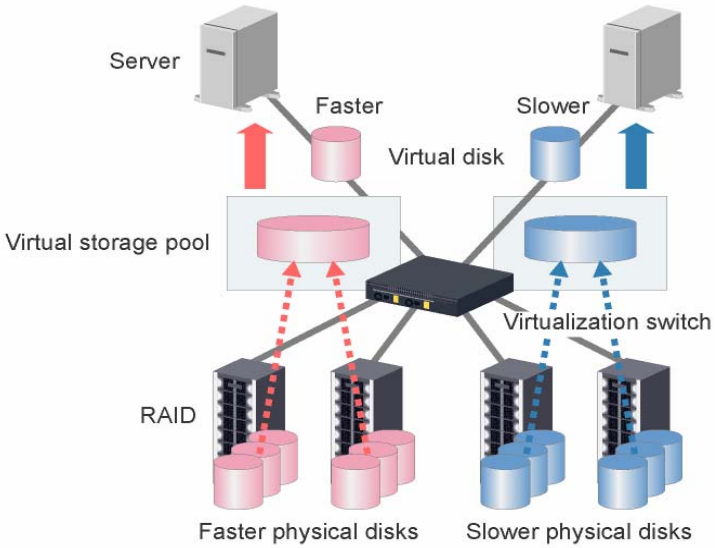

Fig.4. Creating virtual disks

Once created you can expand the size of any virtual disk. This lets you assign a small virtual disk at the start of a new business system and later expand the size according to how the business expands. This eliminates cost by removing the need to prepare a large volume from the very beginning.

3.2. Hardware

As described in 3.1, every port on VS900 has a specialized ASIC for virtualization processing. These ASICs execute READ/WRITE operations for the virtual disks and maintain high throughput. In addition, VS900 has its own CPU for copy functions and operations other than READ/WRITE. It also communicates with a management server, receives virtual disk configuration information and sends it to the port ASICs. ASICs store such information in memory and automatically execute all READ/WRITE operations based on these instructions, without further direction from the CPU. This combination of CPU and ASICs enables very high performance out-of-band virtualization.

VS900 also supports redundant configuration. That is, two VS900 units working in active-active configuration to avoid either being a single point of failure. In redundant configuration, a multipath driver installed on the servers manages redundant data paths and load balancing, further adding to the high availability and throughput. In addition, the management LAN can also be dual-pathed, making VS900 highly available in every respect.

3.3. Software

The VS900 virtualization system combines two software functions. Virtualization software on the VS900 CPU and management software on the management server. (Fig. 5) The virtualization software communicates with the management software and receives information such as virtual disk configuration and direction of copy function. Management software, which is executed on the management server external to the VS900 receives user operational instructions through a GUI and translates them into information sent to VS900.

By splitting these software functions into two parts, one of which undertakes the minimum virtualization functions and the other which manages the more complicated functions such as configuration; future enhancements can be introduced without performance degradation. Importantly, once the virtual disks are configured, the VS900 CPU does not need to communicate with the management software. It can therefore continue operating even in the event of a management server failure.

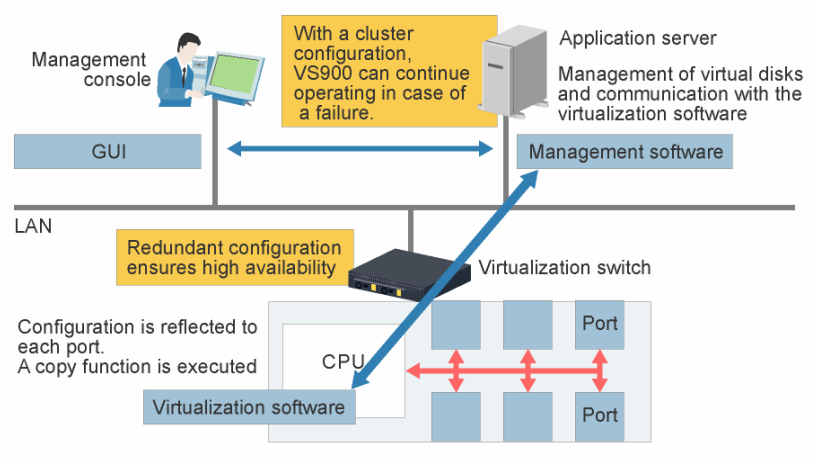

Fig.5. Software components

3.4. User Interface

VS900 management software works as a component of Systemwalker Resource Coordinator, Fujitsu's total resource management software. The combined resource relationship management functions of Systemwalker Resource Coordinator enable you to not only view the connectivity of servers, SAN and RAID systems, but also easily command the creation of virtual disks from physical disks. (Fig. 6)

