ETERNUS SF
AdvancedCopy Manager V15.1

Operation Guide

Linux

Preface

Purpose

This manual describes the functionality of ETERNUS SF AdvancedCopy Manager (hereafter referred to as "AdvancedCopy Manager").

Intended Readers

This manual is intended for a system administrator who performs storage management using AdvancedCopy Manager.

Organization

This manual is composed as follows:

Notation

The names, abbreviations, and symbols shown below are used in this manual.

Operating systems

Formal name	Abbreviation	
Microsoft(R) Windows Server(R) 2003, Standard Edition Microsoft(R) Windows Server(R) 2003, Standard x64 Edition Microsoft(R) Windows Server(R) 2003, Enterprise Edition Microsoft(R) Windows Server(R) 2003, Enterprise x64 Edition Microsoft(R) Windows Server(R) 2003 R2, Standard Edition Microsoft(R) Windows Server(R) 2003 R2, Standard x64 Edition Microsoft(R) Windows Server(R) 2003 R2, Enterprise Edition Microsoft(R) Windows Server(R) 2003 R2, Enterprise x64 Edition	Windows Server 2003	Windows
Microsoft(R) Windows Server(R) 2008 Standard (32-bit)(64-bit) Microsoft(R) Windows Server(R) 2008 Standard without Hyper-V(TM) (32-bit) (64-bit) Microsoft(R) Windows Server(R) 2008 Enterprise (32-bit)(64-bit) Microsoft(R) Windows Server(R) 2008 Enterprise without Hyper-V(TM) (32-bit) (64-bit) Microsoft(R) Windows Server(R) 2008 Datacenter (32-bit)(64-bit) Microsoft(R) Windows Server(R) 2008 Datacenter without Hyper-V(TM) (32-bit) (64-bit)	Windows Server 2008	
Microsoft(R) Windows Server(R) 2008 R2 Foundation Microsoft(R) Windows Server(R) 2008 R2 Standard Microsoft(R) Windows Server(R) 2008 R2 Enterprise Microsoft(R) Windows Server(R) 2008 R2 Datacenter	Windows Server 2008 R2	
Microsoft(R) Windows(R) XP Professional Edition Microsoft(R) Windows(R) XP Home Edition	Windows XP	
Windows Vista(R) Home Basic Windows Vista(R) Home Premium Windows Vista(R) Business Windows Vista(R) Enterprise Windows Vista(R) Ultimate	Windows Vista	
Windows(R) 7 Home Basic Windows(R) 7 Home Premium Windows(R) 7 Professional Windows(R) 7 Enterprise Windows(R) 7 Ultimate	Windows 7	
Solaris(TM) 9 Operating System	Solaris 9	Solaris or Solaris OS
Oracle Solaris 10	Solaris 10	
Oracle Solaris 11	Solaris 11	
Red Hat(R) Enterprise Linux(R) AS (v.4 for x86) Red Hat(R) Enterprise Linux(R) AS (v.4 for EM64T)	RHEL-AS4	Linux
Red Hat(R) Enterprise Linux(R) ES (v.4 for x86) Red Hat(R) Enterprise Linux(R) ES (v.4 for EM64T)	RHEL-ES4	

Formal name	Abbreviation	
Red Hat(R) Enterprise Linux(R) 5 (for x86) Red Hat(R) Enterprise Linux(R) 5 (for Intel64)	RHEL5	
Red Hat(R) Enterprise Linux(R) 6 (for x86) Red Hat(R) Enterprise Linux(R) 6 (for Intel64)	RHEL6	
SUSE(R) Linux Enterprise Server 11 (for x86) SUSE(R) Linux Enterprise Server 11 (for EM64T)	-	
HP-UX 11.0 HP-UX 11i v1 HP-UX 11i v2 HP-UX 11i v3	HP-UX	
AIX 5L(TM) V5.1 AIX 5L(TM) V5.2 AIX 5L(TM) V5.3 AIX(R) V6.1 AIX(R) V7.1	AIX	
VMware(R) Infrastructure 3 Foundation VMware(R) Infrastructure 3 Standard VMware(R) Infrastructure 3 Enterprise	VMware Infrastructure 3	VMware
VMware(R) vSphere 4 Essentials Kit VMware(R) vSphere 4 Essentials Plus Kit VMware(R) vSphere 4 Standard Edition(TM) VMware(R) vSphere 4 Standard Plus Data Recovery VMware(R) vSphere 4 Advanced Edition(TM) VMware(R) vSphere 4 Enterprise Edition(TM) VMware(R) vSphere 4 Enterprise Plus Edition(TM)	VMware vSphere 4	
VMware(R) vSphere 5 Essentials Kit VMware(R) vSphere 5 Essentials Plus Kit VMware(R) vSphere 5 Standard Edition(TM) VMware(R) vSphere 5 Standard Plus Data Recovery VMware(R) vSphere 5 Enterprise Edition(TM) VMware(R) vSphere 5 Enterprise Plus Edition(TM)	VMware vSphere 5	

Oracle Solaris might be described as Solaris, Solaris Operating System, or Solaris OS.

Related products with Fujitsu Storage System ETERNUS and Storage Management Software ETERNUS SF

Formal name	Abbreviation			
ETERNUS DX60/DX60 S2 ETERNUS DX80/DX80 S2 ETERNUS DX90/DX90 S2	-		ETERNUS DX series	ETERNUS Disk storage system
ETERNUS DX410 ETERNUS DX440	ETERNUS DX400 series	ETERNUS DX400/DX400 S2 series		
ETERNUS DX410 S2 ETERNUS DX440 S2	ETERNUS DX400 S2 series			
ETERNUS DX8100 ETERNUS DX8400 ETERNUS DX8700	ETERNUS DX8000 series	ETERNUS DX8000/DX8000 S2 series		
ETERNUS DX8100 S2 ETERNUS DX8700 S2	ETERNUS DX8000 S2 series			

Formal name	Abbreviation	
ETERNUS2000 ETERNUS4000 ETERNUS8000	-	
Web GUI of ETERNUS DX series	ETERNUS Web GUI	
ETERNUSmgr		

Software products

Formal name	Abbreviation
Microsoft(R) Internet Explorer(R)	Internet Explorer
Mozilla(R) Firefox(R)	Firefox
Microsoft(R) Cluster Service	MSCS
Microsoft(R) Windows Server(R) Failover Clustering	WSFC
Microsoft(R) Exchange Server	Exchange Server
Microsoft(R) SQL Server(TM)	SQL Server
PRIMECLUSTER Global Disk Services	GDS
PRIMECLUSTER Global File Services	GFS
Symfoware Server Enterprise Extended Edition Symfoware Server Enterprise Edition	Symfoware
VMware(R) ESX(R)	VMware ESX
VMware(R) ESXi(TM)	VMware ESXi
VMware(R) vCenter(TM) Server	VMware vCenter Server or vCenter Server

Manuals

Formal name	Abbreviation
ETERNUS SF Express / ETERNUS SF Storage Cruiser / ETERNUS SF AdvancedCopy Manager Installation and Setup Guide	ETERNUS SF Installation and Setup Guide
ETERNUS SF Express / ETERNUS SF Storage Cruiser / ETERNUS SF AdvancedCopy Manager Migration Guide	ETERNUS SF Migration Guide
ETERNUS SF Express / ETERNUS SF Storage Cruiser / ETERNUS SF AdvancedCopy Manager Web Console Guide	ETERNUS SF Web Console Guide
ETERNUS SF Express / ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module	ETERNUS SF Operation Guide for Copy Control Module
ETERNUS SF Storage Cruiser / ETERNUS SF AdvancedCopy Manager Cluster Environment Setup Guide	ETERNUS SF Cluster Environment Setup Guide
ETERNUS SF Express / ETERNUS SF Storage Cruiser / ETERNUS SF AdvancedCopy Manager Messages	ETERNUS SF Messages
ETERNUS SF Express / ETERNUS SF Storage Cruiser Event Guide	ETERNUS SF Event Guide
ETERNUS SF Express / ETERNUS SF Storage Cruiser / ETERNUS SF AdvancedCopy Manager Glossary	ETERNUS SF Glossary

Export Administration Regulation Declaration

This document may contain specific technologies that are covered by International Exchange and International Trade Control Laws. In the event that such technology is contained, when exporting or providing the document to non-residents, authorization based on the relevant law is required.

Trademarks

- Microsoft, Windows, Windows Server, Windows Vista and Internet Explorer are registered trademarks of Microsoft Corporation in the United States and other countries.

- UNIX is a registered trademark exclusively licensed for X/Open Company Limited.

- Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

- Linux is a trademark or registered trademark of Mr. Linus Torvalds in the United States and other countries.

- Red Hat, Red Hat Shadow Man logo, RPM, RPM logo and Glint are registered trademarks of Red Hat, Inc.

- SUSE is a registered trademark of Novell, Inc, in the United States and other countries.

- HP-UX is a registered trademark of Hewlett-Packard Company in the United States and other countries.

- AIX and AIX 5L are trademarks or registered trademarks of International Business Machines Corporation in the United States and other countries.

- VMware, VMware logo, Virtual SMP and VMotion are the registered trademarks or trademarks of VMware, Inc in the US and in other countries.

- All other trademarks and product names are the property of their respective owners.

Shipment date and revision history

Shipment date	Revision	Manual code
January 2012	1	B1FW-5971-01ENZ0(00) / B1FW-5971-01ENZ2(00)
January 2012	1.1	B1FW-5971-01ENZ0(01) / B1FW-5971-01ENZ2(01)
February 2012	1.2	B1FW-5971-01ENZ0(02) / B1FW-5971-01ENZ2(02)
March 2012	1.3	B1FW-5971-01ENZ0(03) / B1FW-5971-01ENZ2(03)
April 2012	2	B1FW-5971-02ENZ0(00) / B1FW-5971-02ENZ2(00)
May 2012	2.1	B1FW-5971-02ENZ0(01) / B1FW-5971-02ENZ2(01)
August 2012	3	B1FW-5971-03ENZ0(00) / B1FW-5971-03ENZ2(00)

Notes

- No part of this manual may be reproduced without permission.

- This manual is subject to change without advance notice.

Copyright

Update history

Content of update	Updated section	Revision
Added a command.	8.3	2.1
Added the operation on Model upgrade for ETERNUS Disk storage system.	7.5.8	

Content of update	Updated section	Revision
Modified the explanation of pre-processing and post-processing.	A.2.1, A.2.2, A.3.1, A.3.2, C.2.2, C.2.3	
Added explains operation using the KVM Support Function.	Appendix F	
Added Oracle Solaris 11, AIX V6.1 and AIX V7.1 in the operating systems.	Notation in Preface	3
Added the explanation for backup operations for Symfoware.	Chapter 4	
Added the explanation for the maintenance for database of ETERNUS SF Manager.	8.7	
Added the notes on combining cascade copy and Concurrent OPC operations.	10.1.2	

Manual organization and reading suggestions

Manual organization

The following table describes the Manual organization of Express, Storage Cruiser and AdvancedCopy Manager.

When to read	Related manuals (abbreviated)	Related products (**NOTE**)			Explanation
		EXP	SC	ACM	
Before installation	Quick Reference	Yes	Yes	Yes	This manual is unique for each product. The following manuals are available: - Express Quick Reference - Storage Cruiser Quick Reference - AdvancedCopy Manager Quick Reference
	Overview	No	No	Yes	This manual is unique for each product.
During installation	Installation and Setup Guide	Yes			This manual is common for all products.
	Cluster Environment Setup Guide	No	Yes		This manual is common for Storage Cruiser and AdvancedCopy Manager.
	Migration Guide	Yes			This manual is common for all products.
During operation	Operation Guide	Yes	Yes	Yes	This manual is unique for each product. The following manuals are available: - Express Operation Guide - Storage Cruiser Operation Guide - Storage Cruiser Operation Guide for Optimization Option - AdvancedCopy Manager Operation Guide (for Windows) - AdvancedCopy Manager Operation Guide (for Solaris) - AdvancedCopy Manager Operation Guide (for Linux) - AdvancedCopy Manager Operation Guide (for HP-UX) - AdvancedCopy Manager Operation Guide (for AIX)

When to read	Related manuals (abbreviated)	Related products (**NOTE**)			Explanation
		EXP	SC	ACM	
	Operation Guide for Copy Control Module	Yes	No	Yes	This manual is common for Express and AdvancedCopy Manager.
	Web Console Guide	Yes			This manual is common for all products.
Anytime	Event Guide	Yes		No	This manual is common for Express and Storage Cruiser.
	Messages	Yes			This manual is common for all products.
	Glossary	Yes			This manual is common for all products.

NOTE: "EXP" indicates Express, "SC" indicates Storage Cruiser and "ACM" indicates AdvancedCopy Manager.

How to read manuals

Please use the following table to find the most useful information in the Express, Storage Cruiser and AdvancedCopy Manager manuals to answer your inquiry.

Purpose	Related products (**NOTE**)	Manual	Main contents	How to read
Acquiring a product overview and basic operation knowledge	EXP	- Express Quick Reference	- Product overview - Installation decision - Overview of the necessary tasks from installation to first use	Please read if you want to acquire a fundamental knowledge of the product and its operation in order to decide to install it or not.
	SC	- Storage Cruiser Quick Reference		
	ACM	- AdvancedCopy Manager Quick Reference		
		- AdvancedCopy Manager Overview	- Main functions - Linkable applications - Procedure overview for Advanced Copy of ETERNUS Disk storage system	
Deciding if a version upgrade is required	common	- ETERNUS SF Migration Guide	- Incompatibilities with previous version - Notes and cautions about version upgrade - Version upgrade procedure	Please read if you want to upgrade from a previous version.
Installing and correctly operating the product Setting up operating environment depending on purpose	common	- ETERNUS SF Installation and Setup Guide	- Operating environment - Installation procedure - Setup procedure - Uninstallation procedure	Please read if you want to install and setup the product.
	SC, ACM	- ETERNUS SF Cluster Environment Setup Guide	- Supported cluster software - Installation procedure for a clustered system - Setup procedure for a clustered system	Please read if you want to install and setup the product on a clustered system.

Purpose	Related products (**NOTE**)	Manual	Main contents	How to read
			- Uninstallation procedure for a clustered system	
Administration and operation of the installed system	EXP	- Express Operation Guide	- Starting and stopping the software - Device monitoring - Data copy inside the storage system - Necessary tasks after an architectural modification of the system as well as product maintenance	Please read if you want to start or shutdown the system, monitor the operation status, do backup/restore operations, etc.
	SC	- Storage Cruiser Operation Guide	- Starting and stopping the software - Device monitoring - Necessary tasks after an architectural modification of the system as well as product maintenance - Command reference	
		- Storage Cruiser Operation Guide for Optimization Option	- Operating environment construction - Operating status monitoring - Necessary tasks after an architectural modification of the system as well as product maintenance - Command reference	
	EXP, ACM	- ETERNUS SF Operation Guide for Copy Control Module	- Starting and stopping the software	
	ACM	- AdvancedCopy Manager Operation Guide (for Windows) - AdvancedCopy Manager Operation Guide (for Solaris) - AdvancedCopy Manager Operation Guide (for Linux) - AdvancedCopy Manager Operation Guide (for HP-UX) - AdvancedCopy Manager Operation Guide (for AIX)	- Data backup/restore inside the storage system - Necessary tasks after an architectural modification of the system as well as product maintenance - Command reference	
	common	- ETERNUS SF Web Console Guide	- Operating environment - Screen layout description	Please read if you want to understand the ETERNUS SF Web Console.

Purpose	Related products (**NOTE**)	Manual	Main contents	How to read
Dealing with messages issued by the software	common	- ETERNUS SF Messages	- Messages and their explanations - Parameter (variable information) description - System action - Countermeasures	Please read if you want a practical way of investigating and dealing with messages issued by the software.
Dealing with events issued by the software	EXP, SC	- ETERNUS SF Event Guide	- Phenomenon of event - Countermeasures	Please read if you need to find a practical way of investigating and dealing with events.
Researching the meaning of specific terms related to the products and other important terms	common	- ETERNUS SF Glossary	- Product specific terminology explanation - Explanation of important terminology appearing in the manual - Synonyms and related terms - Proper form of abbreviated terms	Please read if you want to learn the meaning of important terms, product specific terms or abbreviations used in the manuals.

NOTE: "EXP" indicates Express, "SC" indicates Storage Cruiser and "ACM" indicates AdvancedCopy Manager.

Contents

Chapter 1 Overview of AdvancedCopy Manager

This chapter describes the overview of AdvancedCopy Manager.

1.1 What is AdvancedCopy Manager?

AdvancedCopy Manager is a product that provides data storage backup and recovery management in an open system and incorporates the following features.

- 1.1.1 High-speed backup (24-hour basis)
- 1.1.2 Non-stop database backup
- 1.1.3 High-speed replication

1.1.1 High-speed backup (24-hour basis)

The Advanced Copy function of ETERNUS Disk storage system is used to back up data quickly. The Advanced Copy function provides fast data copy regardless of data volume, and can maintain a fixed backup time even if the amount of data to be backed up increases. This dramatically reduces the duration for which the application must be stopped to perform backup.

In addition, since operations management is linked to the following applications, high-speed backup can be performed without interruption to live operational processing:

- Oracle
- Symfoware

If other databases or a general file system are being used, downtime or interruption to normal throughput is still significantly reduced.

1.1.2 Non-stop database backup

By using AdvancedCopy Manager, some databases can be backed up without stopping the application. Please refer to the following for the backup method for Oracle or Symfoware databases:

- For Oracle: "B.3 Sample Scripts for Oracle"
- For Symfoware: "Chapter 4 Backup Operations (Symfoware)"

AdvancedCopy Manager can also back up independently of PRIMECLUSTER GDS, including the mirror volume (SDX object) environment.

1.1.3 High-speed replication

The Advanced Copy function of ETERNUS Disk storage system enables data to be used for purposes other than the original purpose and replicated data (copy) to be collected as a protection against data corruption.
The Remote Advanced Copy function of ETERNUS Disk storage system enables copy operations to an ETERNUS Disk storage system at remote locations.

When using a script for linkage with below applications, high-speed copying (backup) of user data can be performed without interruption to live operational processing:

- Oracle

1.2 Product Configuration of AdvancedCopy Manager

1.2.1 Function system

AdvancedCopy Manager consists of the following three functions:

Table 1.1 Functions

Function name	Description
AdvancedCopy Manager's manager	Operates on the Management Server, and manages two or more Managed Servers. Backup management and replication management are performed.
AdvancedCopy Manager's agent	Operates on the Managed Server with which the storage device is connected. Backup management and replication management are performed.
	This (Agent) can be managed by AdvancedCopy Manager's manager running on Solaris, Windows or Linux operating system.
ETERNUS SF Web Console	Monitors and/or manages the AdvancedCopy Manager environment by connecting to the Management Server.

 Note
....

The Manager program has also the Agent function.
....

The AdvancedCopy Manager's server type depends on the operating environment.

Managed Server

> The Managed Server is a server which executes the backup operation and the replication operation that AdvancedCopy Manager offers. The Agent component of AdvancedCopy Manager is installed on the Managed Server. If a Managed Server is also to be used as the Management Server then the Manager, as well as the Agent, will need to be installed. The Agent installation is included as part of the Manager installation.

Management Server

> The Management Server manages the Managed Servers and only one Management Server can be set in the system.
> The Management module of AdvancedCopy Manager is installed on the Management Server.

Three different operating scenarios for the AdvancedCopy Manager are shown below:

- When only one Managed Server is installed

  In this operating environment, only the Management module of AdvancedCopy Manager needs to be installed (Management module of AdvancedCopy Manager includes the Agent functions of AdvancedCopy Manager).

Figure 1.1 When only one Managed Server is installed

- When two or more Managed Servers are installed

  When two or more Managed Servers are available, set only one of them up as the Management Server or set up a new server with no ETERNUS Disk storage system connected as the Management Server. Each environment is explained below.

  a. When a Managed Server is operated as a Management Server:

     One of the servers functions as a Management Server. The Management Server has a repository where the information on operating and management of all Managed Servers is stored.

     When many Managed Servers exist, they all compete for access to the repository. To ensure overall performance and other applications running on the same server are not adversely affected, the Management Server functionality should be installed on a dedicated server, or on a server on which the load of other applications is relatively low.

     The following diagram illustrates the configuration of the Managed Server when it is also operating as the Management Server.

Figure 1.2 When a Managed Server is also operated as a Management Server

b. Setting up a dedicated Management Server:
This setup comprises two or more Managed Servers with a dedicated Management Server.

Figure 1.3 When a dedicated Management Server manages Managed Servers

1.2.2 System configuration

The system configuration of AdvancedCopy Manager is shown below.

Figure 1.4 System configuration of Management Server and two or more Managed Servers

1.2.3 Combining the Agent and Manager

The following product level combinations are possible when using different levels of Agent and Manager.

Table 1.2 Combinations the Agent and Manager

Version level of Manager (Management Server)	Version level of Agent (Managed Server)		
	13.x	14.x	15.x
13.x	A	C	C
14.x	B	A	C
15.x	B	B	A

A: Connection possible

B: Connection possible (however, functions supported by the Managed Server version or later cannot be used)

C: Connection not possible

Note

- There is a difference in functionality between some product revision levels. Installation of the manager and agent of the same revision level is recommended.

- The following restrictions apply to combinations of Manager and Agent products at different revision levels:

  - If the Manager is a newer level than the Agent

    Connection possible (but functions supported by levels later than that of the Managed Server cannot be used)

  - If the Agent is a newer level than the Manager

    Connection not possible

1.3 Function Overview of AdvancedCopy Manager

1.3.1 What is the backup function?

The backup function of AdvancedCopy Manager provides operation management for high-speed backup by using the Advanced Copy function of ETERNUS Disk storage system.

The following sub-sections provide more information concerning the backup function.

High-speed backup operation

The Advanced Copy function of ETERNUS Disk storage system is used to perform high-speed backup between volumes regardless of volume capacities. This function greatly reduces the impact of a backup operation.

Faster backups can be made by performing update differential copies using the Advanced Copy function.

Backup operation in conjunction with Symfoware

If, in your transactions, you are using databases that were configured with Symfoware, you can implement fast backups without stopping your transactions.

Backup operation in conjunction with Oracle (sample scripts)

The following sample scripts provide AdvancedCopy Manager commands and Oracle operations combined together:

- Acquisition of DBMS (Oracle) resource information

- Synchronization with DBMS (Oracle) and volume duplication

Additionally, the following functions allow you to manage the backup operation.

Policy management function

The policy management function manages the number of preservation generations of backup and the number of interval days in which backup should be performed (displays the delay when displaying information for a volume to be backed up).

Storage configuration management function

The storage configuration management function manages the configurations of volumes to be backed up and volumes to which a backup should be saved.

On AdvancedCopy Manager, a volume to be backed up is called a transaction volume; a volume to which a backup should be saved is called a backup volume.

Backup and restore functions

The backup and restore functions back up and restore a specified volume.

History information management function

The history information management function manages, out of the information that has been backed up in the past, the information in preservation generations as history.

Consistency recovery function

The consistency recovery function resolves inconsistencies in the AdvancedCopy Manager backup management information caused by a system or other type of failure.

1.3.2 What is the replication function?

The replication function of AdvancedCopy Manager provides operation management of replications by using the Advanced Copy function of ETERNUS Disk storage system.
As a result, the following replication operations become possible.

High-speed replication operation

The Advanced Copy function of ETERNUS Disk storage system is used to perform high-speed replication between volumes regardless of volume capacities. This function greatly reduces the impact of a replication operation.

Faster backups can be made by using update differential copies with the Advanced Copy function.

Consistency recovery function

The consistency recovery function resolves inconsistencies in the AdvancedCopy Manager replication management information caused by a system or other type of failure.

Main differences with backup function

The backup function provides policy and history management. These functions are not supported by the replication function.
The replication function replicates data between two volumes.

1.3.3 Web Console

The ETERNUS SF Web Console (hereafter referred to as "Web Console") can be used to perform environmental setup and employment functions for AdvancedCopy Manager. The main functions are as follows.

- Security management
  Operational restrictions of users of AdvancedCopy Manager can be performed using the GUI interface.

- Backup management
  AdvancedCopy Manager Backup management tasks (such as definition of a server or a device, backup, restoration, etc.) can be performed using the GUI interface.

- Replication management
  AdvancedCopy Manager replication management tasks (such as defining volumes and executing replications) can be performed using the GUI interface.

Backup management and replication management can also be performed at the command line as well as the Web Console.
For more information about the Web Console, refer to "ETERNUS SF Web Console Guide".

