

Application-Oriented Storage Resource Management

● Sawao Iwatani

(Manuscript received November 28, 2003)

Storage Area Networks (SANs) have spread rapidly, and they help customers make use of large-capacity storage resources. Many software products have been developed so that complex SANs can be managed easily and effectively. However, these products are not enough to ensure the service levels of storage resources required by applications. This is because they only provide functionality for managing storage resources from the hardware point of view (e.g., the SAN topology) and it is difficult for administrators to understand the relationships between applications on servers and storage resources when only these products are used.

This paper describes a Fujitsu storage management solution called application-oriented storage resource management that not only has the functionality of traditional SAN management but also maintains the logical relationships between storage resources and applications on servers. This management solution provides highly stable and reliable network storage environments.

1. Introduction

Because of the increasing presence of IT in our lives, there is an explosive increase in the total amount of storage capacity and much more effort is required to manage network storage. In addition, recovery from failures of storage devices such as a disk array system must be done immediately, because they may contain important application data and therefore continuous access to that data must be guaranteed.

In response, more and more alternates to traditional direct attached storage (DAS) technology are being deployed to make storage devices independent of servers and allow the consolidation of data accessed from multiple servers. One such alternative is the storage area network (SAN). DAS devices have to be managed individually on the server to which the device is attached and, as a result, there are often unused areas on some storage devices while other devices have insufficient free area. In contrast, consolidating storage

devices enables customers to reduce management work through centralized management and promotes more effective use of storage capacity. For this reason, the SAN is regarded as a driving force for optimizing an entire system from the viewpoint of storage devices.

However, as the number of storage devices and servers increase and the SAN configuration becomes more complex, it takes more skill to design a configuration that provides the required reliability and capability. This also implies that the relationships between disk drives inside storage devices and applications on a server are seldom tracked manually, because of the complex structure of the SAN. As a result, it becomes more difficult for storage administrators to pinpoint a failed part and identify which data in which operation will be affected by that failure. Therefore, it is impossible to maintain the required service levels of network storage by simply extending traditional management methods for a DAS.

The final target of storage management is to guarantee the service levels of storage resources required by applications. For example, the effect of a disk drive failure on applications should be analyzed promptly and a performance problem of storage resources must be analyzed quickly from the viewpoint of applications. It is necessary to help administrators manage their storage systems from the viewpoint of the relationships between applications and storage resources. This is the basic concept of application-oriented storage management.

Many storage vendors and software vendors provide SAN management software to reduce the complexity of SAN management. However, these products only provide functionality to manage a SAN's topology and do not help administrators understand the relationships between applications on servers and the disk drives of storage devices.

Fujitsu's Softek Storage Cruiser helps administrators understand and manage the relationships between storage resources and applications.

To achieve this, we use a new comprehensive configuration management mechanism. It has a single database that stores all physical link information and logical link information collected from storage resources. To represent the relationships between applications and storage resources, a new type of information called an "application link" is also added. This paper first describes the configuration management mechanism of Softek Storage Cruiser as the basis of application-oriented storage management. Then, some practical management functions, for example, fault management and performance management are explained.

2. Configuration management

This section describes two basic functions of application-oriented storage management. These are "storage resource correlation management" for logical and physical storage resources and "SAN

access path management" for logical connections on a SAN.

2.1 Storage resource correlation management

A SAN generally consists of storage devices, servers that access data in the devices, and FC (Fibre Channel) switches to connect them.

Servers can access any logical volume in any storage device on the network by defining the configuration in individual devices. These configuration definitions are executed and verified using commands and several kinds of management software tools.

When a network storage system is designed, the configuration definition is often made based on a design sheet. Although the configuration can be grasped to some extent by this method, a design sheet cannot easily be updated to follow frequent changes of the logical setting and physical devices. As a result, it is difficult to resolve troubles quickly, because the latest configuration is not known. This is a big problem in current network storage management.

Softek Storage Cruiser solves this problem by making the current configuration of a SAN visible and comprehensible. First, it acquires the detailed configuration and status information about storage devices, FC switches, host bus adapter (HBA) drivers, middleware, and other components related to the storage resources of the servers. Then, it correlates this information and displays it on a single screen (**Figure 1**).

The detailed configuration and status information of storage devices and FC switches are acquired via a LAN through which Softek Storage Cruiser communicates directly to firmware in the devices. To acquire information about FC HBAs and middleware products on servers, resident agents for those storage resources are placed on the servers.