The management functions of Systemwalker Resource Coordinator go beyond the storage system to include other computer system resources such as servers and networks, letting you view and control their physical and logical connections.

For instance, under a virtualized storage environment, when a RAID failure occurs it is difficult to determine which logical disk and hence which application is affected. It can therefore take a longer time to investigate and return to normal operation. Systemwalker Resource Coordinator however provides a visualization of affected resources in a single management console view, aiding the administrator to take swift action and resume operation in the shortest time.

Rleationships between servers, VS900 and physical/logical disks

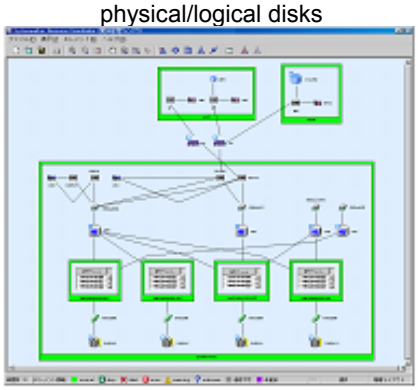

Virtual disk configuration

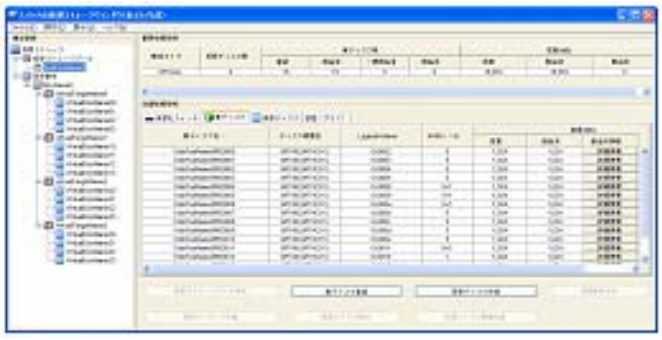

Figure 6. VS900 management GUI

4. Copy Function

For virtual disks to be fully productive they are required to have the same additional functions as RAID disks. Use of such functions allows you to manage virtualized storage systems just like physical RAID disks. VS900 provides the same copy functions as those found on Fujitsu's ETERNUS series RAID systems. Migration and replication functions are provided with VS900. Both functions are processed by the CPU on the virtualization switch, independent of the application servers. They therefore have no affect on the application servers' workloads. Furthermore, VS900 controls the mutual exclusion between copying and access from application servers, therefore business applications can continue while copying.

4.1. Migration

Migration means transferring data between storage devices. Normally, when you transfer data from one physical disk to another, you need to take several steps. These include, stopping the applications using the disk, copying the data, altering the configuration of the server and resuming. With VS900, because applications can access the transferred data transparently during migration, you do not need to stop the application. (Fig. 7)

This migration feature also enables you to leverage your storage more efficiently. For instance, when the frequency of access to data on a high speed virtual disk decreases, migrating it to a slower disk environment frees up space on the source disk making it available for more frequently accessed data.

In addition, you can use migration when replacing old RAID storage with new more efficient devices. Importantly business application processing does not need to stop while VS900 migrates data from the old RAID device to the new.