1.3.4 Copy Control Module

When the AdvancedCopy Manager Copy Control Module is used, it is possible to perform back ups from the server where the Copy Control Module is installed, using the Advanced Copy function of ETERNUS Disk storage system.
As a result, the AdvancedCopy Manager's agent need not be installed.
For details, refer to the "ETERNUS SF Operation Guide for Copy Control Module".

1.4 Operating Environment of AdvancedCopy Manager

This section explains the software requirements for the operation of AdvancedCopy Manager.

For the hardware requirements, refer to "Operating Environment" in the "ETERNUS SF Installation and Setup Guide".
For browsers on which Web Console can run, refer to "Supported web browsers" in the "ETERNUS SF Web Console Guide".

Point

....

To perform the Advanced Copy function of ETERNUS Disk storage system, it is necessary to register the Advanced Copy license to ETERNUS Disk storage system.
For details, refer to "License Management" in the "ETERNUS SF Installation and Setup Guide".

....

1.4.1 Software requirement of AdvancedCopy Manager's manager

The following table lists the software requirements of AdvancedCopy Manager's manager.
For about the supported operating systems, refer to "Operating environment of ETERNUS SF Manager" in the "ETERNUS SF Installation and Setup Guide".

For information on cluster software that can be linked with OSes on which AdvancedCopy Manager's managers are running, refer to "Operating Environment of ETERNUS SF Manager" in the "ETERNUS SF Cluster Environment Setup Guide".

Table 1.3 Software requirements of AdvancedCopy Manager's manager

Product name	Remarks
ETERNUS multipath driver	For multipath disk control and path load balancing.
Symfoware Server Advanced Backup Controller V10.1 or later	When a Management Server operates as a Managed Server, and when backup operation is used with Symfoware.
PRIMECLUSTER GDS Snapshot 4.1A30 or later	When a Management Server operates as a Managed Server, and when backup operation or replication operation is used with GDS Snapshot.

Note that the software version and level differs in accordance with the install destination system environment. Refer to the relevant software manual for details.

Note

....

Note the following points regarding SELinux (Security-Enhanced Linux):

- For RHEL5
  AdvancedCopy Manager's manager is not supported in an environment where SELinux (Security-Enhanced Linux) is enabled.
  Disable SELinux at servers that operate AdvancedCopy Manager's manager.

- For RHEL6

  AdvancedCopy Manager's manager is supported in an environment where SELinux (Security-Enhanced Linux) is enabled with the following updates:
  The supported mode is "enforcing".

  - Red Hat Enterprise Linux 6 Update 2 (for x86) or later

  - Red Hat Enterprise Linux 6 Update 2 (for Intel64) or later

  Any version numbers earlier than the updates stated above are currently not supported. Set SELinux to "disable" for operation.

....

1.4.2 Software requirement of AdvancedCopy Manager's agent

The following table lists the software requirements of AdvancedCopy Manager's agent.
For about the supported operating systems, refer to "Operating environment of AdvancedCopy Manager's agent" in the "ETERNUS SF Installation and Setup Guide".

However, for information on the supported operating systems of KVM guest, refer to "F.3 Software requirements".

For information on cluster software that can be linked with OSes on which AdvancedCopy Manager's agent is running, refer to "Operating Environment of AdvancedCopy Manager's agent" in the "ETERNUS SF Cluster Environment Setup Guide".

Table 1.4 Software requirements of AdvancedCopy Manager's agent

Product name	Remarks
ETERNUS multipath driver	For multipath disk control and path load balancing.
Symfoware Server Advanced Backup Controller V10.1 or later	When backup operation is used with Symfoware.
PRIMECLUSTER GDS Snapshot 4.1A30 or later	When backup operation or replication operation is used with GDS Snapshot.

Note that the software version and level differs in accordance with the install destination system environment. Refer to the relevant software manual for details.

Note

Note the following points regarding SELinux (Security-Enhanced Linux):

- For RHEL-AS4, RHEL-ES4, or RHEL5
  AdvancedCopy Manager's agent is not supported in an environment where SELinux (Security-Enhanced Linux) is enabled.
  Disable SELinux at servers that operate AdvancedCopy Manager's agent.

- For RHEL6

  AdvancedCopy Manager's agent is supported in an environment where SELinux (Security-Enhanced Linux) is enabled with the following updates:
  The supported mode is "enforcing".

  - Red Hat Enterprise Linux 6 Update 2 (for x86) or later

  - Red Hat Enterprise Linux 6 Update 2 (for Intel64) or later

  Any version numbers earlier than the updates stated above are currently not supported. Set SELinux to "disable" for operation.

1.5 Managing a Device on AdvancedCopy Manager

When using AdvancedCopy Manager, you should register the block type devices to operate.

For the supported devices on KVM guest, refer to "F.5 Supported devices".

1.5.1 Normal devices

It explains the each following operating system.

- For Red Hat Enterprise Linux 5, Red Hat Enterprise Linux 6 and SUSE Linux Enterprise Server 11

- For Red Hat Enterprise Linux AS v.4 and Red Hat Enterprise Linux ES v.4

For Red Hat Enterprise Linux 5, Red Hat Enterprise Linux 6 and SUSE Linux Enterprise Server 11

The following devices are supported in backup and replication management.

- sd driver resources

- Multi-path resources (MPLB resources and Device Mapper Multipath resources)

Moreover, the unit of management is a disk partition unit or a Logical Unit (disk) unit.

AdvancedCopy Manager uses the following device names for ordinary devices:

The device name set during the operating system startup (compatible device name)

Example:

```
/dev/sda1 (for disk partition units)
/dev/sda  (for Logical Unit units)
```

The multi-path device name started "/dev/mapper/" which is generated by Device Mapper facility of the operating system (Device Mapper Multipath device name)

Example:

```
/dev/mapper/*p1     (for disk partition units on RHEL5 or RHEL6)
/dev/mapper/*_part1 (for disk partition units on SUSE Linux Enterprise Server 11)
/dev/mapper/*       (for Logical Unit units)
```

The device name generated by the udev facility of the operating system (udev device name)

Example:

```
/dev/disk/by-id/xxxxxxxx-part1 (for disk partition units)
/dev/disk/by-id/xxxxxxxx       (for Logical Unit units)
```

(xxxxxxxx is the information generated from the identification information set in the disk.)

```
/dev/disk/by-path/yyyyyyyy-part1 (for disk partition units)
/dev/disk/by-path/yyyyyyyy       (for Logical Unit units)
```

(yyyyyyyy is the information generated from the disk location information.)

The udev device name takes priority and is used in environments that have udev device names.
The by-id name and by-path name can be used in udev device names, but the by-id name takes priority and is used in environments that have both by-id names and by-path names.

If the environment contains udev device names but you want to use the compatible device names for operations, then before the device information collection and reflection processing is performed, you must change the information collection mode so that udev device names are not used.
Refer to "stgxfwcmsetmode (Information collection mode configuration command)" for details.

🖐 Note
. . . .

- Refer to "10.1.1.9 Notes on executing Logical Unit (disk) unit copy" if operation is in Logical Unit (disk) units.

- The information collection mode can be set for each Managed Server. For the information collection mode, the udev device name and compatible device name cannot be the same in the same Managed Server.

- If Device Mapper Multipath device names are used in an environment in which Device Mapper Multipath is installed, configure them so that Device Mapper Multipath device names are used.
  Please don't use the compatible device name or udev device name that indicates each path for the multi-path.

- If operation is in the Logical Unit (disk) units, the Device Mapper Multipath device cannot be configured.

- If using the EC (Equivalent Copy) in the backup management, the Device Mapper Multipath device cannot be defined as the transaction volume and the backup volume.

- If using the EC (Equivalent Copy) or REC (Remote Equivalent Copy) in the replication management, the Device Mapper Multipath device cannot be defined as the copy source volume and the copy destination volume.

- If udev device names are used in an environment in which ETERNUS multi-path drivers are installed, configure them so that by-id names are used.
  If by-id names are not generated, change the settings so that by-id names are generated.

- Note the following if udev device names are used in an environment in which ETERNUS multi-path drivers are not installed:
  The udev device name may change if by-id is used and disks are exchanged, or if by-path is used and the connection configuration is changed. If the udev device name has been changed, the steps described in "8.5.5 Changing the device configuration" are required.

- -

For Red Hat Enterprise Linux AS v.4 and Red Hat Enterprise Linux ES v.4

The following devices are supported in backup and replication management.

- sd driver resources

- Multi-path resources (MPLB resources)

Moreover, the unit of management is a disk partition unit or a Logical Unit (disk) unit.

AdvancedCopy Manager uses the following device names for ordinary devices:

The device name set during the operating system startup (compatible device name)

Example:

```
/dev/sda1 (for disk partition units)
/dev/sda  (for Logical Unit units)
```

The device name generated by the udev facility of the operating system (udev device name)

Example:

```
/dev/disk/by-id/xxxxxxxx-part1 (for disk partition units)
/dev/disk/by-id/xxxxxxxx        (for Logical Unit units)
```

(xxxxxxxx is the information generated from the identification information set in the disk.)

```
/dev/disk/by-path/yyyyyyyy-part1 (for disk partition units)
/dev/disk/by-path/yyyyyyyy        (for Logical Unit units)
```

(yyyyyyyy is the information generated from the disk location information.)

The udev device name takes priority and is used in environments that have udev device names.
The by-id name and by-path name can be used in udev device names, but the by-id name takes priority and is used in environments that have both by-id names and by-path names.

If the environment contains udev device names but you want to use the compatible device names for operations, then before the device information collection and reflection processing is performed, you must change the information collection mode so that udev device names are not used.
Refer to "stgxfwcmsetmode (Information collection mode configuration command)" for details.

🖰 Note

- -

- Refer to "10.1.1.9 Notes on executing Logical Unit (disk) unit copy" if operation is in Logical Unit (disk) units.

- The information collection mode can be set for each Managed Server. For the information collection mode, the udev device name and compatible device name cannot be the same in the same Managed Server.

- If udev device names are used in an environment in which ETERNUS multi-path drivers are installed, configure them so that by-id names are used.
  If by-id names are not generated, change the settings so that by-id names are generated.

- Note the following if udev device names are used in an environment in which ETERNUS multi-path drivers are not installed:
  The udev device name may change if by-id is used and disks are exchanged, or if by-path is used and the connection configuration is changed. If the udev device name has been changed, the steps described in "8.5.5 Changing the device configuration" are required.

1.5.2 SDX objects of PRIMECLUSTER GDS

PRIMECLUSTER GDS (hereinafter referred to as GDS) is software whose main function is mirroring disk units.
AdvancedCopy Manager can backup and replicate the SDX objects of GDS.
SDX objects are virtual resources managed by GDS, and they include classes, groups, disks, volumes, and slices.
AdvancedCopy Manager can backup and replicate SDX objects in units of logical volumes or slices.
The format used to specify a device name in AdvancedCopy Manager is different depending on whether units of logical volumes or units of slices are to be used. Specify a device name as follows:

Logical volume unit operations

```
/dev/sfdsk/<class name>/dsk/<volume name>
```

Slice unit operations

```
/dev/sfdsk/<class name>/dsk/<volume name>:<device name used for volume configuration>
```

Either of the following is used as the device name for volume configuration:

sdxinfo DEVNAM values

Example:

```
sda
```

The device name generated by the facility indicated by the sdxinfo DEVNAM value (udev device name)

Example:

```
/dev/disk/by-id/xxxxxxxx
```

(xxxxxxxx is the information generated from the identification information set in the disk)

```
/dev/disk/by-path/yyyyyyyy
```

(yyyyyyyy is the information generated from the disk location information.)

The udev device name takes priority and is used in environments that have udev device names. The by-id name and by-path name can be used in udev device names, but the by-id name takes priority and is used in environments that have both by-id names and by-path names. If the environment contains udev device names but you want to use the sdxinfo DEVNAM value for operations, then before the device information collection and reflection processing is performed, you must change the information collection mode so that udev device names are not used.
Refer to "stgxfwcmsetmode (Information collection mode configuration command)" for details.

Example:

When the class name is 'class1', the volume name is 'volume1', the sdxinfo DEVNAM value is 'sda', and the udev device name indicated by sda is /dev/disk/by-id/SFUJITSU_MAP3367NC_UPP4P4307J95

Logical volume unit operations

```
/dev/sfdsk/class1/dsk/volume1
```

Slice unit operations (the sdxinfo DEVNAM value)

```
/dev/sfdsk/class1/dsk/volume1:sda
```

Slice unit operations (the udev device name)

```
/dev/sfdsk/class1/dsk/volume1:/dev/disk/by-id/SFUJITSU_MAP3367NC_UPP4P4307J95
```

🛈 Note
. . . .
To use the udev device name in an environment in which the ETERNUS multi-path driver is installed, enable use of the by-id name.
If the by-id name has not been generated, modify the settings so that it is generated.
. . . .

🛈 Note
. . . .
Note the following points when using the udev device name in an environment in which the ETERNUS multi-path driver is not installed.

- The udev device name may change if by-id is used and disks are exchanged, or if by-path is used and the connection configuration is changed. If the udev device name has been changed, perform the steps described in "8.5.5 Changing the device configuration".
. . . .

🛈 Note
. . . .

- PRIMECLUSTER GDS 4.1A30 or later and PRIMECLUSTER GDS Snapshot 4.1A30 or later are required to perform separate operations for each logical volume.

- Operation in units of logical volumes and operation in units of slices at the same time on a single Managed Server is not possible. Note that for cluster operations, operating in units of logical volumes or in units of slices is possible in separate Managed Server transactions.
. . . .

For details on backup of SDX objects, refer to "3.8 SDX Object Operations". For details on replication of SDX objects, refer to "7.8 SDX Object Replication".

Note that the software versions and levels vary in accordance with the environment at the installation destination system. Refer to the relevant software manuals for details.

1.5.3 Volumes under control of VERITAS Volume Manager (VxVM)

AdvancedCopy Manager can back up or replicate volumes under the control of VxVM (referred to as VxVM volumes in this document) in the following units:

- Physical disk that constitutes a volume group

The VM disks (physical disks) on which a sub-disk used by a logical volume exists are used as the management units.

Figure 1.5 Example of VxVM configuration

Since AdvancedCopy Manager copies an entire VM disk, not only the shared area that is a sub-disk but also the public region containing VxVM-internal configuration information is copied.
When backup or replication is performed, therefore, the integrity of VxVM configuration information must be maintained during copying.

Specify a device name as follows:

  - /dev/vx/dmp/sd#

Note

. . . .

  - Operations are not supported under naming convention operations based on enclosures in VxVM.

  - For operation in a cluster configuration, the device name (/dev/sd#) of the physical disks that comprise the volume group must be the same at all servers that comprise the cluster, and the ETERNUS Disk storage system's disk indicated by the device name must be the same.

. . . .

Chapter 2 Starting and Stopping Daemons

This chapter explains how to activate and stop daemons provided by AdvancedCopy Manager.

For more information on how to activate and stop the daemons described in this chapter, refer to "9.5 Start/Stop Daemon Commands".

2.1 Daemons of AdvancedCopy Manager

The Manager is configured with the following daemons:

Table 2.1 Daemons of Manager

Daemon type	Description
The communication daemon	Communicates with the Managed Server
The RMI daemon	Manages the GUI display
The RDB daemon	Accesses the repository

The Agent is configured with the following daemon:

Table 2.2 Daemons of Agent

Daemon type	Description
The communication daemon	Communicates with the Management Server and other Managed Servers

When the Management Server is in a clustered system, the following daemons compose Management Server transactions:

Table 2.3 Daemons used for Management Server transactions

Daemon type	Description
The transaction-intended communication daemon	Communicates with the Managed Server
The RMI daemon	Manages the GUI display
The RDB daemon	Accesses the repository

When the Managed Server is in a clustered system, the following daemon composes Managed Server transactions:

Table 2.4 Daemons used for Managed Server transactions

Daemon type	Description
The transaction-intended communication daemon	Communicates with the Management Server and other Managed Servers

Note

The following daemon cannot be used for a Management Server transaction or Managed Server transaction:

- Communication daemon (local daemon)

2.2 Starting Daemons

Daemons must be running on the Management Server and Managed Servers.

If for whatever reason daemons stop, the daemons must restarted on each server.

```
# /opt/swstorage/bin/startacm
```

Note

. . . .

- Before executing the startacm command, confirm that the SWSTGNODE environment variable is not defined.

- Confirm that /var/tmp directory as a work directory exists.

- If the message below is displayed when trying to start the daemon, then stop it and start it again:
  "Symfoware/RDB for AdvancedCopy Manager failed to start !!(version missing)"

- If starting daemon fails, confirm that IP address and name of the local server in /etc/hosts file are validate.

. . . .

When the Management Server or Managed Server is in a clustered system, start up the daemon(s) from the clustered system. For the start procedure, refer to the relevant cluster software manual.

To start the daemon separately

To start the daemon separately on AdvancedCopy Manager transactions in a clustered system, follow the steps described below:

- Communication daemon:

  You must start the daemon for each AdvancedCopy Manager transaction (including local transactions) using the following procedure for the respective transactions.

    1. Execute the following command.

       ```
       # ps -elf | grep stgxfws
       ```

    1. For environment variable SWSTGNODE, specify the name of the logical node on which the transaction is to be performed.

    2. Start the communication daemon.

       For the startup command, refer to "9.5.1 Starting and stopping the communication daemon".

- For daemons other than the communication daemon:

    1. For environment variable SWSTGNODE, specify the name of the logical node on which the transaction is to be performed.

    2. Refer to the following sections to start the daemon.

        - RMI daemon: "9.5.2 Starting and stopping the RMI daemon"

        - RDB daemon: "9.5.3 Starting and stopping the RDB daemon"

2.3 Stopping Daemons

Stop daemons from the Management Server and Managed Server. Normally, they are automatically stopped on system shutdown.

To stop daemons on a server, execute the following command on the server:

```
# /opt/swstorage/bin/stopacm
```

When the Management Server or Managed Server is in cluster operation, stop the daemon(s) from the clustered system. For the stop procedure, refer to the relevant cluster software manual.

To stop the daemon separately

To stop the daemon separately on AdvancedCopy Manager transactions in a clustered system, follow the steps described below:

- Communication daemon:

  You must stop the daemon in each AdvancedCopy Manager transaction (including local transactions) using the following procedure for the respective transactions.

  1. Execute the following command.

     ```
     # ps -elf | grep stgxfws
     ```

  1. For environment variable SWSTGNODE, specify the name of the logical node on which the transaction is to be performed.

  2. Stop the communication daemon.

     For the stop command, refer to "9.5.1 Starting and stopping the communication daemon".

- For daemons other than the communication daemon:

  1. For environment variable SWSTGNODE, specify the name of the logical node on which the transaction is to be performed.

  2. Refer to the following sections to stop the daemon.

     - RMI daemon: "9.5.2 Starting and stopping the RMI daemon"

     - RDB daemon: "9.5.3 Starting and stopping the RDB daemon"

Chapter 3 Backup Operation

This chapter provides details of the backup operation performed by AdvancedCopy Manager AdvancedCopy Manager backups are created in units of partitions.

For information on the relationships between the configurations of supported devices and the copy operation, refer to "1.5 Managing a Device on AdvancedCopy Manager".

3.1 Overview

This chapter describes the Backup and Restore functions of AdvancedCopy Manager.

AdvancedCopy Manager performs high-speed backup of volumes of ETERNUS Disk storage system, regardless of volume capacities, using the Advanced Copy functions, OPC and EC.

Backup that uses the OPC function is called "snapshot fast backup", while backup that uses the EC function is called "synchronized high-speed backup".

Automatic snapshot backup can be scheduled without first estimating the time required, because it does not use synchronous backup processing. However, when the save number is 1, a state may occur in which there is no backup data because a snapshot backup deletes the history prior to copy processing.

As synchronized high-speed backup deletes the history after backup processing is complete, there is no status where backup data does not exist, even though the generation number is 1, but synchronized high-speed backup does requires one more backup volume generation than is required by snapshot fast backup. Synchronized high-speed backup also requires synchronized backup processing before the backup execution command is entered. Therefore, execution of automatic synchronized high-speed backup must be scheduled to allow enough time for synchronized processing to finish first. Synchronized backup processing is then started based on the estimated time, and the backup execution command is entered after equivalency maintained status has been achieved.

Backing up an Oracle database can be performed using the sample script which AdvancedCopy Manager offers.

AdvancedCopy Manager automatically manages and selects a medium (backup volume) to be used for backup.

AdvancedCopy Manager also manages the backup carriers and generations, and allows you to restore data from the latest or previous generations of backed up data.

Figure 3.1 Managing backup volumes
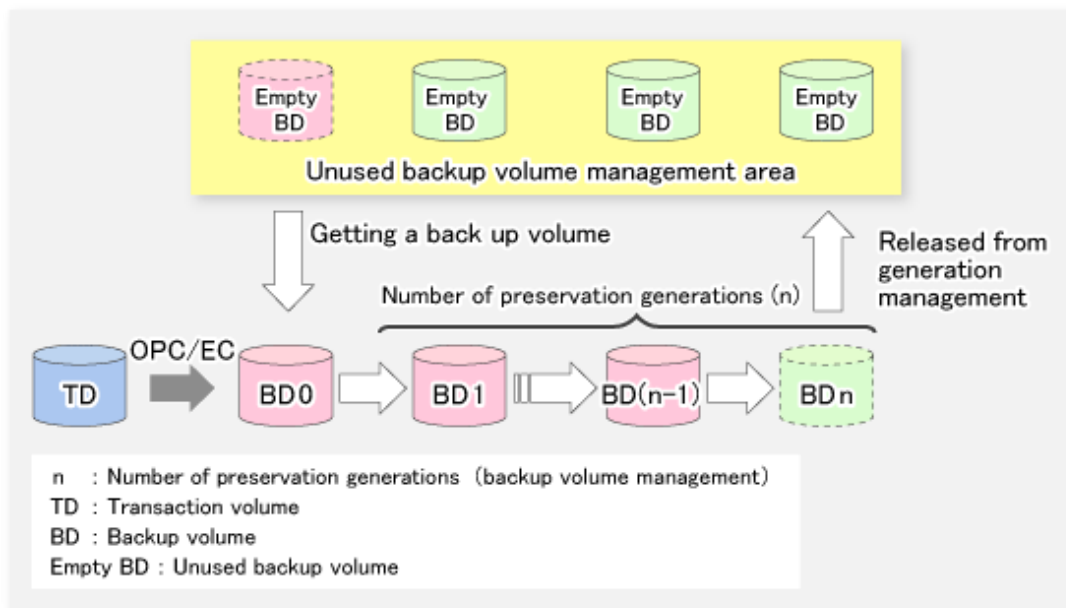
3.1.1 Snapshot backup processing

The OPC(One Point Copy) function of ETERNUS Disk storage system is used to copy data from a transaction volume to an unused backup volume.

To perform a snapshot fast backup, take the following steps:

1. When swstbackup (Backup execution command) is executed ((1) in the following figure), an unused backup volume with the same size as the relevant transaction volume is selected from the unused backup volume management area.
   If the number of backup volumes exceeds the number of preservation generations by performing backup, the history of the oldest backup volume is deleted, and then the volume is selected as an unused backup volume.

2. Data on the transaction volume will be copied to the selected unused backup volume. The logical copy is instantly performed and the physical copy is started on the hardware ((2) to (4) in the following figure).

3. Backup carrier information will be set.

Figure 3.2 Snapshot fast backup processing

The snapshot fast backup is completed as soon as you confirm the execution of the backup. The actual data copying will be performed internally by the OPC function of ETERNUS Disk storage system.
The QuickOPC function of the snapshot high-speed backup creates a snapshot image by only copying the data updated since the start of previous backup. This reduces the time for physical copying significantly.

For information on the QuickOPC function, refer to "Chapter 6 Backup Operation using QuickOPC Function".

📝 Note
. . . .
If all the generations of backup volumes are being managed, and if no backup volume can be obtained from the unused region in the backup volume management area, backup copies cannot be created.
. . . .

3.1.2 Synchronized high-speed backup processing

The EC function of ETERNUS Disk storage system is used to copy data from a transaction volume to an unused backup volume.

To perform synchronized high-speed backup, take the following steps:

1. When swststartsync (Backup synchronous processing start command) is executed ((1) in the following figure), an unused backup volume with the same size as the relevant transaction volume is selected from the unused backup volume management area and backup synchronous processing begins.

2. The transaction and backup volumes reach an equivalent state. From this point on, equivalence between the transaction volume and backup volume is maintained. This status is referred to as equivalency maintenance status ((3) in the following figure).

3. When swstbackup (Backup execution command) is executed ((4) in the following figure), the backup synchronization processing stops and the backup history information is registered.

4. When the generation number of the newly backed up data exceeds the number of generations of backup history being kept, the oldest backup volume is treated as an unused backup volume.