The information about the logical and physical relationships between devices is included in the data acquired via LAN or by agents. Softek

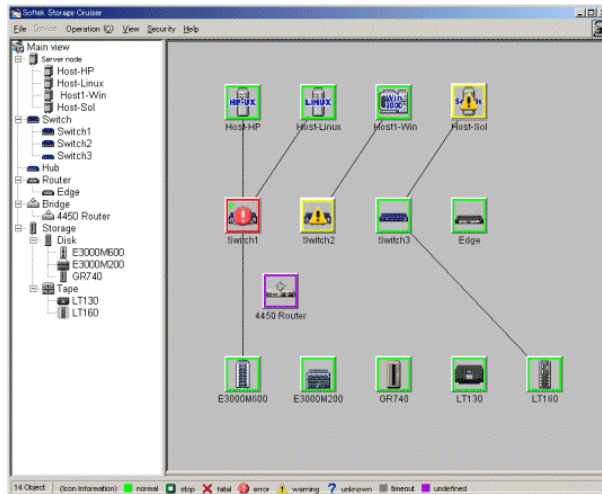


Figure 1
Main view.

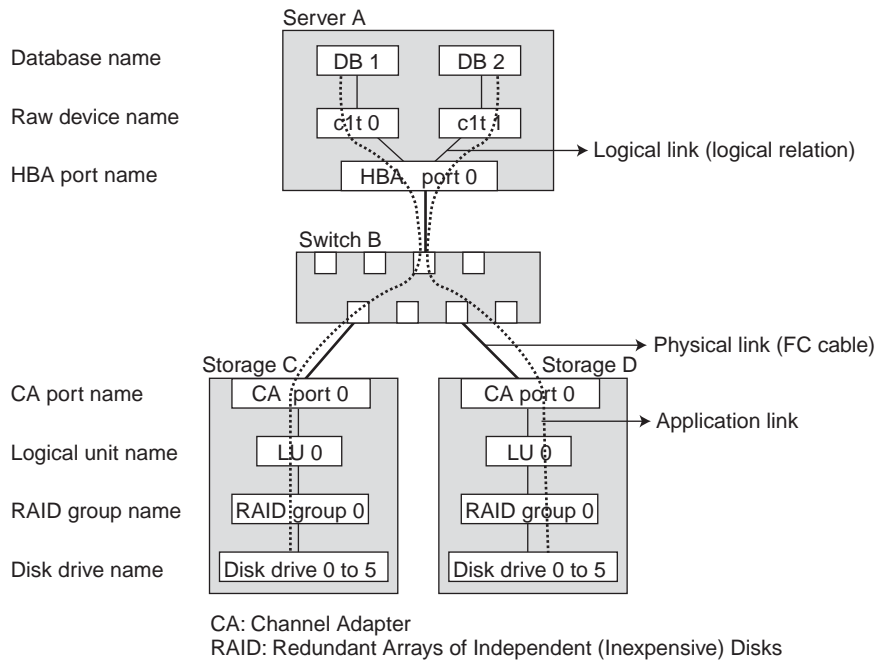


Figure 2
Logical/physical/application links.

Storage Cruiser utilizes this data to construct the overall resource relationship.

Softek Storage Cruiser has three types of link information to represent resource relationships (**Figure 2**):

- 1) Logical links that represent the logical relationships between storage resources.
- 2) Physical links that represent FC cable con-

nections. These links are created using the WWPN (World Wide Port Name) information of FC ports.

- 3) Application links that represent the relationships between applications and storage resources. This information includes data access routes from applications to disk drives inside storage devices. An application link