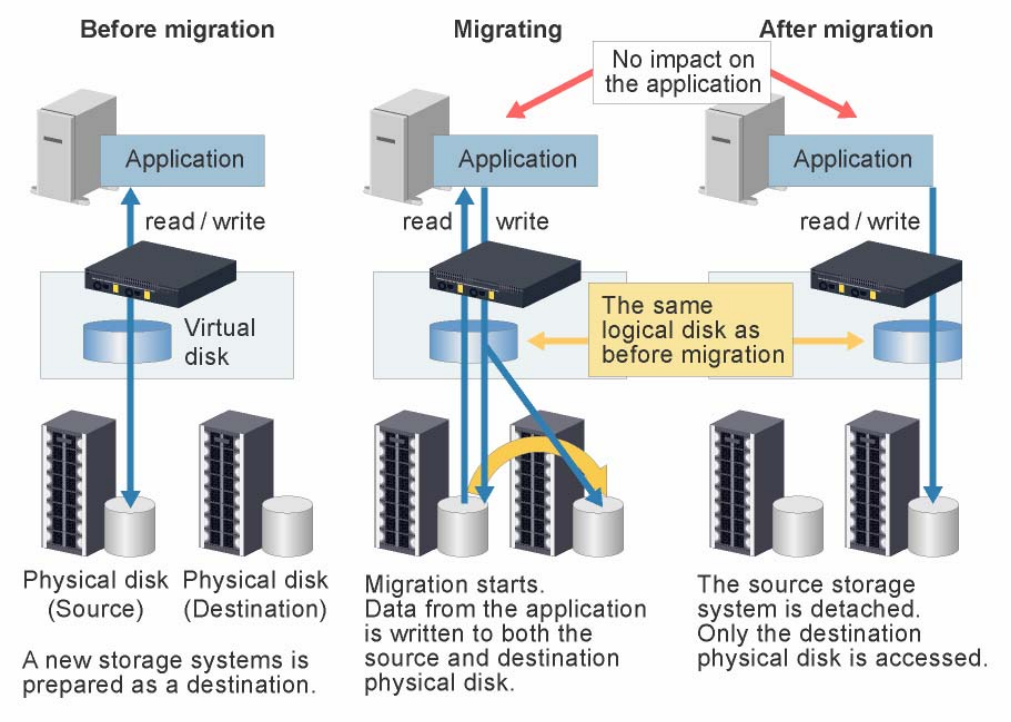

Fig.7. Migration

4.2. Replication

Replication means creating a copy of a virtual disk. The copy is created in background by specifying the source and destination virtual disks. Once the copy is created, you can automatically keep the source and copy synchronized until you command them to split.

By combining VS900 with ETERNUS SF AdvancedCopy Manager Fujitsu adds this replication function, as an additional asset of your data backup solutions. (Fig. 8)