Note

The swstbackup (Backup execution command) cannot be executed before the transaction and backup volumes have become equivalent.

Figure 3.3 Backup using the synchronized high-speed backup function

During synchronous high-speed backup using the Suspend/Resume function, the equivalency maintenance status of EC can be set to Suspend or Resume. The synchronous backup process is therefore faster because the function can create a differential copy from the Suspend state.

For information on the Suspend/Resume function, refer to "Chapter 5 Backup Operation using Suspend/Resume Function".

Note

If all generations of backup volumes are being managed, and if no backup volume can be obtained from the unused region in the backup volume management area, backup copies cannot be created.

3.1.3 Restoring Data

This section describes how data is restored in AdvancedCopy Manager.
The restoration operation refers to restoring data from a backup volume to a transaction volume using backup carrier management information.

Data restore from a backup volume refers to using the OPC function provided by ETERNUS Disk storage system to restore data from a backup volume to a transaction volume.

If the data on a backup volume under carrier management is restored to a transaction volume, the carrier management information will not change.

Note that data will be restored to a transaction volume that was the backup source. However, you can change the restore destination.

📁 Note

Data updated between the last backup and a subsequent restoration will be lost as the transaction volume is overwritten by the backup volume.

Figure 3.4 Notes on restore operations in AdvancedCopy Manager
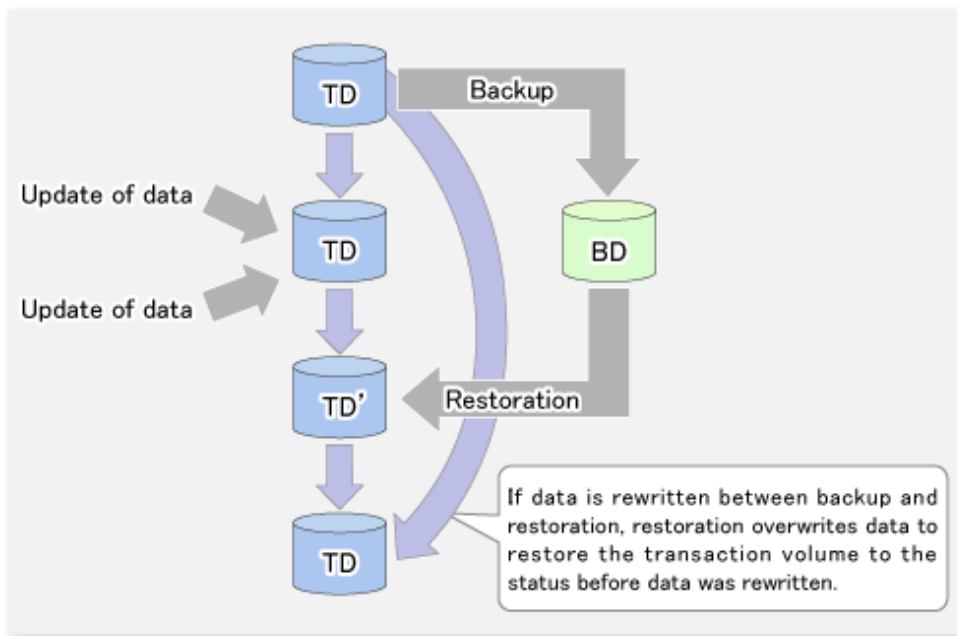
TD: Transaction volume
BD: Backup volume

3.2 Operation Design

This section describes how to configure a backup system in AdvancedCopy Manager.

When configuring a backup system, you will need to consider the following issues:

- 3.2.1 Selecting a server on which backups are performed

- 3.2.2 Selecting a backup source

- 3.2.3 Configuring a backup policy

- 3.2.4 Preparing a backup volume

- 3.2.5 Preparing the automatic backup operation

3.2.1 Selecting a server on which backups are performed

The backup operation can be performed on one of the following servers:

Management Server

Multiple Managed Servers are unified and managed centrally. The Management Server can also operate as a Managed Server.

Managed Server

AdvancedCopy Manager operations can be performed on the Managed Server.

Figure 3.5 Configuration of servers on which the backup operation should be performed

3.2.2 Selecting a backup source

AdvancedCopy Manager backs up data in units of partition. A partition that stores work data is called a transaction volume.

A transaction volume to be backed up must reside on the ETERNUS Disk storage system.

For the supported units, refer to "1.5 Managing a Device on AdvancedCopy Manager".
For information on setting up a transaction volume, refer to "3.4.6 Setting the operation type for a device".

Note

- Do not back up the system disks or partitions on which AdvancedCopy Manager is installed.

- Do not select the logical volume and volume group of LVM (Logical Volume Manager) as a backup volume.

- Refer to "" for points to consider when deciding on a candidate volume for backup. Some points to consider include but may not be limited to:

    -

    -

    -

    -

3.2.3 Configuring a backup policy

To configure a backup policy on a transaction volume, set the following items:

The number of preservation generations

The number of preservation generations refers to the maximum number of generations of backup data that should be kept.

The snapshot fast backup first releases the oldest generation from generation management. If a system failure or any other failure occurs while the backup is running, the oldest generation may have already been released and there may not be as many generations of backup data as required. Therefore, Fujitsu recommends that you re-execute backup immediately to create the required number of generations of backup data.

When performing snapshot type high-speed backup using only one preservation generation, we recommend you also back up data to a tape.

A synchronous high-speed backup operation releases the oldest generation from generation management after backup of the latest generation is completed. Therefore, the required number of backup volumes is equal to the maximum number of backup generations + 1.

Interval days

Interval days refers to the time interval (in days) between successive backups.

A notification that backup is overdue is displayed if a period longer than the specified interval has passed since the last backup date, but a backup will not performed automatically in this instance even if you have specified the interval days between backups.

For information on setting a backup policy, refer to "".

3.2.4 Preparing a backup volume

This section describes how to prepare the backup volume.

The backup unit of AdvancedCopy Manager is a partition. In AdvancedCopy Manager, the partition which stores the transaction data used as the candidate for backup is called the backup volume.

A destination backup volume must be placed on the same ETERNUS Disk storage system.

A backup volume must have the same size as a transaction volume. For more information, refer to "" under "".

The number of backup volumes is determined by whether snapshot fast backup or synchronized high-speed backup is performed:

Table 3.1 Number of backup volumes required for backup operations

Backup method	Number of required backup volumes
Snapshot fast backup	Number of backup generations
Synchronized high-speed backup	Number of backup generations + 1

When backup policies have already been set for multiple transaction volumes, the registration of a backup policy for a newly registered transaction volume requires that the following number of backup volumes be registered beforehand:

Table 3.2 Number of backup volumes required for backup operations (if a backup policy has been set)

Backup method	Number of required backup volumes
Snapshot fast backup	Total number of backup generations in the registered backup policies + number of generations in the new backup policy being set
Synchronized high-speed backup	Total number of backup generations in the registered backup policies + number of registered transaction volumes + number of generations in the new backup policy being set + 1

For information on configuring a backup volume, refer to "3.4.6 Setting the operation type for a device".

🔰 Note

- Do not configure AdvancedCopy Manager to backup a system partition or a partition where AdvancedCopy Manager is installed.

- Do not select the logical volume and volume group of LVM (Logical Volume Manager) as a backup volume.

🅿 Point

If the number of backup volumes that can be prepared in synchronized high-speed backup is only the number of preservation generations, the backup operation can be carried out by saving the backup volumes to secondary media, such as tapes.

To do this, follow the procedure below:

1. Start backup synchronization.

2. Confirm the equivalency maintain status.

3. Execute backup.

4. Save the backup volumes to secondary media.

5. Delete history information.

6. Return to step 1.

Be aware that, if any history information is deleted for any reason, the restore execution command cannot be used to restore data from the period beginning with the time of deletion until the time the next backup is completed inclusive. In this event, the data must be restored directly from the secondary media where applicable.

3.2.5 Preparing the automatic backup operation

The backup operation of AdvancedCopy Manager can be automated by using Systemwalker Operation Manager.

Automatic snapshot fast backup

Automatic snapshot fast backup can be executed by using one of the following two methods:

- Start the backup command as part of a batch job.

Figure 3.6 Start in a job net

- Start the backup command at the specified time.

Figure 3.7 Start at a specified time

Automatic synchronized high-speed backup

For synchronized high-speed backup, the transaction and backup volumes must have maintained equivalency status at the scheduled time for backup.

For this reason, if synchronized high-speed backup is automated, the synchronized backup processing start command and the backup execution command must be started separately, as described below.

Starting the synchronized backup processing start command

Estimate the expected time required for the synchronized backup processing.

Execute the synchronized backup processing start command. To do this, specify the time obtained by subtracting the estimated time (or more) from the time that the backup execution command is to start. The transaction and backup volumes will then have maintained equivalency status when backup is executed.

The time between the start of synchronized backup processing and the transition to maintenance equivalency status depends on the capacity of the volume to be copied.

Starting the backup execution command

Use either of the following methods to automatically execute the backup execution command:

- Start the backup execution command as part of a batch job.

Figure 3.8 Start in a job net

Note

If there are other jobs (jobs A and B in the figure) to be executed before the backup execution command in the job net, the start time of the backup execution command cannot be determined. Assume that the start time of the backup execution command is the time obtained by adding the shortest time required for jobs A and B to the job net start time. Determine the start time of the synchronized backup processing start command by subtracting the estimated time from the start time of the backup execution command.

- Start the backup execution command with the time specified.

Figure 3.9 Start at a specified time

3.3 Flow of Operation

The following diagram shows the flow of the backup operation on a transaction volume:

Figure 3.10 Flow of backup operation on a transaction volume

| Activating a daemon |
| Starting the Web console |
| When a command performs subsequent operations, starting of an initial screen is unnecessary |
| Registering a Managed Server |
| For the first time only |
| Fetching device information on a Managed Server |
| For the first time or only when a device is changed |
| Configuring the environment for a backup operation server |
| Only when operation is changed |
| Configuring the operation type for a device |
| Only when operation is changed |
| Configuration options |
| Only when inter-box backup is performed |
| Configuring backup policies |
| Only when operation is changed |
| Customizing pre-processing and post-processing scripts |
| Only when the operation is changed |
| Preparing a device map file |
| Only when the backup destination is changed |
| Operation |
| Stopping a daemon |

3.4 Preparing to Start the Backup

This section describes the preparations that are required to perform the backup operation.

3.4.1 Activating daemons

Before the backup operation is started, the daemons of AdvancedCopy Manager must be started on both the Management Server and the Managed Servers. Normally, these daemons are automatically started up when the system is started. If the startup fails the system for some reason or if a daemon has been stopped, you will need to start the daemon on each server.

For information on starting a daemon, refer to "Chapter 2 Starting and Stopping Daemons".

3.4.2 Starting the Web Console

Start the Web Console. For details of this, refer to "How to start Web Console" in the "ETERNUS SF Web Console Guide".

3.4.3 Registering a Managed Server

The Management Server registers the Managed Server that is to be managed.

This process is carried out using Web Console.

3.4.4 Fetching device information on a Managed Server

This process can be carried out using either Web Console or stgxfwcmsetdev (Device information collection/reflection command).

📒 Note

- As part of replication management, device information on a Managed Server must first be stored in the repository.
  There may not be sufficient area available for updating the repository if information is to be obtained for devices under two or more Managed Servers. Check if the area required when updating the repository is available before obtaining the device information. If there is insufficient area, allocate the required area, then obtain the device information.

  Refer to step 2 of "8.2.3 Troubleshooting: Failure in accessing a repository database" for details on checking the required area.

- The time taken to add new devices is in proportion to the total number of devices defined in the selected Managed Server. For this reason, if a large number of devices have been defined, execute the process only when the CPU load on the system is low. As a guide, each operation takes about 0.5 seconds per device (ie, partition) with no load.

3.4.5 Configuring the environment for a backup operation server

Configuring the environment for a backup server can be carried out using either Web Console or swstsvrset (Managed Server configuration information setting command).
When performing backup using commands only, the swstsvrset (Managed Server configuration information setting command) must be executed at the time that operations start.

🅿 Point

If the Management Server is combined with the Managed Server, these environment settings are also required for the Management Server. If the Managed Server environment setup has already been carried out, this process is not required.

3.4.6 Setting the operation type for a device

Set up the volume type for a device that is connected to the Managed Server. The setting of the type is explained below.

- Defining the volume to be backed up as a transaction volume.

- Preparing backup volumes. Prepare the number of volumes shown below in accordance with the operation type, each volume having the same partition size as the transaction volume, and define them as backup volumes.

Table 3.3 Number of backup volumes required for backup operations

Backup method	Number of required backup volumes
Snapshot fast backup	Number of backup generations
Synchronized high-speed backup	Number of backup generations + 1

This process can be carried out using either Web Console or swstdevinfoset (Device information setting command).

⚑ Note

. . . .

- To change the configuration of a partition that has been registered as either a transaction volume or a backup volume, take the following steps:

    1. Delete the registered configuration, and then change the configuration.

    2. Attempt to "3.4.4 Fetching device information on a Managed Server".

    3. Register the information by executing swstdevinfoset (Device information setting command).

- The backup volume must be the same size (have the same number of bytes) as the transaction volume.

- Backup and backup volumes are not supported for the system partition or the partition where AdvancedCopy Manager is installed.

- The mounted device cannot be set to backup volume. However, it is possible to set it to the transaction volume.

- Do not select the logical volume and volume group of LVM (Logical Volume Manager) as a backup volume.

- Logical Units (disks) cannot be set for a transaction volume or a backup volume.

. . . .

ⓟ Point

. . . .

The registered volumes can be confirmed using either Web Console or swstdevdisp (Device usage status display command).

. . . .

3.4.7 Setting a backup policy

For all transaction volumes registered on a Managed Server, set the following backup policy parameters:

- Preservation generations

    This refers to how many generations of backup data should be retained.

- Interval days

    This is the number of days until the next backup should be performed. If the specified number of days is exceeded and no backup has been performed, you will be prompted to perform a backup in the Web Console.

This process can be carried out using either Web Console or swstbkpolset (Backup policy setting command).

⚑ Note

. . . .

- Even if the number of interval days has been specified, AdvancedCopy Manager does not automatically invoke the backup when the date for the next backup is reached. This value is used only to prompt the user to perform backup when the number of backup interval days is exceeded.

- When you set a backup policy, there must be as many registered backup volumes as required to perform backup according to the specified backup policy. For information on the necessary number of backup volumes, refer to "3.2.4 Preparing a backup volume".

- Even though the number of backup volumes required for the backup may not have been registered when backup policies for synchronized high-speed backup were set, this number can be set by default provided the number of backup volumes required for snapshot fast backup has been previously registered. Otherwise, it may not be possible to execute synchronized high-speed backup.

 Point

The registered backup policy can be confirmed using either Web Console or swstbkpoldisp (Backup policy display command).

3.4.8 Customizing pre-processing and post-processing

In AdvancedCopy Manager, backup or restore processing must be performed while transaction volumes are unmounted. These essential steps, mount and unmount, are typically performed using pre-processing and post-processing scripts which may need to be customized to include the mount/unmount operation and any other actions that may be deemed necessary at the time.
The backup or restore processing is not executed when the transaction volume cannot be unmounted.

Pre-processing and Post-processing scripts for Backup need to be customized in the following conditions.

- For operational reasons, to avoid unmount and mount processing of transaction volumes that comprise a file system

- To add special pre-processing and post-processing to the corresponding scripts.

For information about the customization method, refer to "Appendix A Pre-processing and Post-processing of Backup and Restoration" for more information.

 Point

When AdvancedCopy Manager is upgraded

Where there has been an upgrade to existing versions of AdvancedCopy Manager or its components, any scripts used with AdvancedCopy Manager or its components should be reviewed and customized accordingly after upgrading rather than continuing to reuse scripts from previous versions.

3.4.9 Preparing a device map file

When using AdvancedCopy Manager to back up data, a backup volume with the same capacity as a transaction volume is automatically selected from the available registered backup volumes.

If a specific backup volume is to be used to back up a given transaction volume, a "Device map file" must be created. A "Device map file" is a file defining the relationship between the transaction and a specific backup volume.

A device map file may be created anywhere on the Managed Server that performs backup. If you specify this file while performing a backup, the relationship as defined in the device map file is automatically recognized for that backup operation.

 Point

To manage multiple generations of backups using defined transaction volume to backup volume relationships, multiple device map files must be prepared.
The device map file used when backing up or starting synchronous processing should meet either of the following requirements:

- Either an unused backup volume must be specified

- A backup volume detailed in the history information that was deleted by this backup must be specified.

Care should taken to ensure that the device map file is configured and used correctly where the backup system uses two or more backup volumes.

If the number of backup volumes prepared for synchronized high-speed backup is no more than the number of preservation generations, the backup can be performed by saving the backup volumes to secondary media, such as tape, and then deleting corresponding history information.

In such cases, the device map file specified at the start of synchronous backup processing is the file specifying the backup volumes that become available for reuse when history information is deleted.

. . . .

3.4.9.1 Describing a device map file

The following provides an example of a device map file:

```
# A comment line starts with "#".
# Business volume name      Output destination backup volume name

# Normal Disks
/dev/sdc1            /dev/sdv1
/dev/sdc2            /dev/sdv2
/dev/sdc3            /dev/sdv3

# A comment can be described here also.
/dev/sdd1            /dev/sdz1
/dev/sdd2            /dev/sdz2
/dev/sdd3            /dev/sdz3


/dev/sdg4            /dev/sdl4
/dev/sdg5            /dev/sdl5
:
```

The rules for creating a device map file are as follows:

- Describe a transaction volume name and a corresponding target backup volume name in one line. Separate the transaction volume name and backup volume name with one or more space or tab characters.
One or more space or tab characters may be used as delimiters between the beginning of a line and a transaction volume name or between the end of a target backup volume name and the end of a line (new line character).

- A blank line (Space or tab characters) may be included in a file.

- Anything after the symbol '#' is a comment. If the comment uses more than a single line, then each line must be preceded by a '#' character.

- You cannot specify more than one backup volume for the same transaction volume in one device map file. If you do, only the first occurrence is valid. While the device map file will still be read, subsequent occurrences will be ignored.

- A device map file may include descriptions of transaction volumes other than the processing target. Refer to the description example provided in the beginning of this chapter.

3.5 Operation

This section describes how to perform a backup in AdvancedCopy Manager.

Before performing the backup operation, configure the environment settings that are required for the backup according to "3.4 Preparing to Start the Backup".

There are two kinds of backup in AdvancedCopy Manager:

- Snapshot type high-speed backup

- Synchronous high-speed backup

3.5.1 Performing snapshot high-speed backup

Snapshot high-speed backup can be carried out using either Web Console or swstbackup (Backup execution command).

To check the backup status, use swstbackstat (Backup execution status display command).

🛈 Note
. . . .

When the QuickOPC function is used for snapshot high-speed backup, the command options listed below cannot be performed from the Web Console.
Enter the appropriate command for each operation.

- Execute the backup command with the -T option specified swstcanceltrk (Tracking cancel command).

. . . .

3.5.2 Performing synchronized high-speed backup

This process can be carried out using either the Web Console or commands.
If operating with commands, perform the backup synchronous process, then carry out the backup for the transaction volume that you need to backup. For details on commands, refer to the following:

- Starting backup synchronous process: swststartsync (Backup synchronous processing start command)

- Checking the status of backup synchronous process: swstsyncstat (Backup synchronous processing progress display command)

- Executing backup: swstbackup (Backup execution command)

- Checking backup status: swstbackstat (Backup execution status display command)

🛈 Note
. . . .

- The swstbackup (Backup execution command) causes an error if it is executed before the transaction and backup volumes have entered equivalency maintenance status.

- The device map cannot be specified while backup is processing.

- When the "Chapter 5 Backup Operation using Suspend/Resume Function" is used for synchronous high-speed backup, the operations listed below cannot be performed from the Web Console. Enter the appropriate command for operation.

  - If one transaction volume has multiple synchronous processing (ie, EC sessions), specify the backup volume and cancel only the specified EC session swstcancelsync (Backup synchronous processing cancel command).

. . . .

3.5.3 Restoring Data

This section describes how to restore a volume or other data that has been backed up.

3.5.3.1 Restoring an entire volume

All of the data from a volume that has been saved with the snapshot high-speed backup can be restored using either Web Console or swstrestore (Restore execution command).
To check the restore status, use swstreststat (Restore execution status display command).

🛈 Note
. . . .

If restoring for a general file system, stop operations.

. . . .

3.5.3.2 Restoring an individual file

Restore an individual file as shown below. Note that you cannot restore an individual file from the Web Console.

1. Mount the backup volume. Check the backup volume using the Web Console or swsthistdisp (History information display command).

2. Copy files to restore using the cp command.

3. Unmount the backup volume.

3.5.4 Managing a backup history

Backup history information can be displayed to check information such as the backup volume names and backup date of backup data under generation management in a specific transaction volume.

P Point

- Backup history information is only used to manage information stored on a backup volume.

- Any backup history information exceeding the number of preservation generations defined in a backup policy will automatically be deleted.

3.5.4.1 Displaying a backup history

Backup history can be displayed using either Web Console or swsthistdisp (History information display command).

3.5.4.2 Deleting a backup history

Backup history can be deleted using either Web Console or swsthistdel (History information deletion command).

3.6 Changing the setting information

It will be necessary to change the configuration of AdvancedCopy Manager if changes are made to the configuration of a Managed Server or a device associated with a Managed Server used in the backup. This section describes how to perform these changes.

3.6.1 Adding a device

After a device is added to any Managed Server, incorporate it into the backup as described below.
This process can be carried out using either the Web Console or commands.
For information on adding a device using commands, refer to the following steps:

1. Use stgxfwcmsetdev (Device information collection/reflection command) to acquire information for the newly added device, and accept management information.

2. Use swstdevinfoset (Device information setting command) to set up the operation type of the added device.

3. If the added device was defined as a transaction volume, use swstbkpolset (Backup policy setting command) to set up the backup policy.

Note

- - System resources such as disk and memory resources may be insufficient when a device is added. Re-estimate the system resource requirements before adding a device.

    Refer to "Required resources for operation" under "Operating environment of AdvancedCopy Manager's agent" in the "ETERNUS SF Installation and Setup Guide" for details.

- - There may not be sufficient area available for updating the repository when a device is added. Check if the area required when updating the repository is available before obtaining the device information. If there is insufficient area, free an area, then obtain the device information.

    Refer to step 2 of "8.2.3 Troubleshooting: Failure in accessing a repository database" for details on checking the required area.

- - The repository (directory for the DB space that stores repository data) may be too small when a device is added. Re-estimate the repository size before adding a device. Refer to "Estimation of database size" in the "ETERNUS SF Installation and Setup Guide" for details.
    If the estimation result is less than 65 megabytes, the size does not need to be increased. If it is more than 65 megabytes, extend the size before adding a device. Refer to "8.2.1 Troubleshooting: Insufficient free space in a repository" for details.

3.6.2 Deleting a device

This section describes how to delete different types of devices.

3.6.2.1 Deleting a transaction volume

This process can be carried out using either the Web Console or commands.
For information on the operation method using commands, refer to the following steps.

1. Use swsthistdel (History information deletion command) to delete the backup history.

2. Use swstbkpoldel (Backup policy deletion command) to delete the backup policy.

3. Use swstdevinfoset (Device information setting command) to delete the device information for the deleted transaction volume.

Point

Ensure that all of the steps are complete before removing the transaction volume for deletion.

3.6.2.2 Deleting a backup volume

This process can be carried out using either the Web Console or commands.
For information on the operation method using commands, refer to the following steps:

1. If deleting the backup volume, ensure that you first use swstbkpoldisp (Backup policy display command) to check for a previously defined backup policy.

Note

When a backup policy is deleted, it may result in a previously defined backup policy for an existing transaction volume no longer being valid (for example, if the number of backup volumes is less than the number of preservation generations defined in the policy). As a result, those backup operations may not be able to continue. If there is no space in the backup volume, first register an alternative backup volume to replace the volume to be deleted, and then delete the target backup volume.

2. Use swsthistdel (History information deletion command) to delete the backup history, removing the relevant volume from the backup history. If there is no backup history, proceed to the next step.

3. Use swstdevinfoset (Device information setting command) to delete the device information for the deleted backup volume.

🅿 Point
. . . .

Ensure that all of the steps are complete before removing the backup volume for deletion.

. . . .

3.6.3 Adding the Managed Server

The process of adding a Managed Server to the backup operation is carried out using Web Console.

The required steps for adding a Managed Server are described below.

1. Register a Managed Server. For more information, refer to "3.4.3 Registering a Managed Server".

2. Fetch all the information on the devices connected to the newly added Managed Server. For more information on this, refer to "3.4.4 Fetching device information on a Managed Server".