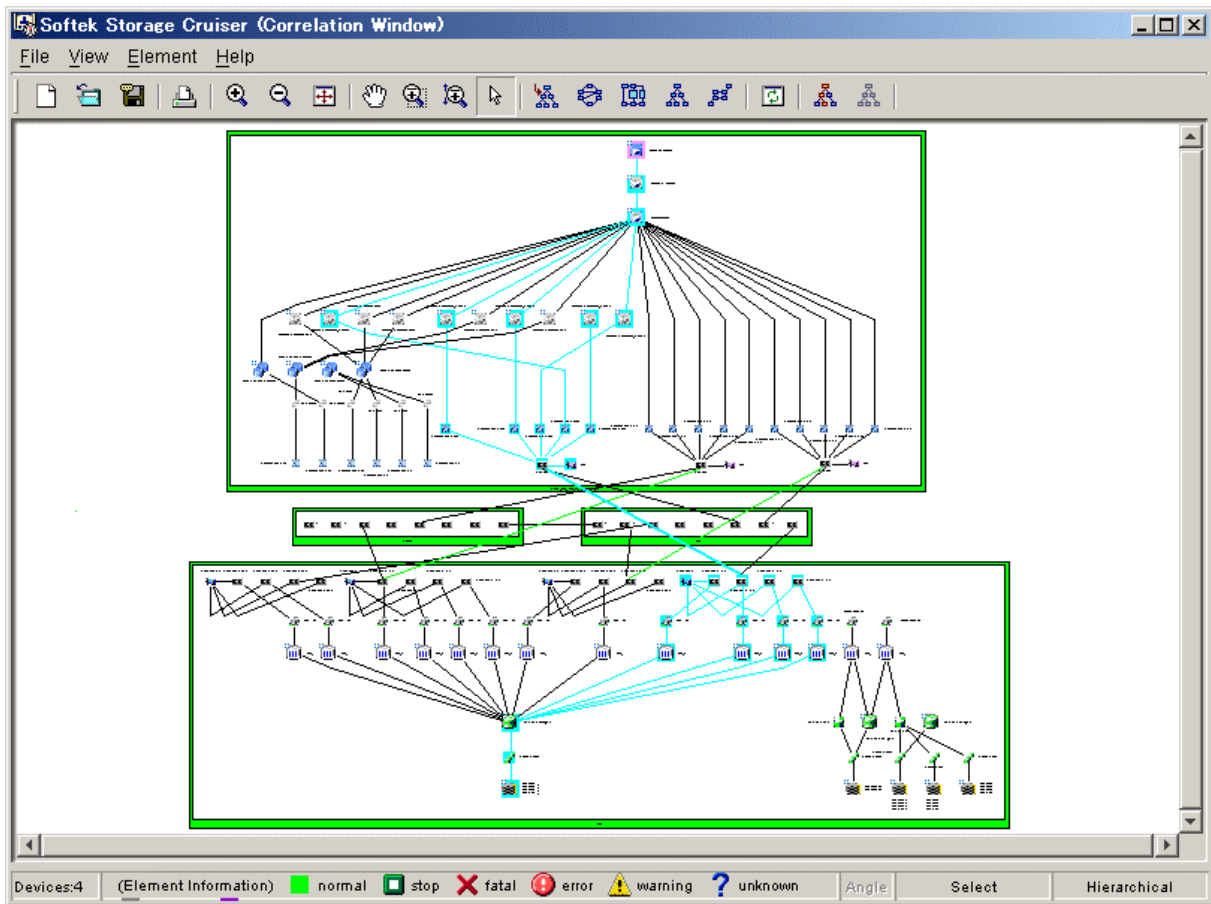


Figure 3
Correlation window.

includes a special ID to identify a data access route. For example, in Figure 2, the dotted line (an application link) shows that DB 1 accesses disk drives 0 to 5 in storage device C through raw devices c1t 0, HBA port 0, switch B, CA port 0 of storage device C, LU 0, and RAID group 0. Application links are the key information for implementing Softek Storage Cruiser and realizing application-oriented storage management.

By combining this information, all the storage resource relationships and data access routes from database applications on the server to the disk drives in storage devices are displayed on the screen. By clicking on a resource, administrators can display information about that resource and its related data access routes (**Figure 3**).

2.2 SAN access path management

A SAN is a scalable and flexible FC network through which any server can access any storage device on the network. However, unless access paths are adequately managed, important data written by a server can be corrupted by other servers. To prevent this serious problem, the access paths should be managed correctly.

An access path is logically defined on a physical fibre channel of a SAN. One way to define an access path is to use the zoning function of FC switch devices. By using this function together with the host affinity function of storages and the logical volume mapping function of servers, a SAN can be used more effectively. However, management is complicated because the definitions for an access path exist among multiple kinds of de-