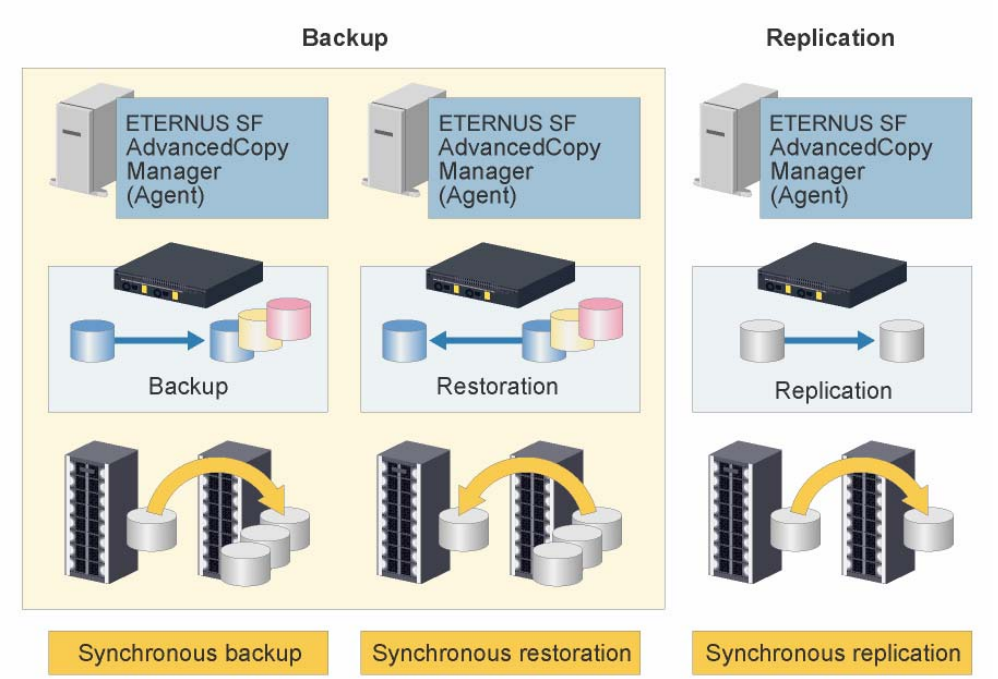

Fig.8. Replication and ETERNUS SF AdvancedCopy Manager

5. Conclusion

VS900 lets you mitigate the storage system management workload through storage virtualization. VS900 provides high performance flexible and reliable storage environments with hardware-processed virtualization, high availability focused system configuration and copy functions with the familiar functionality as of traditional RAID systems. Furthermore, Systemwalker Resource Coordinator enables unified system management, allowing flexible assignment of storage resources in line with business application requirements, while displaying all physical and logical connections between servers, networks and storage devices.

Under Fujitsu's autonomous and virtualized system infrastructure vision "TRIOLE", all system infrastructure including servers and storage are able to be virtualized flexibly assigned, and managed. Our goal is to continually advance system management to reduce the workload of system administrators through automatic management of resources. VS900, through the virtualization of storage resources, is one of the products in the realization of this "TRIOLE" vision. Fujitsu will continue to achieve its goals through ongoing development, enhancement and improvement of such functional devices.

6. Specification (As of June 2006)

6.1. Specifications

		VS900 model 200
Fibre Channel Standards		FC-PH-3 9.4, FC-SW-2 4.9, FC-FG 3.5, FC-FLA 2.7
Performance		Fibre Channel 2Gbit/sec or 1Gbit/sec (full duplex)
Total Throughput		Max. 32Gbit/sec
Architecture		Non-blocking shared memory switch
Latency		< 50μ sec
Maximum Frame Size		2112 bytes
Classes of Services		Class 3, Class F
Media Types	SWL (*1)	SFP
	LWL (*1)	Not supported
Number of Ports		16
Maximum Distance	SWL SFP (*1)	500m(1Gbit/sec), 300m(2Gbit/sec)
	LWL SFP (*1)	Not supported
Fabric Services		Simple Name Server, Zoning, WEB TOOLS
External Interfaces		Ethernet (10/100BaseT), RS232C
Size (Width × Depth × Height)		427 × 635 × 88mm
Weight		18.1kg
Power	Voltage	200 to 240VAC
	Phase	Single Phase
	Frequency	50Hz or 60Hz
Maximum Power Consumption		320W
Heat Dissipation		1152kJ/h
Hot Swappable Components		Power Supply Modules and Fans
Redundant Components		Power Supply Modules (2), Fans(2)
Environmental Conditions	Temperature	10 to 35°C
	Humidity	20 to 80%RH

*1 SWL: Short Wave Length; LWL: Long Wave Length

6.2. Specifications on Virtual Storage

		Maximum Value
Management Server	Number of manageable switches	128
	Number of storage pools	128
	Number of manageable LUNs	65,535
VS900 model 200	Number of virtual disks	2,048
	Size of a virtual disk	2TB
	Capacity of virtual storage	3.5PB *

* When connecting 14 ports to servers, and 2 ports to storage systems.

6.3. Supported Servers/Storage Systems

■ Supported Servers

Manufacturer		Product Name
Mission critical IA server	Fujitsu	PRIMEQUEST
UNIX server	Fujitsu	PRIMEPOWER
	HP	HP9000 series, Integrity*
	IBM	p Series、RS/6000 series*
IA server	Fujitsu	PRIMERGY
	Other	Other vendors' IA servers*

*We are planning to expand supports for other vendors' servers.

■ Supported Storage Systems

Manufacturer		Product Name
Disk Array	Fujitsu	ETERNUS6000 ETERNUS3000 ETERNUS GR Storage

About This White Paper

This white paper is devoted to providing technical information and an overview of the basic facilities of ETERNUS3000 storage systems. The contents of this document may be modified without any prior notice.
Please contact FUJITSU LIMITED if you find any error in descriptions.

FUJITSU LIMITED is not responsible for indemnity that might be caused by the contents in this documentation or any damage related to contents in this documentation.

FUJITSU LIMITED
http://www.fujitsu.com/storage/