3. Set the environment information for the Managed Server. For more information on this, refer to "3.4.5 Configuring the environment for a backup operation server".

4. Set the operation type for the device for which information has been fetched in Step 2. For more information on this, refer to "3.4.6 Setting the operation type for a device".

5. Configure a backup policy for all the transaction volumes registered in Step 4. For more information on this, refer to "3.4.7 Setting a backup policy".

🅖 Note
. . . .

- System resources such as disk and memory resources may be insufficient when a Managed Server is added. Re-estimate the system resource requirements before adding a Managed Server.

  Refer to "Required resources for operation" under "Operating environment of ETERNUS SF Manager" in the "ETERNUS SF Installation and Setup Guide" for details.

- There may not be sufficient area available for updating the repository when a Managed Server is added. Check if the area required when updating the repository is available before adding the Managed Server. If there is insufficient area, free an area, then add the Managed Server.

  Refer to step 2 of "8.2.3 Troubleshooting: Failure in accessing a repository database" for details on checking the required area.

- The repository (directory for the DB space that stores repository data) may be too small when a Managed Server is added. Re-estimate the repository size before adding a Managed Server. Refer to "Estimation of database size" in the "ETERNUS SF Installation and Setup Guide" for details.
  If the estimation result is less than 65 megabytes, the size does not need to be increased. If it is more than 65 megabytes, extend the size before adding a Managed Server. Refer to "8.2.1 Troubleshooting: Insufficient free space in a repository" for details.

. . . .

3.6.4 Deleting the Managed Server

The process of deleting a Managed Server can be carried out using either Web Console or stgxfwcmdelsrv (Server information deletion command).

Point

Ensure that all of the steps are complete before removing the Managed Server that is to be deleted.

The required steps for deleting a Managed Server are described below.

1. Delete the backup history of all transaction volumes on the Managed Server to be deleted. For more information, refer to "3.5.4.2 Deleting a backup history".

2. Delete the backup policy of all transaction volumes on the Managed Server to be deleted.

   The process of deleting a backup policy can be carried out using either Web Console or swstbkpoldel (Backup policy deletion command).

3. Delete the volume type that was registered using "3.4.6 Setting the operation type for a device" for the device connected to the Managed Server.

4. Delete the Managed Server.

3.7 Stopping the Operation

To stop a backup operation, stop the daemons running on the Managed Server. Normally, they are automatically stopped when the system stops.

You can stop a specific daemon if you need to do so for some reason. For more information, refer to "Chapter 2 Starting and Stopping Daemons".

Note

- When all daemons on the Managed Server have stopped, all functions of AdvancedCopy Manager running on the Managed Server will stop.

- Stop daemons on a Management Server only after ensuring that all the Managed Servers that it manages have also been stopped.

3.8 SDX Object Operations

This section describes the backup and restore operations of SDX objects.
A good understanding of the basic operation of normal volumes is required for these operations.

For details on GDS, refer to the "PRIMECLUSTER(TM) Global Disk Services User's Guide (Linux)"

3.8.1 Units of SDX object operations

There are two types of backup operations (with different copy units for each type):

1. Backup operation in units of logical volumes (method with GDS Snapshot)

2. Backup operation in units of slices (method without GDS Snapshot)

For details on how to confirm and select the type of operation, refer to "3.8.1.1 Configuring the unit of SDX object operations"

Note

The two types of operation cannot be performed on the same Managed Server.

3.8.1.1 Configuring the unit of SDX object operations

The unit of SDX object operations is normally set as follows:

Units of Logical volumes

When PRIMECLUSTER GDS Snapshot has been included PRIMECLUSTER GDS 4.1A30 or later

Units of Slices

In cases other than the above

For operation in units of slices in an environment with PRIMECLUSTER GDS 4.1A30 or later installed, use the command described in stgxfwcmsetsdx (SDX object management unit configuration command), and then set operation in units of slices by taking action as described in "3.4.4 Fetching device information on a Managed Server".

🛈 Note
. . . .

- If PRIMECLUSTER GDS has already been upgraded from version 4.1A20 or earlier to 4.1A30 or later, follow the procedure below to change the setting to operation in units of logical volumes.

    1. Refer to "3.6.4 Deleting the Managed Server", and delete the Managed Server.

    2. Use the command described in stgxfwcmsetsdx (SDX object management unit configuration command), to specify management in units of logical volumes.

    3. Perform the steps as described in "3.4.3 Registering a Managed Server", and perform the subsequent operations.

        For operation in units of slices as before, the above procedure need not be followed.

- If AdvancedCopy Manager has been upgraded from version 10.3 or earlier in an environment that includes PRIMECLUSTER GDS 4.1A30 or later, the operation described in "3.4.4 Fetching device information on a Managed Server" results in a change to the unit of an SDX object operations will detect a change to a managed device.
  To keep the operation in units of slices, abort the fetching operation without accepting the detected change, and use the command described in stgxfwcmsetsdx (SDX object management unit configuration command) to set operation in units of slices, and re-specify the required values as necessary.

. . . .

3.8.2 Backup operation in units of logical volumes (method with GDS Snapshot)

Backup operation in units of logical volumes is possible in collaboration with GDS Snapshot.
The physical structure of the logical volumes is not important.

Figure 3.11 Backup operation in units of logical volumes

To specify a GDS logical volume in the AdvancedCopy Manager command, specify the GDS logical volume name.

```
/dev/sfdsk/<class-name>/dsk/<volume-name>
```

Note

Configuration and conditions of unsupported SDX objects

AdvancedCopy Manager does not support the following SDX objects (volumes):

- Volume of a shadow class

- Stripe volume

- Switch volume

- Volume in a concatenation group
  Stripe groups and concatenation groups, however, can be used as subordinate groups of a mirror group.

Copying cannot be processed in the following circumstances:

If the total number of the slices that make up the transaction volume plus all slices that make up all backup volumes associated with the transaction volume at the start of synchronous processing or snapshot-type backup processing is 33 or greater.

If either the source volume (the transaction volume for backup processing or the backup volume for restore processing) are in an invalid state.

Point

To confirm the type of volume being used, (that is: single, mirror, stripe, concatenation, or switch), execute "sdxinfo -e long" at the command line and check the volume type attribute values returned.

The value in the TYPE column on the line indicated for the volume in the OBJ column will match one of the following:

Volume object type attributes

- single : Single volume (supported)

- mirror : Mirror volume (supported)

- stripe : Stripe volume (not supported)

- concat : Volume in a concatenation group (not supported)

- switch : Switch volume (not supported)

・・・

3.8.2.1 Designing the backup operation

The following information about performing the backup operation should be noted.

For a backup performed in collaboration with GDS Snapshot, AdvancedCopy Manager processes binding, separation, and release of volumes by using the GDS Snapshot function.

Therefore, an SDX object used by AdvancedCopy Manager must be set to a state in which the object can operate as either a master object or proxy object.

A transaction volume operates as a master object, and a backup volume as a proxy object.

Configure the GDS volume structure accordingly before registering a volume with AdvancedCopy Manager.

When configuring SDX objects of GDS for use by AdvancedCopy Manager, pay careful attention to the relevant structure attributes of "3.8.2.2 Transaction volume" and "3.8.2.3 Backup volumes" described in the next sections.

For details on GDS Snapshot and premises of the proxy structure, refer to the "PRIMECLUSTER(TM) Global Disk Services User's Guide (Linux)".

3.8.2.2 Transaction volume

To use an SDX object (logical volume) for transactions, register it as a transaction volume.
The transaction volume operates as a master object of GDS Snapshot.

Therefore, note the following about configuring the backup operation:

- An SDX object that belongs to a shadow class cannot be registered as the transaction volume.

- The only types of logical volumes that can be registered as the transaction volume are single volumes and mirror volumes.

- A stripe volume, a volume in a concatenation group, or a switch volume cannot be registered as the transaction volume.
  Stripe groups and concatenation groups, however, can be used as subordinate groups of a mirror group.

- A volume in a disk group to which a backup volume also belongs cannot be registered as the transaction volume.

- A logical volume that belongs to the single disk group with a backup volume cannot be registered as the transaction volume.

- If any of the slices that make up the transaction volume are separated temporarily or are being copied, backup cannot be performed.

- If the transaction volume (logical volume) is in the invalid state, backup cannot be performed.

- A proxy object of GDS Snapshot cannot be used as the transaction volume.

3.8.2.3 Backup volumes

Register an SDX object (logical volume) as a backup volume.
The backup volume operates as a proxy object of GDS Snapshot.

Therefore, note the following about configuring the backup operation:

- An SDX object that belongs to a shadow class cannot be registered as a backup volume.

- The types of logical volumes that can be registered as backup volumes are single volumes and mirror volumes. A stripe volume, a volume in a concatenation group, or a switch volume cannot be registered as a backup volume.
  Stripe groups and concatenation groups, however, can be used as subordinate groups of a mirror group.

- A volume in a disk group to which a transaction volume also belongs cannot be registered as a backup volume.

- A logical volume that belongs to the single disk with the transaction volume cannot be registered as a backup volume.

- An object whose class is different from that of the transaction volume cannot be used as a backup volume.

- A master object of GDS Snapshot cannot be used as a backup volume.

- A backup volume for backup by using EC or OPC must be a non-mirroring volume. If the backup volume is in a mirroring structure, only software-copy can be used.

- If a backup volume is in use or is accessible by any application, backup cannot be performed.

- If any slices of the slices that make up a backup volume is separated temporarily or being copied, backup cannot be performed.

- As a backup volume, a logical volume must be the same size as the transaction volume.

- For the backup operation, the total number of the slices that make up the transaction volume and slices that make up all backup volumes associated with the transaction volume must not exceed 32.

3.8.2.4 Copy functions available when used with GDS Snapshot

The following copy functions can be used in linkage with GDS Snapshot:

- Advanced copy: Copying with the hardware function of ETERNUS Disk storage system

- Software-copy: Copying with the disk driver of PRIMECLUSTER GDS

The following table shows the relationship between SDX object structures and their available copy functions.

Table 3.4 Copy functions available in linkage with GDS Snapshot (backup)

Unit	Volume	SDX object structure(*1)		Available copy function
Logical volume	Transaction	Single		Advanced Copy / Software-copy
		Mirror	No subordinate group exists.	Advanced Copy / Software-copy
			A subordinate group exists.	Software-copy
	Backup	Single		Advanced Copy / Software-copy
		Mirror	This consists of one disk.	Advanced Copy / Software-copy
			This consists of multiple disks.	Software-copy

(*1) Stripe- and concatenation-type objects cannot be copied.

Table 3.5 Copy functions available in linkage with GDS Snapshot (restore)

Unit	Volume	SDX object structure(*1)		Available copy function
Logical volume	Backup	Single		Advanced Copy / Software-copy
		Mirror	No subordinate group exists.	Advanced Copy / Software-copy
			A subordinate group exists.	Software-copy
	Transaction	Single		Advanced Copy / Software-copy
		Mirror	This consists of one disk.	Advanced Copy / Software-copy
			This consists of multiple disks.	Advanced Copy / Software-copy

(*1) Stripe-, concatenation-, and switch-type objects cannot be copied.

3.8.2.5 Configuration options

Specify options when performing backup to a backup volume in a box that does not contain the transaction volume.
Backup with GDS Snapshot does not require a setting to indicate inter-box copying.

3.8.2.6 Fetching device information from the Managed Server

To perform the backup operation with a Managed Server and obtain information of a subordinate device related to the Managed Server, you must register the server.
For details on how to fetch the device information, refer to the section titled "3.4.4 Fetching device information on a Managed Server".

3.8.2.7 Configuring the operation type for a device

Register a GDS logical volume as a transaction volume or backup volume.

```
# /opt/FJSVswsts/bin/swstdevinfoset -t /dev/sfdsk/CLS1/dsk/vol1
swstdevinfoset completed
# /opt/FJSVswsts/bin/swstdevinfoset -b /dev/sfdsk/CLS1/dsk/bvol1
swstdevinfoset completed
```

📝 Note

. . . .

The following SDX objects cannot be registered as transaction volumes or backup volumes:

  - Object of a shadow class

  - Stripe-type object

  - Concatenation-type object

  - Switch-type object

. . . .

3.8.2.8 Configuring a backup policy

Checking the number of backup generations

Before configuring a backup policy, the required number of backup volumes for operation must be prepared according to the backup policy. (For details, refer to the section titled "3.4.7 Setting a backup policy".)

If the transaction volume is an SDX object, verify that there are enough backup volumes registered as proxy objects to match the required number of backup generations.

If the transaction volume is a GDS logical volume, each backup volume must be also a GDS logical volume.

```
#/opt/FJSVswsts/bin/swstbkpolset -s 3 /dev/sfdsk/CLS01/dsk/VOL01
/dev/sfdsk/CLS01/dsk/VOL01 swstbkpolset completed
#
```

📝 Note

. . . .

In the following cases, some SDX objects are not counted as available backup volumes:

  - Where the backup volume is an object whose class is different from that of the transaction volume

  - Where the size of the backup volume is different from that of the transaction volume

. . . .

Limit on the number of slices

Where an SDX object is to be backed up, the transaction volume is described as the master and its associated backup volume is described as the proxy. Therefore, the total number of the slices that make up the transaction volume plus the number of slices that make up the associated backup volume must not exceed 32. Consequently, the number of backup generations that can be used is limited according to the GDS structure. This will be important when configuring a backup policy.

For example, if both the transaction volume and backup volume are single volumes, you can set the number of backup generations to 31, which is the maximum value. However, if both the transaction volume and backup volume have a two-slice structure, the number of backup generations is limited to 15. If 16 or a higher value is set, the 16th generation and subsequent generations are not backed up.

3.8.2.9 Backup

Backup in units of logical volumes

If the transaction volume is an SDX object, select a backup volume that can be a valid proxy object for backup.

AdvancedCopy Manager automatically selects it from the registered backup volumes.

```
# /opt/FJSVswsts/bin/swstbackup /dev/sfdsk/CLS1/dsk/vol1
/dev/sfdsk/CLS1/dsk/vol1 swstbackup completed
```

Available copy processing methods

The available copy methods vary depending on the hardware environment or configuration supporting the GDS logical volumes.

EC/OPC cannot be used in the following cases:

- The backup volume is an SDX object in a mirroring structure

- A subordinate group is connected to the disk group to which the transaction volume and backup volume belong

Table 3.6 Hardware and the available copy processing methods

Hardware	Available copy processing method
EC/OPC can be used.	Synchronous high-speed backup: EC, software-copy
	Snapshot-type high-speed backup: OPC
Only OPC can be used.	Synchronous high-speed backup: Software-copy
	Snapshot-type high-speed backup: OPC
EC/OPC cannot be used.	Synchronous high-speed backup: Software-copy
	Snapshot-type high-speed backup: Unavailable

3.8.2.10 Restore

Restoration to a transaction volume in units of logical volumes

This process can be carried out using either the Web Console or a command.

For restoration to a transaction volume in units of GDS logical volumes, copy the proxy to the master.

```
# /opt/FJSVswsts/bin/swstrestore /dev/sfdsk/CLS1/dsk/vol1
/dev/sfdsk/CLS1/dsk/vol1 swstrestore completed
```

Restoration with the restore destination specified

If the transaction volume is a GDS logical volume, restoration cannot be performed with the restore destination specified.

Available copy processing methods

The available copy processing methods vary depending on the hardware environment or configuration supporting the GDS logical volumes.

OPC cannot be used in the following case:

- A subordinate group is connected to the disk group to which the transaction volume and backup volume belong.

Table 3.7 Hardware and the available copy processing methods

Hardware	Available copy processing method
OPC can be used.	OPC, software-copy
OPC cannot be used.	Software-copy

3.8.2.11 Operational notes

Do not use the "sdxproxy" command directly to change the state of a volume registered for backup management.
Use of the "sdxproxy" command may create inconsistencies between the backup management information and the actual volume state. In such cases, perform recovery by executing swstsrsemtch (Resource match command).

3.8.3 Backup operation in units of slices (method without GDS Snapshot)

This backup operation uses only the AdvancedCopy Manager function without GDS Snapshot.
For reliable configuration and operation, you need to know the type of physical volumes that make up the logical volume.

Figure 3.12 Backup operation in units of slices

When specifying a name in an AdvancedCopy Manager command, combine the logical volume name and the AdvancedCopy Manager device name using the format shown below. Refer to "1.5.2 SDX objects of PRIMECLUSTER GDS" for details.

```
/dev/sfdsk/<class-name>/dsk/<volume-name>:<DEVNAM value of sdxinfo>
```

🔔 Note

- For notes on SDX object operations, refer to "10.1.5.1 Notes on SDX object operations (backup management)", and "10.1.5.2 Notes on SDX object operations (replication management)".

- If mirroring between boxes is enabled and restoration by OPC is required in the event of a box failure, both mirrored systems must be backed up. In such cases, the required capacity of the backup volume is not the size of the logical volume but the size of the physical volume.

   If a Symfoware DBSP is being mirrored within the ETERNUS Disk storage system, only one of the paired devices (mirrored device or mirroring device) can be registered as a transaction volume. Therefore, if a hardware error occurs on the device that carries out backup, backup and recovery can only be implemented after restoring the device from the hardware error.

- The SDX objects that can be used for operations in slice units are as follows:

   - Single volumes

   - Non-layered mirror volumes

🅿 Point

You can find out if the volume being used is the single, mirror, stripe, concatenation, or switch volume type by executing "sdxinfo -e long". (Check the value in the TYPE column in the line displayed for the volume in the OBJ column.)

Volume object type attribute

   - single: Single volume (operation possible)

   - mirror: Mirror volume (operation possible)

   - stripe: Stripe volume (operation not possible)

   - concat: Volume in a concatenation group (operation not possible)

   - switch: Switch volume (operation not possible)

3.8.3.1 Configuring the backup operation

For notes on configuring the backup operation, refer to "10.1.5.1 Notes on SDX object operations (backup management)".

3.8.3.2 Fetching device information from the Managed Server

To perform the backup operation with a Managed Server and obtain information of a subordinate device of the Managed Server, you must register the server.
For details on how to fetch the device information, refer to "3.4.4 Fetching device information on a Managed Server".

3.8.3.3 Configuring the operation type for a device

Transaction volume

Register the slices that make up the logical volume used for transactions as a transaction volume.

```
#/opt/FJSVswsts/bin/swstdevinfoset -t /dev/sfdsk/CLS01/dsk/VOL01:sda
swstdevinfoset completed
#
```

Backup volume

Slices of an SDX object cannot be registered as a backup volume.
General slices must be used for a backup volume.

```
#/opt/FJSVswsts/bin/swstdevinfoset -b /dev/sdd6
swstdevinfoset completed
#
```

3.8.3.4 Backup

Backup example

```
#/opt/FJSVswsts/bin/swstbackup /dev/sfdsk/CLS01/dsk/VOL01:sda
/dev/sfdsk/CLS01/dsk/VOL01:sda swstbackup completed
#
```

Backup states

Backup can only be performed when the SDX object that makes up the applicable logical volume is in one of the states listed below.
If the object is in any other state, backup cannot be performed.
AdvancedCopy Manager checks the state of the SDX object by using the "sdxinfo" command of PRIMECLUSTER GDS.

1. The state of the volume is "ACTIVE" or "STOP".

2. The state of the physical disk is "ENABLE".

3. The state of the slice is "ACTIVE" or "TEMP".

Backup pre-processing and post-processing

In the following case, the pre-processing and post-processing scripts are not executed:

- A slice of an SDX object is in the TEMP state.

3.8.3.5 Restore

This process can be carried out using either the Web Console or a command.

For operation in units of slices, restoration can be performed only if the transaction volume is a single volume.

An SDX object that is not a single volume cannot be restored.

For details on restoration of an SDX object that is not a single volume, refer to "10.1.5.1 Notes on SDX object operations (backup management)".

Restoration example

```
#/opt/FJSVswsts/bin/swstrestore  /dev/sfdsk/CLS01/dsk/VOL01:sda
/dev/sfdsk/CLS01/dsk/VOL01:sda swstbackup completed
#
```

3.9 VxVM Volume Operation

This section describes the backup operation for volumes under VERITAS Volume Manager (VxVM).

3.9.1 Backup operation in units of physical disks

If a VxVM volume is the backup target, execute backup in the physical disk units that comprise the VxVM volume.

When backup operation is performed in units of physical disks, disk group integrity must be maintained and, therefore, all the physical disks in the disk group must be processed synchronously.

P Point

. . . .

For the physical volume which should be synchronized, confirm it by using either Web Console or stgxfwcmdispdev (Device information display command).

. . . .

Note

. . . .

- Before starting this operation, be sure to understand the basic operation of ordinary volumes.

- Snapshot backup is recommended for backing up in units of physical disks. In the case of synchronized backup, commands that access the destination volume, such as VxVM commands cannot be used during full copy or differential copy.

- For operation in a cluster configuration, the device name (/dev/sd#) of the physical disks that comprise the disk group must be the same at all servers that comprise the cluster, and the ETERNUS Disk storage system's disk indicated by the device name must be the same.

. . . .

3.9.1.1 Operational configuration

Configure disk groups that are to be used as transaction volumes or backup volumes.

Observe the following conditions when configuring the disk groups:

- The number, sizes, and types of VM disks must be the same.

- The configurations of logical volumes must be the same.

Figure 3.13 Transaction volume and backup volume
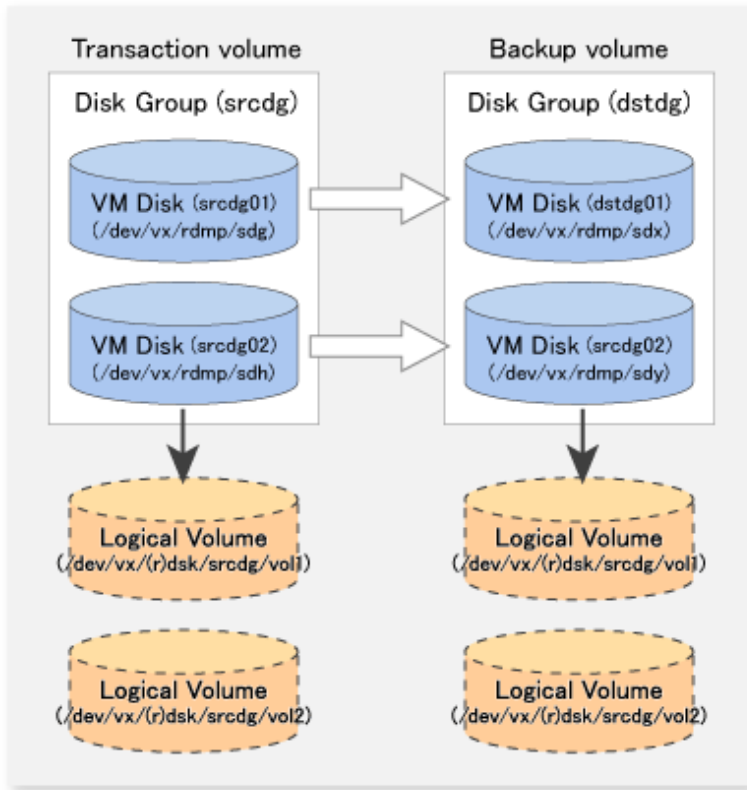
3.9.1.2 Preparations

3.9.1.2.1 Confirming the disk group configuration information file

The disk group must be reconfigured in backup post-processing. Confirm that a volume group configuration information file has been saved in the following format.

```
/etc/vx/cbr/bk/<disk group name>.<disk group ID>
```

3.9.1.2.2 Setting the transaction volume and backup volume

When configuring the transaction and backup volumes, all disks in the volume group must be specified.

Example:

```
# /opt/FJSVswsts/bin/swstdevinfoset -t /dev/vx/dmp/sdg
swstdevinfoset completed
# /opt/FJSVswsts/bin/swstdevinfoset -t /dev/vx/dmp/sdh
swstdevinfoset completed
# /opt/FJSVswsts/bin/swstdevinfoset -b /dev/vx/dmp/sdx
swstdevinfoset completed
# /opt/FJSVswsts/bin/swstdevinfoset -b /dev/vx/dmp/sdy
swstdevinfoset completed
#
```

3.9.1.2.3 Preparing a device map file

For the backup operation of a VxVM volume, a device map file must be created because a backup volume in the same volume structure as the transaction volume must be specified.

Example of a device map file

```
# Transaction volume        Backup volume
/dev/vx/dmp/sdg              /dev/vx/dmp/sdx
/dev/vx/dmp/sdh              /dev/vx/dmp/sdy
```

For details on the device map file, refer to "3.4.9 Preparing a device map file".