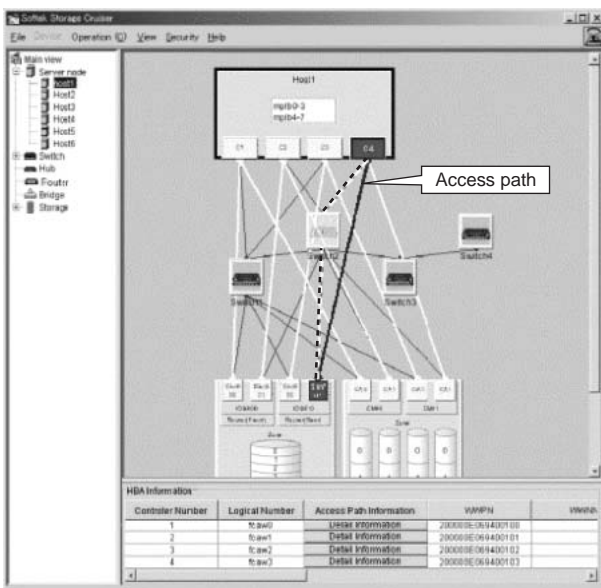


Figure 4
Access path.

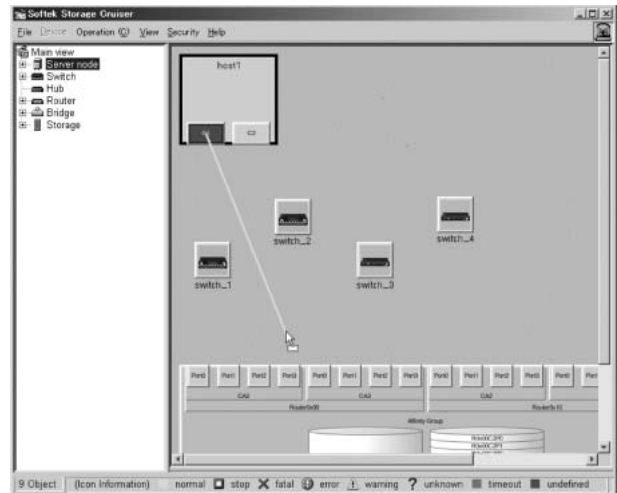


Figure 6
Access path setting.

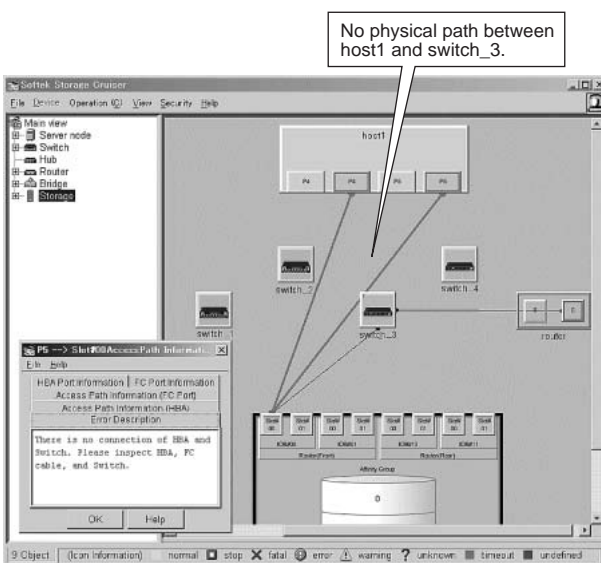


Figure 5
Access path fault.

vices and they have to be defined consistently.^{1),2)} Softek Storage Cruiser makes it possible to integrate and manage access path definition information existing in servers, FC switch devices, and storage devices by utilizing application links. Once the definitions in each device are collected, Softek Storage Cruiser recognizes and displays the access paths by integrating them

(**Figure 4**). The consistency of each piece of access path information is verified, and the details of detected inconsistencies are displayed. In addition, if a physical trouble such as a cut cable occurs, it is automatically detected and the inconsistent part is displayed (**Figure 5**).

Softek Storage Cruiser also makes it easy for administrators to define access paths. By simply connecting an FC port of a server and an FC port of a storage device on the screen of Softek Storage Cruiser, the appropriate information for each device is automatically displayed and the access path is set (**Figure 6**).

This function reduces the time needed to configure a SAN to less than a quarter of that using current tools. In addition, this function reduces the many mistakes that are normally made when setting an access path to nearly zero.

3. Storage fault management

Next, we describe an example in which Softek Storage Cruiser is used for application-oriented fault management.