3.9.1.3 Backup

Before performing backup operation, all physical disks in the disk group must be synchronized.

Perform the required pre-processing and/or post-processing work for each volume group. Disable pre-processing and post-processing when operating individual physical disks.

Example of snapshot backup

```
(Perform pre-processing for the transaction and backup volumes.)

# /opt/FJSVswsts/bin/swstbackup /dev/vx/dmp/sdg -Xdevmap /acm/devmap.txt
/dev/vx/dmp/sdg swstbackup completed
# /opt/FJSVswsts/bin/swstbackup /dev/vx/dmp/sdh -Xdevmap /acm/devmap.txt
/dev/vx/dmp/sdh swstbackup completed
#

(Perform post-processing for the transaction and backup volumes.)
```

Example of synchronous backup

```
(Perform pre-processing for the backup volumes.)

# /opt/FJSVswsts/bin/swststartsync /dev/vx/dmp/sdg -Xdevmap /acm/devmap.txt
/dev/vx/dmp/sdg swststartsync completed
# /opt/FJSVswsts/bin/swststartsync /dev/vx/dmp/sdh -Xdevmap /acm/devmap.txt
/dev/vx/dmp/sdh swstsstartsync completed
#

(After state of equivalency upkeep)
(Perform pre-processing for the transaction volumes.)

# /opt/FJSVswsts/bin/swstbackup /dev/vx/dmp/sdg
/dev/vx/dmp/sdg swstbackup completed
# /opt/FJSVswsts/bin/swstbackup /dev/vx/dmp/sdh
/dev/vx/dmp/sdh swstbackup completed
#

(Perform post-processing for the transaction and backup volumes.)
```

The table below summarizes the pre-processing and post-processing work to be performed before and after backup.

Table 3.8 Pre-processing and post-processing for backup

Volume type	Pre-processing	Post-processing
Transaction volume	1. Secure data integrity by stopping access to all logical volumes in the disk group. 2. If file systems are included, unmount all file systems in the disk group.	1. If file systems are included, mount the volumes that were unmounted during preprocessing.

Volume type	Pre-processing	Post-processing
	3. Import the disk group, when the disk group is not imported.	
Backup volume	1. Stop access to all logical volumes in the disk group. 2. If file systems are included, unmount all file systems in the disk group. 3. Deport the disk group. 4. A disk group subordinate's physical disk is set to offline.	1. The physical disk set to offline with preprocessing is set to online. 2. Reconfigure the disk group. 3. If file systems are included, remount the volumes that were unmounted during preprocessing.

Reconfiguring the disk group

Reconfigure the disk group as follows:

1. Pre-commit analysis for restoration

```
# /etc/vx/bin/vxconfigrestore -p dstdg
Diskgroup dstdg configuration restoration started ......

Installing volume manager disk header for sdx ...
Installing volume manager disk header for sdy ...
-
dstdg's diskgroup configuration is restored (in precommit state).
Diskgroup can be accessed in read only and can be examined using
vxprint in this state.

Run:
  vxconfigrestore -c dstdg ==> to commit the restoration.
  vxconfigrestore -d dstdg ==> to abort the restoration.
#
```

2. Commit the change required for restoring the configuration of the copy destination disk group.

```
# /etc/vx/bin/vxconfigrestore -c dstdg
Committing configuration restoration for diskgroup dstdg ....

dstdg's diskgroup configuration restoration is committed.
#
```

🖝 **Note**

- In the case of a clustered system, when a disk group or a mount resource has been defined, instead of using the import/deport command for the disk group use the online/offline process.
  If a mount point is defined as a cluster resource, instead of using the file system mount/unmount commands use the mount resource online/offline processing.

- When performing a system disk exchange, there are cases when conflicting backup configuration information may exist.
  In such cases, the disk group ID needs to be reset after executing the above command.

- After this operation, where the volumes within a disk group are required to be run in synchronous mode in background, synchronous processing may take some time depending on the volume configuration.
  It is also possible to use the volumes during this time.

3.9.1.4 Restoration

All physical disks in the disk group must first be synchronized to perform this operation.

Perform the required pre-processing or post-processing work for each disk group as necessary. Disable pre-processing and post-processing when using individual physical disks.

Example of restoration

```
(Perform pre-processing for the transaction and backup volumes.)

# /opt/FJSVswsts/bin/swstrestore /dev/vx/dmp/sdg
/dev/vx/dmp/sdg swstrestore completed
# /opt/FJSVswsts/bin/swstrestore /dev/vx/dmp/sdh
/dev/vx/dmp/sdh swstrestore completed
#

(Perform post-processing for the transaction and backup volumes.)
```

The table below summarizes the pre-processing and post-processing work to be performed before and after restoration.

Table 3.9 Pre-processing and post-processing for restoration

Volume type	Pre-processing	Post-processing
Backup volume	1. Secure data integrity by stopping access to all logical volumes in the disk group. 2. Import the disk group, when the disk group is not imported.	Post-processing is not required.
Restoration destination volume	1. Stop access to all logical volumes in the disk group. 2. If file systems are included, unmount all file systems in the disk group. 3. Deport the disk group. 4. A disk group subordinate's physical disk is set to offline.	1. The physical disk set to offline with pre-processing is set to online. 2. Reconfigure the disk group. 3. If file systems are included, remount the volumes that were unmounted by preprocessing.

Reconfiguring the disk group

Reconfigure the disk group as follows:

1. Restoration pre-commit analysis

```
# /etc/vx/bin/vxconfigrestore -p srcdg
Diskgroup srcdg configuration restoration started ......

Installing volume manager disk header for sdg ...
Installing volume manager disk header for sdh ...
-
srcdg's diskgroup configuration is restored (in precommit state).
Diskgroup can be accessed in read only and can be examined using
vxprint in this state.

Run:
  vxconfigrestore -c srcdg ==> to commit the restoration.
  vxconfigrestore -d srcdg ==> to abort the restoration.
#
```

2. Commit the change required for restoring the configuration of the copy destination disk group.

```
# /etc/vx/bin/vxconfigrestore -c srcdg
Committing configuration restoration for diskgroup srcdg ....

srcdg's diskgroup configuration restoration is committed.
#
```

Note

- In case of a clustered system, when a disk group or a mount resource has been defined, instead of using the import/deport command for the disk group use the online/offline process.
  If a mount point has been defined as a clustered system resource, instead of using the file system mount/unmount commands use the mount resource online/offline processing.

- When performing a system disk exchange, there are cases when conflicting backup configuration information may exist.
  In such cases, the disk group ID needs to be reset after executing the above command.

- After this operation, where the volumes within a disk group are required to be run in synchronous mode in background, synchronous processing it may take some time depending on the volume configuration.
  It is also possible to use the volumes during this time.

Chapter 4 Backup Operations (Symfoware)

This chapter describes how to backup the Symfoware database.

4.1 Backup Operations for Symfoware

This section explains operations for backing up Symfoware.

📝 Note
. . . .

- Backup operations for Symfoware are possible only for Managed Servers running Linux or Windows operating systems.

- Only Symfoware database spaces that are allocated to raw devices are available for backup by AdvancedCopy Manager.

- If you are using the Symfoware linkage function under Red Hat Enterprise Linux 5 (for Intel64), AdvancedCopy Manager's agents must be in the 32-bit version.

  Use the following procedure to check that the AdvancedCopy Manager's agents are the 32-bit version.

  1. Execute the file command.

     ```
     # file /opt/FJSVswstf/bin/stgenvdisp
     ```

  2. Check that the result of executing the command contains "ELF 32-bit".

     Example: Execution result when the AdvancedCopy Manager's agents are the 32-bit version

     ```
     # file /opt/FJSVswstf/bin/stgenvdisp
     /opt/FJSVswstf/bin/stgenvdisp: ELF 32-bit LSB executable, Intel 80386, version 1
     (SYSV), for GNU/Linux 2.6.9, dynamically linked (uses shared libs), for GNU/Linux 2.6.9,
     not stripped
     ```

  If the installed AdvancedCopy Manager's agents are not the 32-bit version, uninstall them. After that, refer to "Installation procedure of AdvancedCopy Manager's agent (for Linux)" in the "ETERNUS SF Installation and Setup Guide" for installing the 32-bit version of AdvancedCopy Manager's agents.

. . . .

4.1.1 Overview

AdvancedCopy Manager handles the Symfoware database spaces that are located on the disk volumes in the ETERNUS Disk storage system as transaction volumes. For this reason, AdvancedCopy Manager allows for backup and recovery in physically allocated structures called "database spaces", without you having to be aware of the storage structure for the Symfoware database spaces.

By creating database spaces and log groups allocated on Managed Servers for various transactional purposes, Symfoware can create multiple operating environments. This operating environment can be distinguished by the name "RDB system name". Based on this RDB system name, AdvancedCopy Manager implements backup and recovery without damaging the consistency of relations among database tables in the operating environment.

For a complete overview of Symfoware, refer to the "Symfoware Server RDB Operations Guide".

Figure 4.1 Function Overview

4.1.1.1 Backup

With AdvancedCopy Manager, you can back up Symfoware database spaces in units of either database spaces or log groups.

AdvancedCopy Manager handles the partitions to which the database spaces to be backed up are allocated as transaction volumes.

When backing up in units of database spaces, these transaction volumes are backed up to backup volumes.

When backing up in units of log groups, all database spaces included in a log group are defined as separate transaction volumes, and all transaction volumes are backed up.

Note

For backing up in units of log groups, all partitions to which database spaces included in a log group are allocated must be registered as transaction volumes beforehand. If a single partition remains unregistered, AdvancedCopy Manager cannot back up the unregistered database spaces and, as a consequence, not maintain consistency of relations among tables when recovering the database.

During backup, AdvancedCopy Manager creates a recovery control file which contains the data required for recovery of the database and is managed separately for each backup generation.

4.1.1.2 Recovery

With AdvancedCopy Manager, you can recover Symfoware database spaces in units of either database spaces or log groups.

The recovery process with AdvancedCopy Manager is as follows:

1. Select the appropriate recovery control files and backup volumes.

2. Replicate the backup volumes on the transaction volumes.

3. For restoring the latest condition or a certain point in time, apply the archive log based on the recovery control file.

When the archive log file has become full with archive logs, there may be no more response from Symfoware commands or from user applications that are using the database space. If this happens, you can backup archive logs on external media. AdvancedCopy Manager can implement recovery if you specify the name of the backup file for the archive log that is stored on an external medium (file that includes the name of the archive log's backup location).

Note

Recovery cannot be implemented if the specified "file that includes the name of the archive log's backup location" contains an archive backup file that is backed up on a MT (tape device).

Also, at the time of recovery, the working directory specified in swstsvrset (Managed Server configuration information setting command) will be used.

 P Point

Since you must use AdvancedCopy Manager to recover Symfoware database spaces that were backed up with AdvancedCopy Manager, you cannot directly recover from any backup volumes that were backed up on tape or other secondary media.

Backup Operation for Archive Logs

For backing up archive logs, it is necessary to understand the handling of the backup archive logs that are required for the recovery of database spaces.

AdvancedCopy Manager manages archive log backup processing that is implemented in the required units of log groups as well as the handling of backup and recovery that is implemented in units of database spaces. A recommended procedure is that the database administrator record the backup status of the archive logs. Further, Fujitsu recommends that you create shell scripts that define the backup work for archive logs, and to record the logs in a file on a Managed Server.

Figure 4.2 Backup operation for archive logs
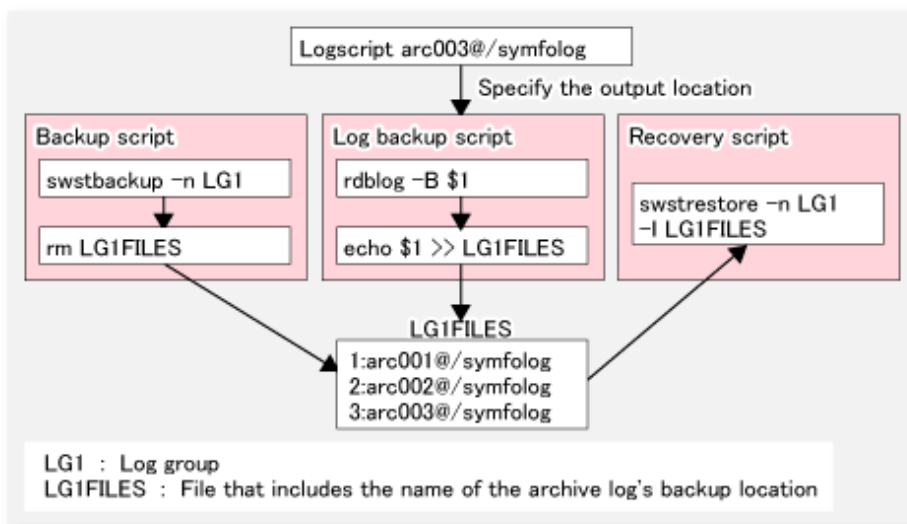
4.1.2 Operation Design

When configuring a backup system for Symfoware, you will need to consider the following issues:

1. "4.1.2.1 Design of backup operations for Symfoware databases"

2. "4.1.2.2 Selecting a server on which backups are performed"

 Note

. . . .

For creating an environment in which AdvancedCopy Manager can implement backup and recovery of Symfoware database spaces in units of either database spaces or log groups, the database administrator must, in advance, create an operating environment for Symfoware that is suitable for the transaction purposes.

. . . .

4.1.2.1 Design of backup operations for Symfoware databases

Refer to the Symfoware Server manuals when you configure a backup system. However, in designing it, you will have to observe the following restrictions regarding the allocation of database space.

Database space on a raw device

For allocating database space on a raw device, since AdvancedCopy Manager's unit of copying (volumes) matches the area of the database space, there are no particular points to take into consideration.

For information on how to create database space on a raw device, refer to the "Symfoware Server RDB Operations Guide (for Database Definitions)".

Only configurations that allocate database space on raw devices are supported. Configurations that allocate database space in file systems are not supported.

4.1.2.2 Selecting a server on which backups are performed

Select a server on which to perform backups.

The backup operation can be performed on one of the following servers.

 - Management Server

   Multiple Managed Servers are unified and controlled centrally. The Management Server can also operate as a Managed Server.

 - Managed Server

   AdvancedCopy Manager operations can be performed on the Managed Server.

 Note

. . . .

The locale of the internal code system (the code system specified when AdvancedCopy Manager is installed on the Management Server) used by AdvancedCopy Manager must be installed on all Managed Servers.

. . . .

4.1.2.3 Selecting a backup source

Select the partition to which the Symfoware database spaces are allocated as the backup source.

In AdvancedCopy Manager, the partitions to which the Symfoware database spaces to be backed up are allocated are called "transaction volumes".

....

- For backing up in units of log groups, all partitions to which database spaces included in a log group are allocated must be registered as transaction volumes beforehand. If a single partition remains unregistered, AdvancedCopy Manager cannot back up the unregistered database spaces and, as a consequence, not maintain consistency of relations among tables when recovering the database.

- Do not designate partitions in which the system is stored or partitions on which AdvancedCopy Manager is installed as backup sources.

- Likewise, do not designate logical volumes and volume groups of LVM (Logical Volume Manager) as backup sources.

....

4.1.2.4 Configuring a backup policy

Select a backup policy to be configured for the transaction volumes or log groups. A backup policy consists of below items.

If you set a backup policy for a log group, the same backup policy will be set for all transaction volumes that are included in that log group.

Number of preservation generations

The number of preservation generations refers to the maximum number of generations of backup data that should be kept.

The snapshot high-speed backup first releases the oldest generation from generation management. If a system failure or any other failure occurs while the backup is running, the oldest generation may have already been released and there may not be as many generations of backup data as required. Therefore, Fujitsu recommends that you re-execute backup immediately to create the required number of generations of backup data.

When performing snapshot type high-speed backup using only one preservation generation, we recommend you also back up data to a tape device.

A synchronized high-speed backup operation releases the oldest generation from generation management after backup of the latest generation is completed. Therefore, the required number of backup volumes is equal to the maximum number of backup generations + 1.

Interval days

This item signifies a reference number of days after which a warning that backup is overdue is displayed.

If a period longer than the specified interval has passed since the last backup date, the delay will be displayed on the transaction list screen and by running the swstexecstat (Execution status display command).

Backups will not be performed automatically, even if you have specified the interval days.

4.1.2.5 Preparing a backup volume

This section describes how to prepare the backup volume.

In AdvancedCopy Manager, the partitions to which transaction volumes are backed up are called "backup volumes". Backup volumes must be located on the ETERNUS Disk storage system.

A backup volume must have exactly the same byte-size as a transaction volume. The number of backup volumes is determined by the type of backup operation as shown below:

Table 4.1 Number of backup volumes required for backup operations

Type of operation	Number of required backup volumes
Snapshot high-speed backup	Number of backup generations
Synchronized high-speed backup	Number of backup generations + 1

When backup policies have already been set for multiple transaction volumes, the registration of a backup policy for a newly registered transaction volume requires that the following number of backup volumes be registered beforehand:

Table 4.2 Number of backup volumes required for backup operations (if a backup policy has been set)

Type of operation	Number of required backup volumes
Snapshot high-speed backup	Total number of backup generations in the registered backup policies + number of generations in the new backup policy being set

Type of operation	Number of required backup volumes
Synchronized high-speed backup	Total number of backup generations in the registered backup policies + number of registered transaction volumes + number of generations in the new backup policy being set + 1

If you specify a log group for setting a backup policy, the following number of backup volumes is required:

Table 4.3 Number of backup volumes required for backup operations (if a log group is specified)

Type of operation	Number of required backup volumes
Snapshot high-speed backup	(Number of transaction volumes in the log group) * (number of generations in the backup policy being set)
Synchronized high-speed backup	(Number of transaction volumes in the log group) * (number of generations in the backup policy being set +1)

Note

. . . .

- Do not designate partitions in which the system is stored or partitions on which AdvancedCopy Manager is installed as backup volumes.

- Do not designate logical volumes and volume groups of LVM (Logical Volume Manager) as backup volumes.

. . . .

4.1.2.6 Preparing the recovery control file and the working file

For backup operation for Symfoware, the following directories need to be set up:

- Target directory for output of the recovery control file

- Working directory

Target directory for output of the recovery control file

The target directory for output of the recovery control file is the storage destination for the recovery control file that is created during backup.

The required capacity for the recovery control file is approximately 1 MB per backup generation of each transaction volume.

For M backup generations of N transaction volumes, the required capacity is as follows:

Table 4.4 Required capacity for the recovery control file

Type of operation	Required capacity (unit: MB)
Snapshot high-speed backup	N * M
Synchronized high-speed backup	N *(M+1)

If, at the time of starting a backup operation, no target directory for output of the recovery control file is set up, the following directory will be used.

- In case of normal operation (non-clustered environment)

```
/etc/opt/FJSVswsts/SymfoWARE/Recovery
```

- In case of operation in clustered environment

```
/etc/opt/FJSVswsts/<logical-node-name>(Note 1)/SymfoWARE/Recovery
```

(Note 1)Logical node name of AdvancedCopy Manager specified during cluster setup

Working directory

The working directory serves to carry out database recovery work when performing a restore.

If, at the time of starting a backup operation, no working directory is set up, the following directory will be used.

```
/var/opt/FJSVswsts/SymfoWARE
```

4.1.3 Flow of Operation

The flow of backup operations in Symfoware is shown below:

Figure 4.3 Flow of backup operations (Symfoware)

4.1.4 Preparing to Start the Backup

This section describes the preparations that are required for backup operation.

4.1.4.1 Activating daemons

In order to carry out backup operations, the daemons of AdvancedCopy Manager must be started on both the Management Server and the Managed Servers. Normally, these daemons are automatically started up when the system is started.

If the startup fails or a daemon has been stopped, start the daemon manually, referring to "Chapter 2 Starting and Stopping Daemons".

4.1.4.2 Starting the Web Console

For carrying out backup operations from the backup management window, start the Web Console, referring to "How to start Web Console" in the "ETERNUS SF Web Console Guide".

4.1.4.3 Registering a Managed Server

On the Management Server, register the Managed Servers that are to be managed.

This process is carried out using the Web Console.

4.1.4.4 Fetching device information on a Managed Server

This process can be executed using either Web Console or the stgxfwcmsetdev (Device information collection/reflection command).

When the process is finished, use the Web Console or the displayed result for the stgxfwcmdispdev (Device information display command) to confirm that the Symfoware information is displayed on the corresponding device. You can also use the stgxfwcmsetdev (Device information collection/reflection command) for confirmation.

📠 Note

⋯⋯

- For fetching the configuration information for a Symfoware database, start the Symfoware database before you fetch the device information from the Managed Server.

- Device information on the Managed Server must briefly be stored in the repository.

- The area required when updating the repository may not be available when information concerning devices under multiple Managed Servers is fetched. Therefore, check if there is sufficient area for updating the repository before fetching the device information. If there is insufficient area, allocate the required area before fetching the device information. Refer to Step 2 in "8.2.3 Troubleshooting: Failure in accessing a repository database" for details on checking the required area.

- In order to be able to fetch Symfoware information, Symfoware Server Advanced Backup Controller must be running. If you have installed Symfoware Server Advanced Backup Controller after fetching the device information, you must fetch it again.
  Also, depending on the system environment on which it is installed, the software's version level may be different. For details, refer to the related software manuals.

- The time required for this operation depends on the total number of devices defined on the selected Managed Server. If the number of devices is large, perform this operation while the CPU load and I/O load are both low. As a guideline, each operation takes about 0.5 seconds per device (or partition) under no-load conditions.

⋯⋯

4.1.4.5 Configuring the environment for a backup operation server

This process can be executed using either Web Console or the swstsvrset (Managed Server configuration information setting command).

When performing backup using commands only, the Managed Server configuration information setting command must be executed at the time that operations start.

🅿 Point

⋯⋯

If the Management Server also works as a Managed Server, these environment settings are also required for the Management Server.

If the Managed Server environment setup has already been carried out, this process is not required.

⋯⋯

4.1.4.6 Setting the operation type for a device

Define the partition on which the Symfoware database space to be backed up is configured as a transaction volume, and the target volume as a backup volume.

Settings for transaction volumes

Define the partition on which the Symfoware database space to be backed up is configured as a transaction volume by running the swstdevinfoset (Device information setting command).

For log groups, all partitions to which database spaces included in a log group are allocated must be registered separately as transaction volumes. If a single partition remains unregistered, AdvancedCopy Manager cannot back up the unregistered database spaces and, as a consequence, not maintain consistency of relations among tables when recovering the database.

Devices on which no Symfoware database space is set up cannot be set as transaction volumes for Symfoware.

📝 **Note**

....

- If you have modified any RDB system names, database space names or log group names that are assigned to transaction volumes for Symfoware, follow the procedure below to register them as transaction volumes:

    1. Run the History information deletion command to delete the backup history information of all registered transaction volumes.

    2. Delete the backup policies of all registered transaction volumes.

    3. Run the Device information setting command to unregister them as transaction volumes.

    4. Redo the procedure described in "4.1.4.4 Fetching device information on a Managed Server".

    5. Run the Device information setting command to re-register devices as transaction volumes.

- For unregistering transaction volumes that are included in a log group, delete all backup policies and backup history information for these transaction volumes before you start to unregister them.

....

Settings for backup volumes

Run the swstdevinfoset (Device information setting command) to configure a backup volume as a backup target. If backup volumes are already registered, this operation is not required.

For the number of backup volumes required by backup management, refer to "4.1.2.5 Preparing a backup volume".

📝 **Note**

....

- Before you modify the configuration of partitions that are registered as backup volumes, unregister them as backup volumes, and after changing the configuration, redo the procedure described in "4.1.4.4 Fetching device information on a Managed Server" before you register them by running the swstdevinfoset (Device information setting command).

....

4.1.4.7 Setting a backup policy

Run the swstbkpolset (Backup policy setting command) to set backup policies for the transaction volumes or log groups.

If you set a backup policy for a log group, the same backup policy will be set for all transaction volumes that are included in that log group.

For details on backup policies, refer to "4.1.2.4 Configuring a backup policy".

📝 **Note**

....

When you set backup policies, the backup volumes required for backup operations must already be registered. For the number of backup volumes required for backup operation, refer to "4.1.2.5 Preparing a backup volume".

....

You can display the set backup policies by running the swstbkpoldisp (Backup policy display command).

If you add a database space to a log group for which a backup policy is specified and set it up as a transaction volume, you have to redo the backup policy settings for the entire log group.

4.1.4.8 Preparing a device map file

From the backup of a group of volumes that are registered as backup volumes, AdvancedCopy Manager automatically selects a volume for use as a backup target volume which has the same capacity as the transaction volume.

However, when it is necessary to recognize a specific backup target volume for any operational reasons, a "device map file" that defines the relationship between a transaction volume and a backup volume must be created in advance.