When a disk drive becomes faulty, the event and the faulty status of related storage resources are displayed using special icons in the correlation window (**Figure 7**). The detailed physical

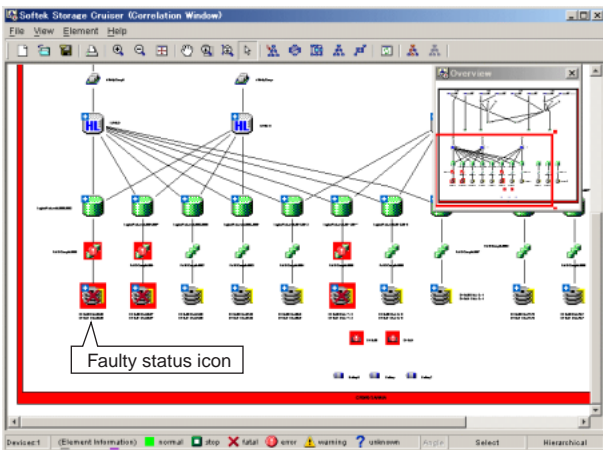


Figure 7
Correlation window – error status.

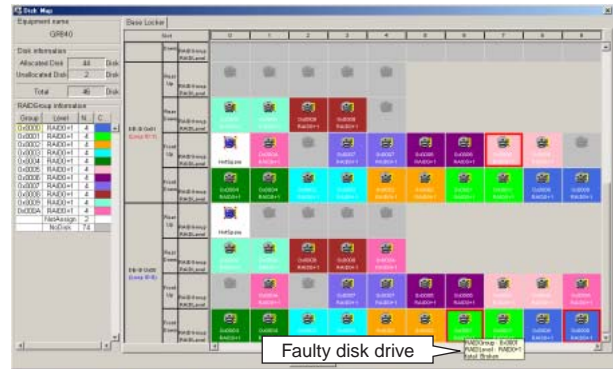


Figure 8
Disk map.

| Status | Date | Device Name | Monitor | EventID | Event log Information |
|---------|---------------------|-------------|---------|---------|---|
| Warning | 2003/01/01 10:19:50 | ETERNUS6000 | SSC | | CM OverLoad Abnormal |
| Warning | 2003/01/02 21:10:00 | LT160 | SSC | | The drive unit0 needs cleaning. |
| Info | 2003/01/03 18:00:02 | Windows2000 | SSC | | Add server unit |
| Info | 2003/01/04 11:05:39 | LT130 | SSC | | The library has detected excessive loader retries |
| Warning | 2003/01/05 15:30:40 | Solaris | SSC | | Some paths could not be detected |
| Warning | 2003/01/06 06:12:15 | Linux | SSC | | Connection timeout |
| Error | 2003/01/07 02:33:09 | GR730 | SSC | | Performance data collecting failure |
| Info | 2003/01/08 09:54:24 | HP-UX | SSC | | Add server unit |
| Error | 2003/01/09 05:36:55 | ETERNUS3000 | SSC | | Port Throughput Threshold value Exceed |
| Warning | 2003/01/10 13:07:59 | Switch | SSC | | Unit status changed: Warning |
| Info | 2003/01/11 12:28:52 | GR840 | SSC | | Performance data collecting success |

Figure 9
Event logs.

location of the faulty disk drive in the storage device is also displayed (**Figure 8**). Using this function, administrators can quickly determine the influence of a fault on applications and immediately start maintenance work if necessary. Since the related resources are displayed as described in Section 2.1, it is possible to learn which applications are related to the fault and deal with the fault appropriately. Softek Storage Cruiser's application-oriented storage resource management therefore improves the MTTR (Mean Time To Repair) of a network storage system.

Another fault management feature of Softek Storage Cruiser is that it can handle SNMP (Simple Network Management Protocol) TRAPs. When a fault is detected in a storage resource, Softek Storage Cruiser receives an SNMP TRAP from it. By using analytical algorithms, Softek Storage

Cruiser accurately analyzes the received fault as an SNMP TRAP and displays it in the event log (**Figure 9**). However, sometimes SNMP TRAPs cannot be delivered due to LAN failures. This will cause a serious problem, particularly in a mission-critical system, because administrators will not be able to detect faults that can lead to system crashes. Softek Storage Cruiser detects LAN failures by polling devices connected to the LAN and reports them as events.

The event information that Softek Storage Cruiser manages can be easily shared by other management software products. For example, Systemwalker Centric Manager, which is a Fujitsu system management infrastructure software, can receive event notifications from Softek Storage Cruiser, which makes it possible to build a centralized system management environment.

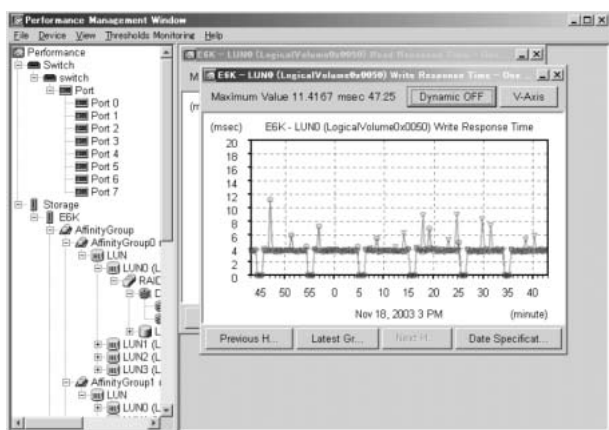


Figure 10
Performance management window.

4. Storage performance management

Softek Storage Cruiser also provides storage performance management functions based on the configuration management features described above.

Because Softek Storage Cruiser maintains information about the relationships between applications on servers and storage resources, performance bottlenecks related to storage resources and their effects on applications are investigated quickly and easily. Even if the cause of a performance problem is in an application program itself, it is easy to identify the primary cause of the problem by analyzing the performance information about related storage resources.

Softek Storage Cruiser regularly acquires performance parameters from all storage devices and summarizes them with respect to each application on the servers. It maintains information about the utilization of resources in storage devices; for example, the RAID throughput, response time, and cache hit ratio, and the utilization of FC ports (**Figure 10**).

Storage devices in a SAN environment often cause performance problems because they are accessed by multiple servers. However, these problems can be prevented by using Softek Storage Cruiser. By monitoring the performance parameters of storage devices, Softek Storage

Cruiser detects performance bottlenecks and indicates the actions that should be taken. Administrators can then respond to ensure stable application performance.

The overhead due to Softek Storage Cruiser acquiring performance parameters from storage devices is very low. Unlike server performance management, storage performance management can be done with minimal degradation of application performance.

5. Storage management in multi-vendor SAN environments

We are now developing management functionalities for Fujitsu's storage resources and also for multi-vendor storage resources. The major issue in supporting multi-vendor SAN environments is that each storage vendor has its own proprietary management interface. Furthermore, these interfaces tend to be more complex as storage devices obtain more advanced functions.

To solve this problem, the SNIA (Storage Networking Industry Association), which is the world's largest industry association for network storage, has proposed SMI-S (Storage Management Initiative Specification) to standardize storage management interfaces. Leading hardware vendors and management software vendors in the industry have already started to develop products based on SMI-S to accelerate the adoption of multi-vendor SAN environments.³⁾

SMI-S includes the following standard technologies:

- 1) Common Information Model (CIM), which is an object model for representing storage resources,
- 2) Web-Based Enterprise Management (WBEM), which enables management software products to collect information,
- 3) CIM events for fault management instead of SNMP TRAPs, and
- 4) SLP as the protocol for discovery of storage resources.

SNIA also tries to include the interfaces for

storage resource management such as backup to SMI-S.

Fujitsu is joining SNIA and plans to adopt SMI-S to achieve storage management of Softek Storage Cruiser in multi-vendor SAN environments.

6. Conclusion

In order to fully utilize a network storage system with rich functionality, it is crucial to establish a powerful management infrastructure. Especially, application-oriented storage resource management is an essential infrastructure for coping with frequent changes and requirements for quick problem-solving in complex SAN environments.

Both customers and the storage industry are

beginning to recognize the potential of network storage. Recently, functional enhancements such as information life-cycle management, virtualized storage pool management, and automatic storage provisioning have been proposed. We believe that application-oriented storage resource management is one of the basic technologies for achieving these advanced functions.

References

- 1) M. Yashiro and S. Iwatani: Fujitsu's Storage Area Network Technologies. *FUJITSU*. (in Japanese), **52**, 2, p.113-117 (2001).
- 2) I. Igarashi: Storage Management Solution in the SAN Environment. *FUJITSU*. (in Japanese), **52**, 2, p.118-122 (2001).
- 3) J. Crandall et al.: SNIA Storage Management Initiative Specification. Version 1.0.1, Storage Networking Industry Association, 12 September (2003).



Sawao Iwatani received the B.E. degree in Information Engineering from Kanazawa Institute of Technology, Ishikawa, Japan in 1991. He joined Fujitsu Limited, Kawasaki, Japan in 1991, where he has been engaged in development of storage device firmware and storage management software.

E-mail: iwa@jp.fujitsu.com