Create the device map file in any suitable location on the Managed Server that carries out backup. Specifying this file when you start the backup enables you to recognize the backup targets during operation.

🅿 Point

•••

For managing multiple backup generations, you need multiple device map files.
Device map files to be used for starting backup or synchronization must fulfill one of the following conditions:

- A backup volume that is not yet in use must be specified.

- A backup volume must be specified that is used in the history which is going to be deleted in the backup process.

Therefore, if you operate with multiple backup volumes, you need different device map files for the different backup volume situations.

•••

4.1.4.8.1 Describing a device map file

The following provides an example of a device map file:

```
# A comment line starts with "#".
# Business volume name   Output destination backup volume name

# Normal Disks
/dev/sdc1           /dev/sdv1
/dev/sdc2           /dev/sdv2
/dev/sdc3           /dev/sdv3

# LOG01/GRP1
/dev/sdd1           /dev/sdz1
/dev/sdd2           /dev/sdz2
/dev/sdd3           /dev/sdz3

# LOG02/GRP2
/dev/sdg4           /dev/sdl4
/dev/sdg5           /dev/sdl5
:
:
```

The rules for creating a device map file are as follows:

- Describe a transaction volume name and a corresponding output target backup volume name in one line. Separate the transaction volume name and backup volume name with one or more space or tab characters. One or more space or tab characters may be used as delimiters between the beginning of a line and a transaction volume name or between the end of a target backup volume name and the end of a line (new line character).

- Blank lines (spaces or tab characters) may be included in the file.

- Anything after the symbol '#' until the end of the line is regarded as a comment.

- You cannot specify more than one backup volume for the same transaction volume in one device map file. If you do, only the first occurrence is valid. While the device map file will still be read, subsequent occurrences will be ignored.

- A device map file may include descriptions of transaction volumes other than the processing target (Device-Name or Log-Group-Name). Refer to the description example provided in the beginning of this chapter.

- If you define a backup volume as the output target for a log group, the descriptions for all transaction volumes must be included in the same single file. The descriptions must not straddle across multiple lines.

4.1.5 Operation

This section explains operations for backing up Symfoware.

Before performing the backup operation, configure the environment settings that are required for the backup according to "4.1.4 Preparing to Start the Backup".

There are two kinds of backup in AdvancedCopy Manager:

- Snapshot high-speed backup

- Synchronized high-speed backup

4.1.5.1 Performing snapshot high-speed backup

For performing a snapshot high-speed backup of Symfoware, specify the transaction volumes or log groups and run the swstbackup (Backup execution command). If you specify backup for a log group, backup will be executed at once for all transaction volumes that are included in that log group.

The following backup methods are available:

Table 4.5 Types of snapshot high-speed backup

Backup method	Description
Normal dump	Databases are backed up while archive log operation (i.e. normal operation) is in progress.
Reference dump	For purposes such as long-term preservation, databases are backed up after detaching them from archive log operation.

Note

- If you specify backup for transaction volumes, you cannot specify reference dump for transaction volumes that include multi-database spaces (physical structures in which a single table or index is assigned to multiple database spaces).

- When you back up with reference dump, you have to execute Symfoware's rdbrtr command in order to set inhibition of updates (disable writing of data) for all DSI (instance that expresses the storage structures for tables) on the transaction volumes to be backed up.

- From the Web Console, you cannot specify backup of individual transaction volumes.

4.1.5.2 Performing synchronized high-speed backup

For performing a synchronized high-speed backup of Symfoware, follow the procedure below:

1. Use the swststartsync (Backup synchronous processing start command) to specify the transaction volumes or log groups and start backup synchronous processing. If you specify a log group, backup synchronous processing will be executed for all transaction volumes that are included in that log group.
   After starting backup synchronous processing, you can cancel it with the swstcancelsync (Backup synchronous processing cancel command).

2. Run the swstsyncstat (Backup synchronous processing progress display command) to check that the transaction volumes are in the equivalency maintenance state while backup synchronous processing is in progress. If you have specified backup synchronous processing for a log group, check that all transaction volumes in the log group are in the equivalency maintenance state.

3. Specify the transaction volumes or log groups, and run the swstbackup (Backup execution command) to collect the backup. With the swstbackup (Backup execution command), you can use the backup methods shown below. On executing the swstbackup (Backup execution command), the current status at that time is registered in the backup history information, and backup synchronous processing is stopped.

Table 4.6 Types of synchronized high-speed backup

Backup method	Description
Normal dump	Databases are backed up while archive log operation (i.e. normal operation) is in progress.

Backup method	Description
Reference dump	For purposes such as long-term preservation, databases are backed up after detaching them from archive log operation.

Note

- If you specify backup for transaction volumes, you cannot specify reference dump for transaction volumes that include multi-database spaces (physical structures in which a single table or index is assigned to multiple database spaces).

- When you back up with reference dump, you have to execute Symfoware's rdbrtr command in order to set inhibition of updates (disable writing of data) for all DSI (function that expresses the storage structures for tables) on the transaction volumes to be backed up.

- If you run the swstbackup (Backup execution command) before the transaction and backup volumes are in an equivalent state, an error will occur.

- From the Web Console, you cannot specify backup of individual transaction volumes.

Point

With synchronized high-speed backup, backups are created by stopping or suspending synchronous processing. If you specify to suspend when you run the swstbackup (Backup execution command), the Suspend/Resume function performs the backup after suspending synchronous processing. For details on the Suspend/Resume function, refer to "Chapter 5 Backup Operation using Suspend/Resume Function".

4.1.5.3 Performing the recovery

For recovering the data backed up by snapshot high-speed backup or synchronized high-speed backup, execute the swstrestore (Restore execution command).

For recovery of Symfoware, from the backup volume where the history of backups by the swstbackup (Backup execution command) is managed, specify the transaction volumes or log groups when you run the swstrestore (Restore execution command).

Note

If you specify recovery for a log group, recovery will be executed at once for all transaction volumes that are included in that log group.

You can also use the -bundle option for batch recovery of multiple transaction volumes that are included in the same log group. This is called "bundle recovery".

If multiple database spaces are allocated in a RAID group, you can recover these database spaces in a single batch, which reduces the time required for applying logs and thereby the entire recovery time.

The following recovery methods are available:

- Recovery to the latest status

- Restoring of a specific point in time by specifying the end point of recovery (specifying the -e option when executing swstrestore (Restore execution command))

- Restoring of the point in time of backup (specifying the -a option when executing swstrestore (Restore execution command))

The available recovery methods depend on the backup method (normal dump or reference dump) that was used for the data to be recovered. You can specify the following combinations:

Table 4.7 Combinations of backup and recovery methods

Backup method	Recovery unit	Recovery method		
		Recovery to the latest status	Restoring of a specific point in time by specifying the end point of recovery (Note)	Restoring of the point in time of backup
Normal dump	Log group	A	A	A
	Multiple transaction volumes in the same log group	A	A	A
Reference dump	Transaction volume	A	N/A	N/A
	Log group	A	N/A	A
	Multiple transaction volumes in the same log group	A	N/A	A
	Transaction volume	A	N/A	N/A

A: available; N/A: not available

Note: For details on recovery points, refer to the Symfoware Server manuals. You have to memorize the recovery points until you carry out recovery of a database.

📝 **Note**

. . . .

- All of the described methods require that access to the database spaces to be recovered is not permitted. In order to prohibit access, execute the rdbinh or the rdbexspc command provided by Symfoware. For details on these commands, refer to the "Symfoware Server RDB Operations Guide".

- For individual recovery of transaction volumes in a log group, the only available method is recovery to the latest status. In such a case, it is required that the database administrator takes charge of matching the relations between tables within this log group.

- For performing "recovery to the latest status" or "restoring of a specific point in time", when the archive log files are stored on external media, you have to create a file that contains the backup file names required during recovery on the Managed Server that includes the transaction volume to be recovered, and specify this file name when you start recovery. For information on how to describe this file, refer to the "Symfoware Server RDB Operations Guide".

- If recovery started by the swstrestore (Restore execution command) fails due to insufficient space in the working directory, you can use the -w option to retry recovery with a different directory serving as a temporary working directory. For details on the -w option, refer to swstrestore (Restore execution command).

- If you perform "restoring of a specific point in time by specifying the end point of recovery" or "restoring of the point in time of backup", processing to restore Symfoware's management information will also be implemented. Since this processing is implemented as a portion of recovery processing, it takes time for commands to be processed.

- If the recovered data were originally backed up mixing both methods (normal dump and reference dump) within a log group, recovery cannot be carried out at once for the entire log group. In such a case, specify the -bundle option in the swstrestore (Restore execution command) in order to recover multiple transaction volumes in the same log group in a batch.

- If you originally backed up in units of database spaces, for recovery in units of log groups or for bundle recovery, you have to specify the relative generation number in the generation specification. This is necessary because, as shown in below example of a backup that was implemented in units of database spaces, the situation may occur that, after updating the history of a specific transaction volume (database space), absolute generation numbers are not complete for the corresponding relative generation numbers.

. . . .

Example: For a log group (LOG1/RDB1) containing Database space 1 (DB1.DBSP1) and Database space 2 (DB1.DBSP2)

Day 1: Backup of Database space 1 (DB1.DBSP1)

```
# /opt/FJSVswsts/bin/swstbackup /dev/sdd1
/dev/sdd1 swstbackup completed
```

Day 2: Separate backup of Database space 1 (DB1.DBSP1) and Database space 2 (DB1.DBSP2)

```
# /opt/FJSVswsts/bin/swstbackup /dev/sdd1
/dev/sdd1 swstbackup completed
# /opt/FJSVswsts/bin/swstbackup /dev/sdd2
/dev/sdd2 swstbackup completed
```

Display of history information

```
# /opt/FJSVswsts/bin/swsthistdisp -n LOG1/RDB1
Server=SV01 Device=/dev/sdd1 Mount-Point=DB1.DBSP1/LOG1/RDB1 (SymfoWARE)
Generation  Version   Backup-Date         Backup-Disk   Status     Execute   ArcSerial
    1          2       2003/12/12 22:00    /dev/sdz1     succeeded   ----        5
    2          1       2003/12/11 22:00    /dev/sdy1     succeeded   ----        5
Server=SV01 Device=/dev/sdd2 Mount-Point=DB1.DBSP2/LOG1/RDB1 (SymfoWARE)
Generation  Version   Backup-Date         Backup-Disk   Status     Execute   ArcSerial
    1          1       2003/12/12 23:00    /dev/sdz2     succeeded   ----        5
```

For recovery in units of log groups or for bundle recovery using a history as shown above, depending on the specified option, different backup data are used. In the case described above, since, by specifying the relative generation numbers, both database spaces use the backup data of Day 2, consistency of databases after recovery can be ensured.

Table 4.8 Specified options and corresponding backup data used in recovery

Type of generation number	Specified option	Backup data used in recovery	
		DB1.DBSP1 (/dev/sdd1)	DB1.DBSP2 (/dev/sdd2)
Relative generation number	-g 1	Backup data of Day 2 (/dev/sdz1)	Backup data of Day 2 (/dev/sdz2)
Absolute generation number	-v 1	Backup data of Day 1 (/dev/sdy1)	

📖 **Note**

. . . .

If you backed up in units of log groups, the situation may occur that absolute generation numbers are not complete for backups that were implemented in units of database spaces or due to deleting a history, but the absolute generation numbers of backup in units of log groups will be complete. As shown in the below example, this is due to the fact that the numbers for transaction volumes (database spaces) with smaller absolute generation numbers are matched with the numbers for transaction volumes (database spaces) with larger absolute generation numbers. Therefore, intermediate absolute generation numbers for transaction volumes (database spaces) with smaller absolute generation numbers are missing.

. . . .

Example: For a log group (LOG1/RDB1) containing Database space 1 (DB1.DBSP1) and Database space 2 (DB1.DBSP2)

Day 1: Backup of Database space 1 (DB1.DBSP1)

```
# /opt/FJSVswsts/bin/swstbackup /dev/sdd1
/dev/sdd1 swstbackup completed
```

Day 2: Backup of Database space 1 (DB1.DBSP1) and Database space 2 (DB1.DBSP2) in units of log groups

```
# /opt/FJSVswsts/bin/swstbackup -n LOG1/RDB1
LOG1/RDB1 swstbackup completed
```

Display of history information

```
# /opt/FJSVswsts/bin/swsthistdisp -n LOG1/RDB1
Server=SV01 Device=/dev/sdd1 Mount-Point=DB1.DBSP1/LOG1/RDB1 (SymfoWARE)
Generation Version  Backup-Date        Backup-Disk   Status     Execute   ArcSerial
    1         2      2003/12/12 22:00   /dev/sdz1     succeeded  ----        5
    2         1      2003/12/11 22:00   /dev/sdy1     succeeded  ----        5
Server=SV01 Device=/dev/sdd2 Mount-Point=DB1.DBSP2/LOG1/RDB1 (SymfoWARE)
Generation Version  Backup-Date        Backup-Disk   Status     Execute   ArcSerial
    1         2      2003/12/12 10:00 PM  /dev/sdz2   succeeded  ----        5
```

For recovery in units of log groups or for bundle recovery using a history as shown above, since all absolute generation numbers corresponding to the relative generation numbers are complete, the same backup data will be used regardless of which type of generation number you specify.

Table 4.9 Specified options and corresponding backup data used in recovery

Type of generation number	Specified option	Backup data used in recovery	
		DB1.DBSP1 (/dev/sdd1)	DB1.DBSP2 (/dev/sdd2)
Relative generation number	-g 1	Backup data of Day 2 (/dev/sdz1)	Backup data of Day 2 (/dev/sdz2)
Absolute generation number	-v 2		

However, if you specify "-v 1", the restore command will result in an error, because backup data for Database space 2 (DB1.DBSP2) do not exist.

P Point
. . . .
Recovery can also be carried out using the Web Console. Bundle recovery, however, cannot be carried out using the Web Console.
. . . .

P Point
. . . .
For executing a bundle recovery, a "Device list file" must be created in advance, defining the transaction volumes that are to be recovered in a batch. For details on device list files, refer to "4.1.5.3.1 Describing a device list file".
. . . .

4.1.5.3.1 Describing a device list file

Create the device list file in any suitable location on the Managed Server that carries out recovery. Specifying this file when you start recovery enables you to recover multiple transaction volumes in a batch.

The following provides an example of a device list file:

```
# A comment line starts with "#".

# LOG01/GRP1
/dev/sdd1    # A comment can be described here also.
/dev/sdd2
/dev/sdd3
/dev/sdd4
/dev/sdd5
:
:
```

The rules for creating a device list file are as follows:

- Describe a single transaction volume name in one line. One or more space or tab characters may be used as delimiters between the beginning of a line and a transaction volume name or between the end of a transaction volume name and the end of a line (new line character).

- Blank lines (spaces or tab characters) may be included in the file.

- Anything after the symbol '#' until the end of the line is regarded as a comment.

Note

If you run the Restore execution command with the -bundle option specified, all transaction volumes described in the device list file will be restored. In the following cases, restore processing will result in an error:

- When there is not a single description for a transaction volume

- When an existing description for a transaction volume contains a formal mistake

- When a transaction volume is not a Symfoware volume

- When transaction volumes are straddling across multiple log groups

- When the description for a transaction volume contains any incorrect lines other than descriptive lines

- When the same transaction volume is specified more than once, as shown in the following example:

```
# Example for multiple descriptions of same transaction volume
/dev/sdd2
/dev/sdd2
:
:
```

4.1.5.4 Managing a backup history

History information for backups that were performed with the swstbackup (Backup execution command) can be displayed and deleted.

4.1.5.4.1 Displaying a backup history

History information for backups that were performed with the swstbackup (Backup execution command) can be displayed using the Web Console or the swsthistdisp (History information display command).

4.1.5.4.2 Deleting a backup history

History information for backups that were performed with the swstbackup (Backup execution command) can be deleted using the Web Console or the swsthistdisp (History information display command).

Point

From the Web Console, you cannot specify deletion of backup history information for individual transaction volumes that are included in log groups.

4.1.6 Stopping the Operation

To stop a backup operation, stop the daemons running on the Managed Server.

Normally, these daemons are automatically stopped when the system is stopped. You can stop daemons manually if you need to do so for some reason. For more information, refer to "Chapter 2 Starting and Stopping Daemons".

Note

- When all daemons on the Managed Server have stopped, all functions of AdvancedCopy Manager running on the Managed Server will stop.

- Stop daemons on a Management Server only after ensuring that all the Managed Servers that it manages have also been stopped.

4.1.7 Operation with Database Spaces Allocated to SDX Objects

With AdvancedCopy Manager, you can back up Symfoware database spaces that are allocated to SDX objects. For more information on this operation, refer to "3.8 SDX Object Operations".

![Note icon] Note

. . . .

For backing up Symfoware database spaces that are allocated to SDX objects in units of logical volumes, observe the following notes.

- For backing up in units of log groups, all Symfoware database spaces belonging to a log group must be created in the SDX object. Otherwise, operation will be available only in units of database spaces.

- Symfoware database spaces that are allocated to the root class are not available for backup operation. Allocate the database spaces to local or shared classes.

. . . .

Chapter 5 Backup Operation using Suspend/Resume Function

This chapter provides details of the backup operation using the Suspend/Resume function.

5.1 Overview

The Suspend/Resume function sets Suspend/Resume for the equivalency maintenance status using EC. Using this function for differential copying from the Suspend state enables higher-speed synchronous backup.

AdvancedCopy Manager provides two backup functions: snapshot high-speed backup and synchronous high-speed backup. In snapshot high-speed backup, OPC is used to create a backup. In synchronous high-speed backup, EC is used to create a backup. The Suspend/Resume function can be thought of as a kind of synchronous high-speed backup that uses EC.

With ordinary synchronous high-speed backup, the synchronous processing start command starts copying using the EC function, then sets the equivalency maintenance status after the copying is completed. The backup execution start command is then executed. This cancels EC and creates a backup.

In synchronous high-speed backup using the Suspend/Resume function, when the backup execution start command is executed (with -suspend specified), EC is suspended, the backup process is completed, and the EC session status is saved. The next time that the synchronous processing start command is executed, EC resumes and only data that has been updated since the suspend state was last set is copied. As a result of this, the preparation time of the backup is shortened.

A comparison between ordinary synchronous high-speed backup and synchronous high-speed backup using the Suspend/Resume function is shown below.

Figure 5.1 Comparison between ordinary synchronous high-speed backup and synchronous high-speed backup using the Suspend/Resume function

5.2 Backup Operation Design

Notes about synchronous high-speed backup using the Suspend/Resume function and backup operation design are provided below, as well as information related to backup management operation design (refer to "3.2 Operation Design").

Transaction volume

Register the backup source device, used for storing transaction data, as the transaction volume.

Backup volume

Register the backup destination device in which transaction volume data is saved as the backup volume.

When the backup command with -suspend option specified is executed, the EC session between the transaction volume and backup volume is maintained in a Suspended state. In this Suspend state, the backup volume cannot be used as the backup destination of another transaction volume.

🅿 Point

. . . .

In the backup process that uses the Suspend/Resume function, each transaction volume requires as many backup volumes as the saved generation number + 1.

. . . .

5.3 Flow of Operations

The flow of backup operations using the Suspend/Resume function is shown below.

Figure 5.2 Flow of backup operations using Suspend/Resume function

5.4 Preparations

This section describes how to prepare the backup operation that uses the Suspend/Resume function.

5.4.1 Configuring the AdvancedCopy Manager environment

For information on the sequence of processes from the installation to the operating environment configuration of AdvancedCopy Manager, refer to the "ETERNUS SF Installation and Setup Guide" and the "ETERNUS SF AdvancedCopy Manager Operation Guide".

5.4.2 Setting the environment for Managed Server

Set the configuration information on the Managed Server processing the backup operations.

This process can be carried out using either Web Console or swstsvrset (Managed Server configuration information setting command).

In addition, once it is set, you can display the Managed Server configuration information by executing swstsvrdisp (Managed Server configuration information display command).

5.4.3 Setting the operation type for a device

Register a transaction volume as the backup source and a backup volume to be used as the backup destination.

This process can be carried out using either Web Console or swstdevinfoset (Device information setting command).

In addition, to display the registered volume information, execute swstdevdisp (Device usage status display command).

🛑 Note
. . . .

- A suspended transaction volume cannot be changed to be used for a different purpose.

- A suspended backup volume cannot be changed to be used for a different purpose.

. . . .

🅿 Point
. . . .

When the device information of a suspended backup volume is displayed, "Backup (used)" is displayed in the **Device-Mode** field.

. . . .

5.4.4 Setting a backup policy

Before a backup policy is set, you need to register as many usable backup volumes as are required to perform the operation in accordance with the specified backup policy.

This process can be carried out using either the Web Console or commands.

- Backup Policy Setup: swstbkpolset (Backup policy setting command)

- Check Backup Policy that has been set up: swstbkpoldisp (Backup policy display command)

- Backup Policy Deletion: swstbkpoldel (Backup policy deletion command)

5.4.5 Preparing a device map file

In its backup operations, AdvancedCopy Manager automatically selects, from a group of devices registered as backup volumes, a device that has the same capacity as the transaction volume. It then uses the selected device as the backup destination.

If there is a suspended backup volume at the time of backup destination selection and the start of the synchronous processing, its corresponding device is selected first, and is used as the backup destination.

To conveniently monitor backup destination devices, create a device map file that maps transaction volumes to backup volumes in advance. Specifying this file at the start of the synchronous processing or during the backup process enables operation with the backup destination specified.

🅿 Point
. . . .

To manage multiple generations of backups, multiple device map files must be prepared.

. . . .

5.5 Backup Operations

For synchronous high-speed backup using the Suspend/Resume function, execute swstbackup (Backup execution command) after the synchronous processing has started and the equivalency maintenance status has been set.

Then suspend the synchronous processing.

When history information is deleted because the number of generations has been exceeded or when swsthistdel (History information deletion command) is executed, the backup volume becomes the Suspend state without any history information.

If new synchronous processing starts at this time, then the suspended backup volume without history information is selected first, and differential copying starts. Only data that has been updated since the Suspend state was last set is copied, therefore the equivalency maintenance status can be reached quickly.

Note that synchronous high-speed backup using the Suspend/Resume function can also be operated from Web Console.

The flow of synchronous high-speed backup operations using the Suspend/Resume function is shown below.

Figure 5.3 Synchronous high-speed backup operations using the Suspend/Resume function

4) Execution of the backup command with -suspend specified

TD ◄—Suspended EC session—► BD1 *With history information* In use

Number of preservation generations=1

No EC session — BD2 Empty

5) Synchronous processing started

TD ◄—Suspended EC session—► BD1 *With history information* In use

Number of preservation generations=1

Full copy EC session ► BD2 In use

6) Equivalency maintain status

TD ◄—Suspended EC session—► BD1 *With history information* In use

Number of preservation generations=1

◄ Equivalent EC session ► BD2 In use

7) Backing up with the suspend specification

TD ◄—Suspended EC session—► BD1 Usable

Number of preservation generations=1

◄—Suspended EC session—► BD2 *With history information* In use

> History information deleted because number of generations exceeded (*1) at backup command execution

*1 "Number of generations exceeded" means that the oldest history information is automatically deleted if existing history information has certain number of saved generations when the backup command is executed.

8) Transaction volume data updated

TD [Data updated] ◄—Suspended EC session—► BD1 Usable

Number of preservation generations=1

◄—Suspended EC session—► BD2 *With history information* In use

5.5.1 Starting synchronous processing

For synchronous high-speed backup, execute swststartsync (Backup synchronous processing start command) to start the synchronous processing. If there is a suspended backup volume at the start of the synchronous processing, then the following devices are selected first, and EC resumes (i.e., the Resume function causes differential copying to start).

   - The backup volume from which history information is deleted because the number of generations has been exceeded; or

   - The suspended backup volume without history information.

If there is no suspended backup volume at the start of the synchronous processing, EC (i.e. a full copy) starts.

The process of backup volume selection at the start of the synchronous processing is shown below.

Figure 5.4 Backup volume selection at the start of synchronous processing

Note

If synchronous processing is started with a device map specified, other transaction volumes and suspended backup volumes cannot be used as the backup destination. To check which transaction volumes are suspended along with their associated backup volumes, execute swstsyncstat (Backup synchronous processing progress display command).

5.5.2 Backing up

For synchronous high-speed backup using the Suspend/Resume function, start synchronous processing. Then, when the copy status of both the transaction volume and backup volume attain equivalency maintenance status, execute swstbackup (Backup execution command) with the suspend specification (using -suspend option).

In synchronous high-speed backup, executing the backup command without the suspend specification releases the EC session. Executing it with the suspend specification suspends the EC session instead of releasing it.

Ordinary backup operations and backup operations with the suspend specification are shown below.

Figure 5.5 Ordinary backup operations and backup operations with the suspend specification

5.5.3 Displaying the synchronous processing execution status

To check the status of synchronous high-speed backup that uses the Suspend/Resume function, execute swstsyncstat (Backup synchronous processing progress display command). If one transaction volume has multiple EC sessions, then this command displays the status of all EC sessions. The statuses displayed for synchronous processing are explained below.

- executing : Synchronous processing (i.e., copying) is in progress

- equivalent : Equivalency maintain status

- suspend : Suspended processing

The command displays the synchronous processing execution status as shown below:

Figure 5.6 Ordinary backup operations and backup operations with the suspend specification
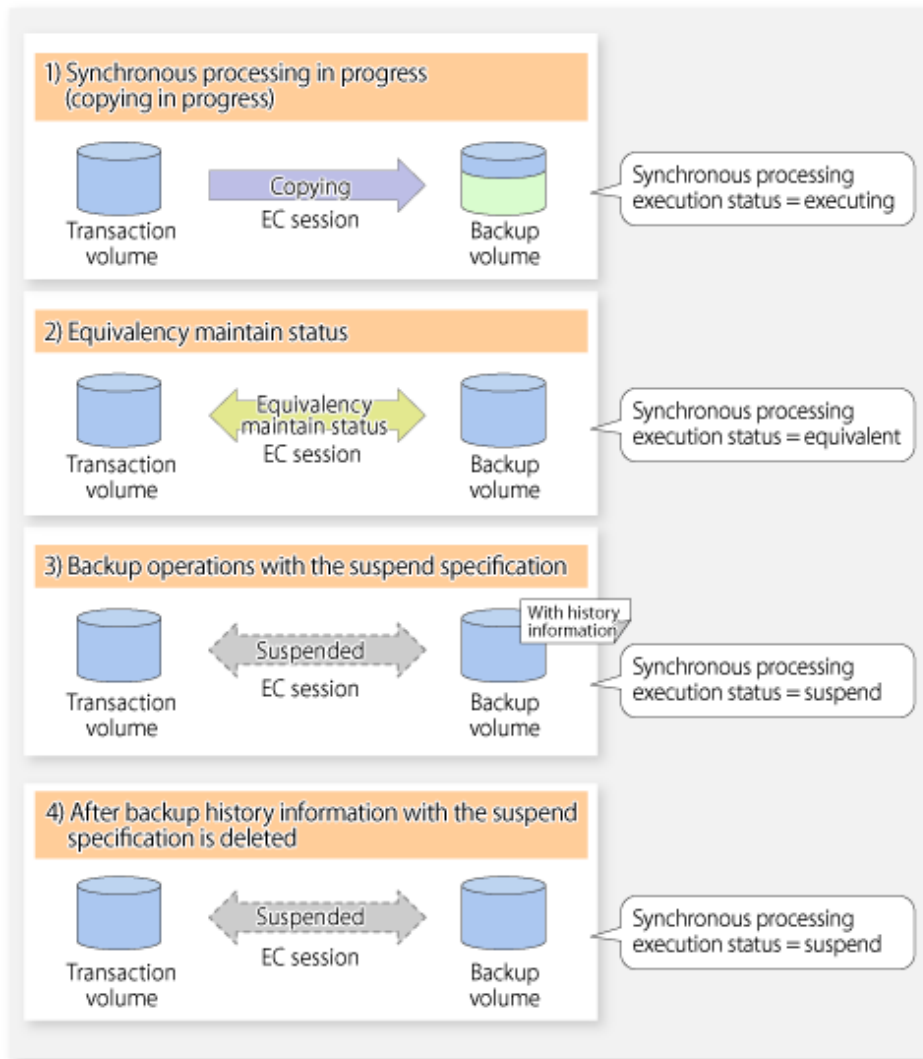
5.5.4 Cancelling synchronous processing

To cancel synchronous processing (during an EC session), execute swstcancelsync (Backup synchronous processing cancel command). The statuses of the different types of synchronous processing cancelled are explained below.

- executing : Synchronous processing (i.e., copying) is in progress

- equivalent : Equivalency maintain status

- suspend : Suspended processing

After the backup command with -suspend option specified is executed, the Suspend state is maintained even if history information is deleted. A suspended EC session cannot be released unless the synchronous processing is cancelled.

If the synchronous processing is cancelled with no options specified, the EC session in the synchronous processing status or in the equivalency maintenance status is cancelled.

If one transaction volume has multiple synchronous processing in progress (i.e., running multiple EC sessions), then only the EC sessions mapped to specific backup volumes can be cancelled. All of the synchronous processing of the transaction volume can be cancelled by specifying the synchronous processing.

🛈 Note

The backup volume and all synchronous processing cannot both be specified simultaneously.

Ⓟ Point

- Cancelling the suspended synchronous processing that has history information does not delete the history information.

- Cancelling synchronous processing (during copying, in the equivalency maintenance status, or in the Suspend state) that does not have any history information sets the backup volume to the unused state.

- Cancelling synchronous processing (in the Suspend state) that has history information leaves the backup volume in the used state.

A range of operational states when synchronous processing is cancelled are shown below:

Figure 5.7 Operational states in synchronous processing caused by cancellation

5.5.5 History management

Backup history information can be referenced using swsthistdisp (History information display command).
Backup history information can be deleted using swsthistdel (History information deletion command).

P Point

. . . .

- Even when the history information created during backup operations with the suspend specification is deleted, an associated EC session is not released.

- Deleting history information during ordinary synchronous high-speed backup releases the backup volume, and "Backup(free)" is displayed in the **Device-Mode** field by swstdevdisp (Device usage status display command).

- Deleting history information during backup operations with the suspend specification does NOT release the backup volume after the information is deleted, and "Backup(used)" is displayed in the **Device-Mode** field by swstdevdisp (Device usage status display command).

. . . .

The process of deleting ordinary backup history information and backup history information with the suspend specification is shown below.

Figure 5.8 Deleting ordinary backup history information and backup history information with the suspend specification

5.5.6 Restoring

Before restoration, swstcancelsync (Backup synchronous processing cancel command) must be executed to release all EC sessions that are set on the transaction volume. To check EC sessions set on the transaction volume, execute swstsyncstat (Backup synchronous processing progress display command).

To restore data from a backup volume, execute swstrestore (Restore execution command).

To cancel restoration or to delete management information associated with restoration that ended abnormally, execute swstcancelrest (Restore cancellation command).

Note

- Restoration is not possible if the transaction volume has a backup volume being copied by EC, a backup volume under the equivalency maintenance status, or a suspended backup volume.

- A suspended backup volume cannot be used as the restore destination volume.

Chapter 6 Backup Operation using QuickOPC Function

This chapter provides details of the backup operation using the QuickOPC function.

6.1 Overview

For a snapshot high-speed backup using OPC, the entire transaction volume must be copied to the backup volume every time the backup is started. However, as long as the same transaction volume and backup volume are being used, the entire transaction volume need not be physically copied every time OPC is run. Instead, physically copying only the data that has been modified since the start of the previous backup can create a complete snapshot image. This can greatly reduce the time taken for physical copying.

The QuickOPC function achieves this by copying only the data that has been changed from the point of the previous physical copy.

The snapshot high-speed backup using the QuickOPC function is referred to as the differential snapshot high-speed backup.

The conventional snapshot high-speed backup and the differential snapshot high-speed backup are compared below:

📒 Note

. . . .

To implement the differential snapshot high-speed backup using the QuickOPC function, ETERNUS Disk storage system that supports the QuickOPC function is required.

The QuickOPC function cannot be used for backing up SDX objects in units of logical volumes.

. . . .

🅿 Point

. . . .

When the QuickOPC function is used, the hardware records the changes that are made to the copy source and destination after completing the OPC logical copy. The state in which the hardware keeps a record of changes is called the "tracking state".

. . . .

Figure 6.1 Comparison between conventional snapshot high-speed backup and differential snapshot high-speed backup
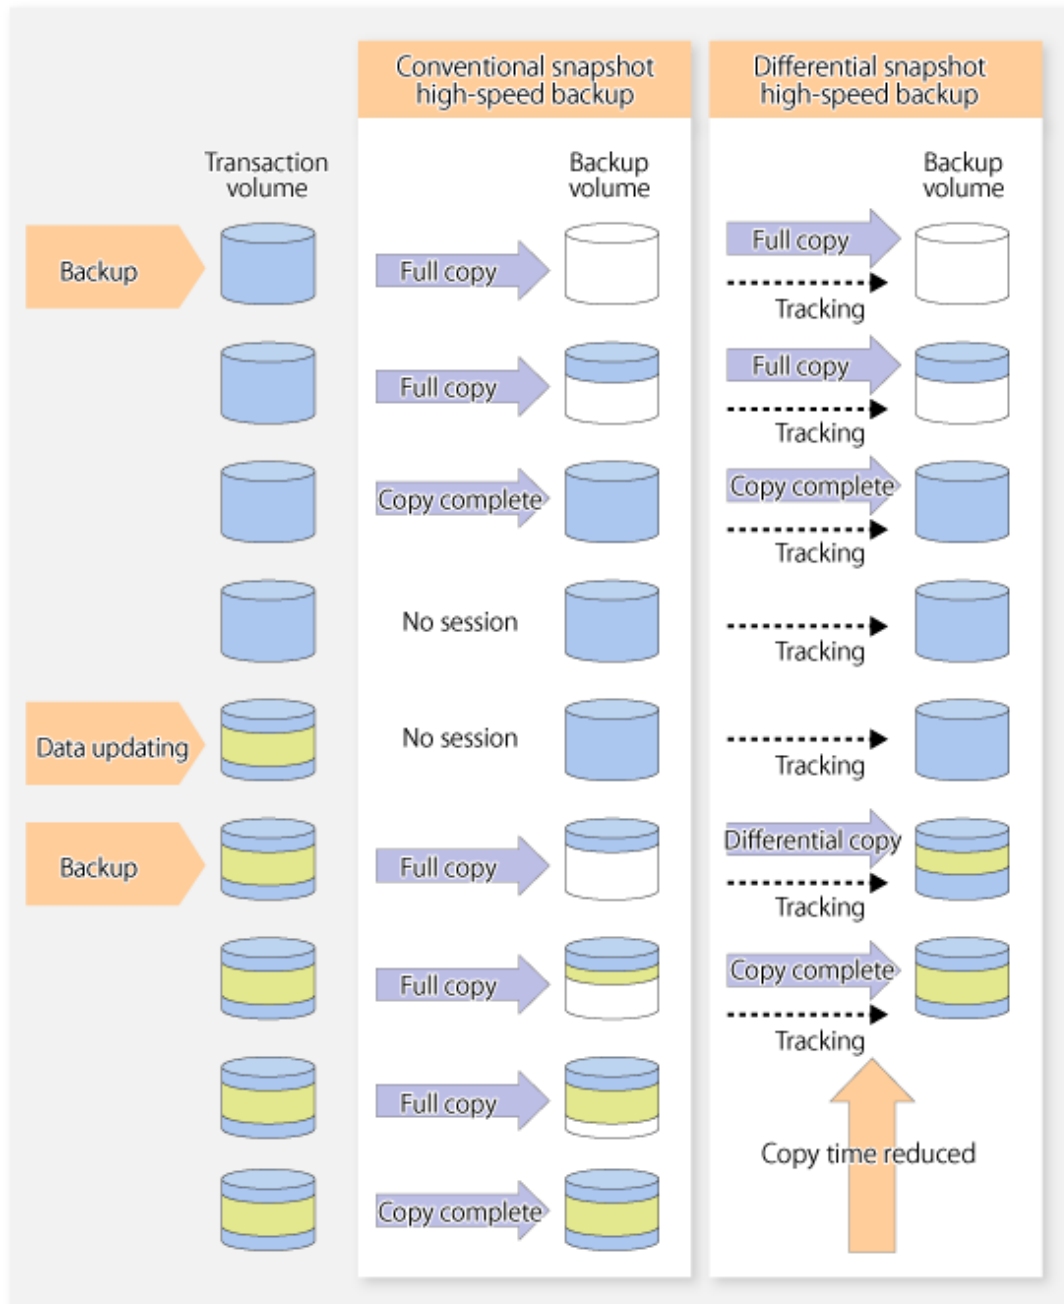
6.2 Configuring Backup Operation

Configure the differential snapshot high-speed backup using the QuickOPC function using information provided in the section relating to "3.2 Operation Design" for backup management with consideration to the following notes:

![Note]

Consider the following when making decisions on the backup policy configuration and backup volume preparation:

- In a backup operation using the QuickOPC function, an OPC session remains active even after the physical copy has been completed. Operators should specify the number of preservation generations so that the maximum number of OPC sessions that can be defined for the same logical unit cannot be exceeded.

- Operators must prepare a backup volume in the same box as the transaction volume to be backed up.

6.3 Operation Flow

The flow of operation using the QuickOPC function is shown below.

Figure 6.2 Flow of backup operations using the QuickOPC function

6.4 Preparation

This section provides details of the preparation for a backup using the QuickOPC function.

6.4.1 Configuring the AdvancedCopy Manager environment

For information on the sequence of processes from the installation to the operating environment configuration of AdvancedCopy Manager, refer to the "ETERNUS SF Installation and Setup Guide" and the "ETERNUS SF AdvancedCopy Manager Operation Guide".

6.4.2 Setting the environment for Managed Server

Set the configuration information of the Managed Server processing the backup operations.

This process can be carried out using either Web Console or swstsvrset (Managed Server configuration information setting command).

In addition, once it is set, you can display the Managed Server configuration information by executing swstsvrdisp (Managed Server configuration information display command).

6.4.3 Setting the operation type for a device

Register a transaction volume as the backup source and a backup volume to be used as the backup destination.

This process can be carried out using either Web Console or swstdevinfoset (Device information setting command).

In addition, to display the registered volume information, execute swstdevdisp (Device usage status display command). The transaction volume in the tracking state can be checked with swsttrkstat (Tracking status display command).

Note

- The settings of the transaction volume in the tracking state cannot be changed.

- The settings of the backup volume in the tracking state cannot be changed.

Point

The both -b and -u options can be specified in swstdevdisp (Device usage status display command) to check for backup volumes, where relevant histories have been deleted, in the suspend or tracking state.

6.4.4 Setting a backup policy

Before a backup policy is set, you need to register as many usable backup volumes as are required to perform the operation in accordance with the specified backup policy.

This process can be carried out using either the Web Console or commands.

- Backup Policy Setup: swstbkpolset (Backup policy setting command)

- Check Backup Policy that has been set up: swstbkpoldisp (Backup policy display command)

- Backup Policy Deletion: swstbkpoldel (Backup policy deletion command)

6.4.5 Preparing a device map file

In the backup, AdvancedCopy Manager automatically selects a device which has the same capacity as the target transaction volume as a backup volume from a group of registered devices.

If a backup volume in the tracking state exists when AdvancedCopy Manager looks for a backup volume for the snapshot high-speed backup, it selects the backup volume in the tracking state in preference to any other and uses it as the backup destination.

When it is necessary to recognize a specific backup destination device for any reason a "device map file" that defines the relationship between a transaction volume and backup volume must be created in advance. By specifying the "device map file" when starting the backup, the backup can always be performed because the backup destination has already been specified.

Point

If multiple generations are to be managed, multiple device map files must be prepared.

6.5 Operation

6.5.1 Invoking backup

Invoke the differential snapshot high-speed backup by specifying the -T option in swstbackup (Backup execution command).

When the QuickOPC starts, the physical copy and tracking processing begins. When the physical copy is complete, only tracking processing is active. When this backup command is executed by specifying the -T option, only new data or data which has changed from the point in time that OPC was started previously is actually copied.

Note that differential snapshot high-speed backup can also be operated from Web Console.

The figure below shows the conventional snapshot high-speed backup and the differential snapshot high-speed backup.

Figure 6.3 Conventional snapshot high-speed backup and differential snapshot high-speed backup

After history information is deleted in a differential snapshot high-speed backup, the backup volume is put into the tracking state without history information. If a new differential snapshot high-speed backup is started under that condition, the backup volume in the tracking state without history information is selected to start the differential data copy. The physical copy can thus be completed in a short period because only the difference in data from where OPC was started previously is copied.

The figure below shows a flow of a differential snapshot high-speed backup in which only one generation is preserved.

Figure 6.4 Flow of differential snapshot high-speed backup operation

6.5.2 History management

Backup history information can be referenced using swsthistdisp (History information display command).
Backup history information can be deleted using swsthistdel (History information deletion command).

Even if the history information created by the differential snapshot high-speed backup is deleted, the tracking state is not released (that is, the physical copy, if it is still being executed, is not yet completed or stopped).

In the conventional snapshot high-speed backup, after the history information is deleted, the backup volume is released and "Backup (free)" is displayed in the **Device-Mode** field by swstdevdisp (Device usage status display command). In the differential snapshot high-speed backup, however, the backup volume is not released even after the history information is deleted and "Backup (used)" is displayed in the **Device-Mode** field by swstdevdisp (Device usage status display command).

To check for backup volumes in the tracking state, specify the -u option in the device use status display command. For a backup volume in the tracking state, "Backup (used-T)" is displayed in the **Device-Mode** field.

The figure below shows a comparison of history deletion between the conventional snapshot high-speed backup and the differential snapshot high-speed backup.

Figure 6.5 Comparison of history deletion between snapshot high-speed backup and differential snapshot high-speed backup
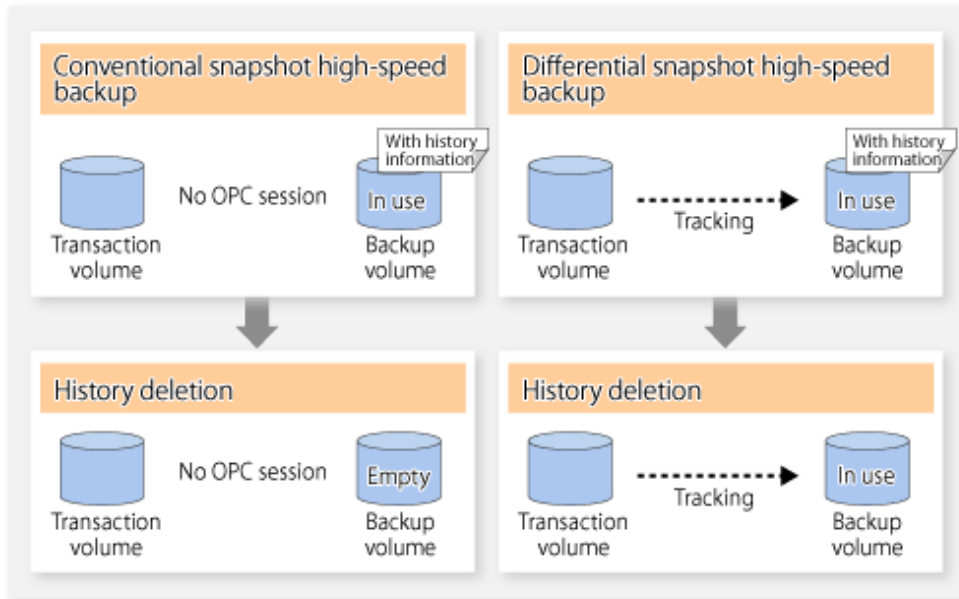
6.5.3 Restoration

The section explains how to use swstrestore (Restore execution command) for restoration.

In a differential snapshot high-speed backup, tracking processing from the transaction volume to the backup volume is performed.

Figure 6.6 Tracking of differential snapshot high-speed backup operations

When restoration is executed, normal OPC from the backup volume to the transaction volume is started while the tracking state from the transaction volume to the backup volume is maintained.

Because OPC for restoration physically copies only the data that has been updated after the backup, the time taken for the physical copy can be reduced for the restoration, as well as for the backup.

Figure 6.7 Restoration from backup volume to transaction volume

While tracking processing is performed between the transaction volume and backup volume, restoration to other volumes cannot be performed. If restoration to other volumes is required, the following operational steps must be performed in advance:

1. If the OPC physical copy from the transaction volume to the backup volume is in progress, wait for the completion of the copy.

2. Stop tracking processing using swstcanceltrk (Tracking cancel command).

Figure 6.8 Restoration to another volume

 Note

If tracking processing is terminated, all data in the transaction volume is physically copied by the next differential snapshot high-speed backup.

6.5.4 Cancelling restoration

Use swstcancelrest (Restore cancellation command) to cancel the restoration.

6.5.5 Displaying the execution status of tracking processing

The execution status of tracking processing by the QuickOPC function can be checked with swsttrkstat (Tracking status display command).

Information on an OPC session that is not in a tracking state is excluded from the display. For an OPC session in tracking, the amount of data that has been updated after execution of the QuickOPC (ie, the logical copy) is displayed as a percentage in the **Update** field.

6.5.6 Stopping tracking processing

Use swstcanceltrk (Tracking cancel command) to stop tracking processing.

Chapter 7 Replication

This chapter explains AdvancedCopy Manager replication in a Linux system.

A partition or Logical Unit (disk) is used as the replication unit.

In this manual, the term "volume" is used instead of "partition" or "Logical Unit (disk)".

7.1 Overview

This chapter explains the operation of the replication function.

Using the Advanced Copy function of ETERNUS Disk storage system, the replication function performs high-speed replication between volumes, regardless of the volume capacities.

Snapshot replication is the term used for the replication function that uses the OPC function of ETERNUS Disk storage system.

Synchronized Replication is where either the EC or REC functions are used.

The replication function creates copies so that the information at a certain point in time can be used for different purposes. It can be used either independently of or in combination with the backup function.

Unlike in the case of backup management, Symfoware database is not supported by replication management.

7.1.1 Snapshot replication processing

Snapshot replication copies from a source volume to a destination volume using the function of ETERNUS Disk storage system.

The ordinary snapshot replication using OPC performs the following two types of processing:

1. When a copy is created, the snapshot processing (OPC) from the source volume to the replica volume is executed with swsrpmake (Replication creation command).

   swsrpmake (Replication creation command) dismounts or mounts the original volume.

   This process determines the scope of original data (for details, refer to "Appendix C Pre-processing and Post-processing of Replication"). Refer to steps (1). and (2). in the figure below.

2. If the user wants to recreate a replica, only swsrpmake (Replication creation command) is required. If the previous snapshot processing is in progress, it is stopped and a new snapshot process is started.

Figure 7.1 Ordinary snapshot replication processing

Snapshot replication is completed when the replication creation command is issued. During snapshot processing performed internally by the ETERNUS Disk storage system, if there is an attempt to access an area where copying is not yet complete at the replication destination, priority is given to copying the accessed area and, after it is copied, access is permitted ((a) in the figure).

If there is an attempt to update the replication source data, the pre-update data at the replication source is first copied to the replication destination, and then the replication source data is updated ((b) in the figure).

These operations make it seem that creation of the replica is completed almost instantly.

7.1.1.1 Types of snapshot replication

The following types of snapshot replication are available, depending on the used Advanced Copy functions:

- Ordinary snapshot replication

- QuickOPC type replication

- SnapOPC type replication

- SnapOPC+ type replication

Ordinary snapshot replication

Ordinary snapshot replication copies all the data at a certain point in time (logical copy) to the copy destination disk area.

Figure 7.2 Ordinary snapshot replication

P Point
. . . .
The copy destination disk must be the same size or larger than the copy source disk.
The copying time is the time taken to copy all the data.
. . . .

QuickOPC type replication

If the ETERNUS Disk storage system supports the QuickOPC function, QuickOPC replication can be performed.
The initial replication copies all the data at a certain point in time (logical copy) to the copy destination disk area.
Subsequent replications copy only the data that has been updated since the previous replication.
This enables large reductions in the time taken for physical copying.

Figure 7.3 QuickOPC replication

Point

. . . .

The copy destination disk must be the same size or larger than the copy source disk.

The copying time for subsequent replications is the time taken to copy the differential data.

. . . .

Note

. . . .

To execute QuickOPC replication, it is necessary to use the ETERNUS Disk storage system that supports the QuickOPC function.

The QuickOPC function cannot be used for replicating SDX objects in the units of logical volumes.

. . . .

Point

. . . .

Hardware using the QuickOPC function records the updates made to the copy source or destination after the completion of OPC logical copying. The status in which the hardware records and displays updates is referred to as the "tracking status".

. . . .

SnapOPC type replication

SnapOPC is a function that enables only data that has been updated in the copy source disk after a certain point in time (a logical copy) to be copied to the destination disk.

Figure 7.4 SnapOPC replication

. . . .

The copy destination disk can be smaller than the copy source disk.

The copy time is the time needed to copy data that has been updated.

. . . .

 Information

. . . .

The SnapOPC copy destination disk is called "Snap Data Volume".

. . . .

Figure 7.5 SnapOPC mechanism

In some ways, SnapOPC is superior to conventional OPC, since copy times are shorter and copy destination volumes can have a smaller capacity. However, there are problems in terms of access performance and the reliability of copy data, as explained in the following:

Access performance

Since the data copy from the source to the destination occurs inside SnapOPC, the access performance of the copy source is sometimes reduced, as well as the access performance of the copy destination.

Reliability of copy data

Copy data is made up of two parts: the "copy source" and "updated part of the copy source". Therefore, a hardware fault with the copy source can cause copy data to be lost.

Additionally, performing a restoration is equivalent to losing the "copy source" data, and so even if there are multiple "updated parts of the copy source", these become invalid from the point when the restoration is performed.

These points suggest that SnapOPC should be used in temporary areas for tape backups, for systems where access performance is not an important consideration.

SnapOPC+ type replication

When data is changed, the SnapOPC+ function copies only the original data of the changed data set, and saves it in snap generation units. The following figure shows the processing flow of the SnapOPC+ function.

Figure 7.6 SnapOPC+ function flow
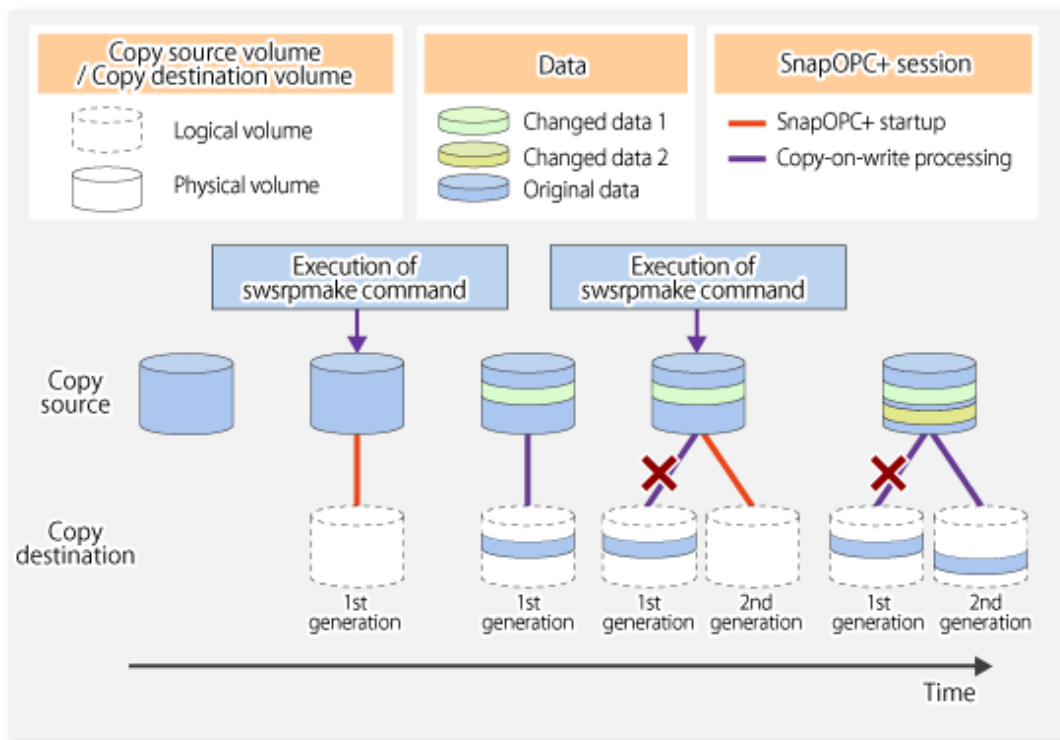
1. Logical copy is performed from the copy source to the copy destination.

2. After the logical copy, a SnapOPC+ session is set between the copy source and the copy destination.

3. After the SnapOPC+ session is set, only the original data of the changed data at the copy source is copied to the copy destination (copy-on-write processing).

4. When the next copy destination volume is created, the copy-on-write processing is stopped.

5. Then, a SnapOPC+ session is set between the copy source volume and the new copy destination volume, and copy-on-write processing is performed.

6. Each time a copy destination volume is created, a snap generation number is assigned to that volume.
   (Snap generation numbers are assigned in sequence from 1, starting with the oldest.)

Since SnapOPC+ copy-on-write processing is performed only between the copy source volume and the latest copy destination volume, there is no effect on access performance to the copy source volume even when the number of copy destination volumes increases.

In addition, space at the copy destination volume is required only for the amount of original data (for data updated immediately prior) and for hardware management. Thus, use of SnapOPC+ enables the saved data area to be kept to the minimum area required.

Snapshot replication that uses SnapOPC+ is called "SnapOPC+ replication".

Point

It is recommended to set the Snap Data Pool using the ETERNUS Web GUI when using SnapOPC+.

For details, refer to "7.1.1.2 Snap Data Volume and Snap Data Pool".

The following figure shows use of SnapOPC+ to perform 1:N replication operations between a copy source volume (1) and multiple snap generations (N).

Figure 7.7 1:3 replication operations using SnapOPC+

Due to the mechanism and features of SnapOPC+, it is assumed that it will be used as a backup to enable recovery from minor errors, such as operator mistakes and software errors. In order to be prepared for hardware errors, Fujitsu recommends that SnapOPC+ be used in conjunction with the copying of all data by means of OPC/QuickOPC/EC/REC.

The following table lists the functional differences between SnapOPC and SnapOPC+.

Table 7.1 List of function differences between SnapOPC and SnapOPC+

Function	SnapOPC	SnapOPC+
Copy method	Copy-on-write method	Copy-on-write method
Copy time (logical copy duration)	Instant	Instant
Copy destination volume type	Snap Data Volume	Snap Data Volume
1:N operation	Available	Available (recommended): Access performance to the copy source volume is better than with SnapOPC. Also, compared with SnapOPC, less physical space is required at the copy destination volume.
Snap generation numbers	No	Yes
Restoration using the replication function	Yes	Yes (**Note**)
Usefulness	- Tape backup temporary area - Backup of file servers or similar that have a small amount of updates	- Tape backup temporary area - Backup of file servers or similar that have a small amount of updates

Note: Only available when using ETERNUS Disk storage system that supports restoration using replication.

7.1.1.2 Snap Data Volume and Snap Data Pool

Snap Data Volume

The Snap OPC/SnapOPC+ copy destination volume is called "Snap Data Volume". Regarding this copy destination volume (Snap Data Volume), the capacity that can be seen in the host is called "logical capacity", and the capacity that has been physically configured is called "physical capacity".
The setting for logical capacity is equal to or greater than for the copy source capacity.
Physical capacity has control information (a conversion table) for managing backup data, as well as capacity for the storage of backup data. The capacity for this control information is 0.1% that of logical capacity.

Note

When there is insufficient Snap Data Volume capacity, it is not possible to access source volume (Snap Data Volume).
When there is insufficient capacity, refer to "8.4.2.3 Troubleshooting when a lack of free space has occurred in the Snap Data Volume or Snap Data Pool".

Snap Data Pool

In ETERNUS Disk storage system, it is possible to set a Snap Data Pool, which is part of the Snap Data Volume group functionality.

When there is insufficient Snap Data Volume capacity and a Snap Data Pool is set, some area is added from the Snap Data Pool automatically.

Note

When there is insufficient Snap Data Pool capacity, it is not possible to access the copy destination (Snap Data Volume).
When there is insufficient capacity, refer to "8.4.2.3 Troubleshooting when a lack of free space has occurred in the Snap Data Volume or Snap Data Pool".

Figure 7.8 Image of Snap Data Volume extension

7.1.2 Synchronized replication processing

Replicas are created by copying from the source volume to the replica volume using the EC or REC function of ETERNUS Disk storage system.

Synchronized replication uses two methods to create a copy: full copying and incremental copying. In the case of full copying, all of the source volume is completely copied. In the case of incremental copying, only the data that has been updated since the last replication is copied.

- Creating replicas with full copying
  Use this method to create the first copy in a replication.

- Creating copies with incremental copying
  Use this method to create copies in the future once a previous replica exists.
  Copies are created by reflecting the data that has been updated after an initial copy has been made (full copying and incremental copying).

The procedure for synchronized replication incorporates the following steps:

1. Full copying starts when synchronized processing that has been invoked using swsrpstartsync (Synchronous processing start command) starts copying from the source volume to the replica volume. Refer to step (1) in the figure below.

2. Full copying is completed when the source volume and replica volume both attain the equivalency maintenance status. From this point on, updates in the source volume are subsequently reflected in the replica volume, so that data equivalence is always maintained (this is called an equivalency maintenance state). Refer to the states in steps (2) and (3) in the figure below.

3. Synchronized processing is temporarily stopped with swsrpmake (Replication creation command) to create a copy of the source volume.

   swsrpmake (Replication creation command) dismounts or mounts the original volume.

   This processing determines the scope of original data (for details of this, refer to "Appendix C Pre-processing and Post-processing of Replication". After execution of the replication creation command, the replica volume can be accessed (copy established status). Refer to the states in steps (3) and (4) in the figure below.

4. To create another copy, the synchronized processing that has been temporarily stopped is restarted with swsrpstartsync (Synchronous processing start command). In the case of this copy, the only data copied to the replica volume is the data on the source volume that has been updated since the previous copy was made. Refer to step (4) in the figure below.
   When the replica volume has been updated after the previous copy is made, the contents of the update of the replica volume are cleared.

5. When the incremental data has been copied, the status returns to the equivalency maintenance state. Refer to the states in steps (5) and (6) in the figure below.

6. When this status has been set, the copy is re-created with swsrpmake (Replication creation command). Refer to step (6) in the figure below.

Figure 7.9 Synchronized replication processing

Note

- When using the EC or REC function in synchronized replication, ETERNUS Disk storage system must support the EC or REC function.

- swsrpmake (Replication creation command) cannot be executed unless both the source volume and replica volume are in an equivalent state.

7.2 Configuration

This section explains the operation of replication.

Configure replication using the following procedure:

1.

Refer to "7.2.3 Notes on the configuration of SnapOPC/SnapOPC+ replication", when SnapOPC/SnapOPC+ snapshot replication is used.

To backup or restore a system volume, refer to "Appendix D System Volume Backup Operation".

7.2.1 Determining the server that performs replication

Determine the server that performs replication.

The following servers are used to perform replication:

Management Server

    Multiple Managed Servers are unified and operated centrally. The Management Server can also be used as a Managed Server.

Managed Server

    AdvancedCopy Manager operations are performed.

Figure 7.10 Configuration of the servers that perform replication

Note

  - The following requirements apply when an inter-box replication is performed using the REC function of ETERNUS Disk storage system:

    - The REC function must be supported in both ETERNUS Disk storage systems.

    - Both ETERNUS Disk storage systems must already be connected to each other using an FC Remote Adapter (FCRA).
    This is a connection that uses an FCRA, in which data flows only from the Initiator side to the Target side. For bi-directional copying, at least two FCRA connections are required.

7.2.2 Determining a replication target

Determine the source and replica volumes to be used in the replication.

....

- The source volume is the volume on which the data to be copied is stored.

- The replica volume is the volume on which a copy of the data is stored.

....

Determine the following attributes to be assigned to source and replica volumes:

- Copy direction

  Determine the direction of copying. Specify one of the following directions:

  bi-directional

    Data can be copied in either direction between source volume and replica volume.

  uni-directional

    Data can only be copied from the source volume to the replica volume. During this operation, copying from the replica to the
    original volume is disabled.

- Operational servers (used only for server-to-server replication)

  In the case of replication between servers, you can specify whether to allow replication to be performed on either the source server or
  destination server, or on both servers.

  Original server

    Sets the Managed Server with the source volume that was connected as the operational server.

  Replica server

    Sets the Managed Server with the replica volume that was connected as the operational server.

  Both servers

    Sets both the source and destination servers as operational servers.

It is possible to create a group by arranging multiple source volumes and destination volumes.
In addition, the operation for multiple source volumes and destination volumes can be performed in a group. For details, refer to "7.4.7
Creating groups".

Figure 7.11 Servers used in performing replication

In this manual, a Managed Server used to perform replication is called the operational server. An original server or replica server is specified
as the operational server.

An operational server can perform all of the replication functions. In contrast, any other type of server can use only the information display function and operational release function.

🔖 Note
・・

- Do not select the logical volume and volume group of LVM (Logical Volume Manager) as a replication target.

- In order to set up a volume as a source volume, there are some points which should be taken into consideration. For details, please refer to "10.1.1 General notes", for notes on determining a replication target.

    - 10.1.1.2 Backup and Replication Exclusions

    - 10.1.1.1 Size of a Partition

    - 10.1.1.5 Maximum Number of Simultaneously Executing Operations

    - 10.1.1.8 Mirroring in GDS

・・

7.2.2.1 Logical Unit (disk) operation

Replication operations can copy data in Logical Units (disk units).
Logical Unit (disk) unit copying supports the following disk formats:

- MBR disk format (DOS format)

- GPT disk format

To execute Logical Unit (disk unit) copy, refer to "10.1.1.9 Notes on executing Logical Unit (disk) unit copy".

7.2.3 Notes on the configuration of SnapOPC/SnapOPC+ replication

7.2.3.1 Notes on session configurations

Only one session can be set up for each Snap Data Volume, which is the copy destination disk for SnapOPC/SnapOPC+.

Figure 7.12 If one session is set for the Snap Data Volume

Accordingly, multiple sessions cannot be set up for a single Snap Data Volume, as shown in the following figure:

Figure 7.13 If multiple sessions are set for the Snap Data Volume

The following additional restrictions apply:

- It is not possible to copy from the Snap Data Volume to a disk other than the copy source disk while SnapOPC/SnapOPC+ is running.

Figure 7.14 If copying from the Snap Data Volume to a disk other than the copy source
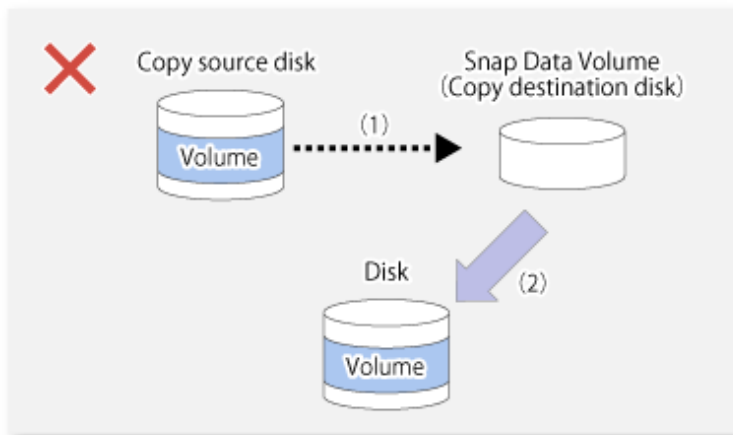
7.2.3.2 Creating a Snap Data Volume

Create a Snap Data Volume by performing the following steps:

1. Calculate the physical capacity for the Snap Data Volume.
2. Define and initialize the Snap Data Volume.
3. Connect to the host.
4. Create partitions and file systems.

Calculate the physical capacity for the Snap Data Volume

The formula for estimating the physical capacity of Snap Data Volume is as follows:

```
Physical capacity = (number of updated blocks for the copy source volume) x (safety factor)
```

Strictly speaking, both the number of blocks that have been updated on the Snap Data Volume and the management area used by the hardware (about 0.1% of the logical capacity) should also be taken into account before the SnapOPC/SnapOPC+ is taken, however this can be covered by increasing the safety factor.

The number of blocks that have been updated on the copy source volume can be estimated using swstestupdate (Update volume measurement command).

Make this measurement using the following procedure:

1. Start measuring the update amount by configuring a pseudo session on the copy source volume.

    ```
    # /opt/FJSVswsts/bin/swstestupdate start /dev/sda1
    /dev/sda1 swstestupdate completed
    #
    ```

2. Start transactions. Blocks that are updated by transactions will be recorded on the hardware.

3. After the measurement period has elapsed, check the number of blocks that have been updated.

    ```
    # /opt/FJSVswsts/bin/swstestupdate status /dev/sda1
    Volume-Name Update
    /dev/sda1   644333
    #
    ```

4. After the measurements are complete, cancel the pseudo session.

    ```
    # /opt/FJSVswsts/bin/swstestupdate stop /dev/sda1
    /dev/sda1 swstestupdate completed
    #
    ```

Define and initialize the Snap Data Volume

Define and initialize the Snap Data Volume using ETERNUS Web GUI. At this point, set the logical capacity to that of the copy source disk (in order to use the same partition configuration for both the copy source disk and the copy destination disk).

Connect to the host

Connect the created Snap Data Volume to the host. Refer to the manual for the disk array device (the Server Connection Guide) for information on this procedure.

Create partitions and file systems

Create partitions and file systems so that a copy source volume can be created on the Snap Data Volume.

📝 Note

··

Once file systems have been created, limit updates to the copy destination volume as much as possible, to avoid unnecessarily using up the physical capacity of the Snap Data Volume.

··

7.2.3.3 Creating a Snap Data Pool

The Snap Data Pool is created using ETERNUS Web GUI.
Note the following points when creating the Snap Data Pool.

🗒 Note

- In the basic operation, copying is completed inside the Snap Data Volume capacity estimated/calculated as appropriate for the copying capacity. The Snap Data Pool is only an auxiliary area; when copying, always make sure that area is added from the Snap Data Pool to the Snap Data Volume.

- Estimate the updated amount for each copy source volume, and then decide the Snap Data Volume and the Snap Data Pool capacity. If an advance estimate is difficult, it is recommended that the Snap Data Volume capacity + the Snap Data Pool capacity is between 30-50% of the total copy source capacity (since this is only a recommended value, the Snap Data Pool capacity must be changed in accordance with the operation status).

7.2.3.4 Snap Data Volume/Snap Data Pool monitoring

When there is insufficient Snap Data Volume or Snap Data Pool capacity, it is not possible to access copy destination volume (the Snap Data Volume) (this status makes it impossible to read copy destination data or write data to the copy destination). In SnapOPC+, it is not possible to write the copy destination or copy destination generations before it.

For this reason, monitoring must be performed so that the Snap Data Volume or the Snap Data Pool capacity is not insufficient at the time of operations.

- Using only the Snap Data Volume (when the Snap Data Pool is not set)
Specify "stat" as swstsdv (Snap Data Volume operation/reference command) subcommand, and regularly monitor whether there is insufficient Snap Data Volume capacity.

- Using the Snap Data Pool (when the Snap Data Pool is set)
Specify "poolstat" as swstsdv (Snap Data Volume operation/reference command) subcommand, and regularly monitor whether there is insufficient Snap Data Pool capacity.

  If there is an area in the Snap Data Pool for which the usage rate exceeds 50%, regardless of encryption, consider increasing the Snap Data Pool. Additionally, increase the Snap Data Pool monitoring frequency.

  If there is an area in the Snap Data Pool for which the usage rate exceeds 70%, regardless of encryption, increase the Snap Data Pool immediately.

When there is insufficient Snap Data Volume or Snap Data Pool capacity, refer to "8.4.2.3 Troubleshooting when a lack of free space has occurred in the Snap Data Volume or Snap Data Pool".

Creating an access volume

When there is insufficient Snap Data Volume or Snap Data Pool capacity, it is not possible to access copy destination volume (Snap Data Volume).
In the following cases, it is not possible to access the Snap Data Volume copy destination volume either:

- Access to the Snap Data Volume after the Snap Data Volume initialization

- Access to the Snap Data Volume after the Disk management information initialization

For this reason, it is recommended that you reserve a volume (access volume) that is separate from the Snap Data Volume.
The volume that is used as the access volume is created on the same ETERNUS Disk storage system as the Snap Data Volume. Since this is not used in the storage of data, there is no problem even if there is little capacity. Any size can be used for the creation of the access volume.

The following operations can be executed using the access volume:

- Snap Data Volume initialization

- Snap Data Volume status display

7.2.3.5 Cluster Operations

Do not use Snap Data Volume as shared disks in a clustered system. This is to avoid the danger of the clustered system failing over when the physical capacity of the Snap Data Volume is insufficient.
In cluster operations, use the following methods to prevent the Snap Data Volume from being used as a shared disk:

- Permit the Snap Data Volume to be referenced from all the nodes in the clustered system.

- Use inter-server replication between the clustered system and a non-clustered system.

7.3 Flow of Operations

The following figure shows the flow of the replication.

Figure 7.15 Flow of replication operations

7.4 Preparations

This section describes the preparations that are required before performing replication.

7.4.1 Starting a daemon

Before replication is started, the AdvancedCopy Manager daemon must be started on the Management Server and Managed Server. Normally, the daemon is configured to be automatically invoked at system startup.

However, if a daemon fails to start for some reason or has stopped, both servers must restart their respective daemons. For information about starting daemons, refer to "Chapter 2 Starting and Stopping Daemons".

7.4.2 Starting the Web Console

Start the Web Console. For details, refer to "How to start Web Console" in the "ETERNUS SF Web Console Guide".

7.4.3 Registering a Managed Server

The Management Server registers the Managed Server to be managed.
This process can be executed using Web Console.

7.4.4 Fetching device information on a Managed Server

This process can be executed using either Web Console or stgxfwcmsetdev (Device information collection/reflection command).

📑 Note

- Device information on the Managed Server must briefly be stored in the repository.
  The area required when updating the repository may not be available when information concerning devices under multiple Managed Servers is fetched. Therefore, check if there is sufficient area for updating the repository before fetching the device information. If there is insufficient area, allocate the required area before fetching the device information.

  Refer to Step 2 in "8.2.3 Troubleshooting: Failure in accessing a repository database" for details on checking the required area.

- The time required for this operation depends on the total number of devices defined on the Managed Server. If the number of devices is large, perform this operation while the CPU load and I/O load are both low. As a guideline, each operation takes about 0.5 seconds per device (or partition) under no-load conditions.

7.4.5 Configuring the source volume and replica volume

Use swsrpsetvol (Replication volume information setting command) to set the source volume and replica volume that will be used in the replication. The information on the registered source volume and replica volume can be checked with swsrpvolinfo (Replication volume information display command).

When specifying the volume of another Managed Server in server-to-server replication, specify the volume in the format "*Volume-Name@Managed-Server-name*".

7.4.6 Customizing pre-processing and post-processing

In AdvancedCopy Manager, replication processing must be performed while volumes are unmounted. Therefore, when replication is executed, processing to mount and unmount the volume is performed by the pre-processing and post-processing scripts respectively.

The replication processing is not executed when the volume cannot be unmounted.

For details of the pre-processing and post-processing, refer to "Appendix C Pre-processing and Post-processing of Replication".

Replication pre-processing and post-processing need to be customized when operational requirements change.

For example:

- For specific operational requirements, you want to avoid the unmount/mount processing of the volume where a file system is constructed.

- If you want to add special pre-processing or post-processing.

Refer to "Appendix C Pre-processing and Post-processing of Replication" for information on the customization.

Note

- The pre-processing and post-processing are not performed in case of replicating a group.
  Unmount all the volumes in the group before replication, and mount all the volumes after replication.

- Pre-processing and post-processing are not executed for the partitions included in a Logical Unit (disk) when Logical Unit (disk) unit copying is executed. Therefore, the partitions included in a Logical Unit (disk) must be unmounted before a replica is created, and the unmounted partitions must be mounted after the replica is created.

Point

When AdvancedCopy Manager is upgraded

The script may need to be updated after an upgrade of AdvancedCopy Manager.
Therefore, you should check the script after any upgrade to ensure it is compatible with newer versions.

7.4.7 Creating groups

7.4.7.1 Conditions for the replication volumes that make up the group

The replication volumes that make up the group (the source volume and the destination volume pairs) must meet the following conditions:

- The copy source servers and destination servers must have matching attributes such as size and device configuration.

- The operational servers and the copy directions must all match. The values for the operational servers and the copy directions can be checked using the **Op-Server** and **Copy** fields displayed by executing swsrpvolinfo (Replication volume information display command).

- The replication volume information being registered must not be registered in any other group.

- Source volumes and destination volumes must not be duplicated within the group.

- A replication group cannot have both GDS logical volumes and other types of volumes.

- If a Logical Unit (disk) is specified, the information of the replication volume in a partition included in the Logical Unit (disk) must not be registered to another group.

- If a Logical Unit (disk) is specified, the information of the replication volume in a partition included in the Logical Unit (disk) must not be registered within the same group.

7.4.7.2 How to create groups

Groups are created using swsrpsetvol (Replication volume information setting command).
Information for created groups can be displayed using swsrpvolinfo (Replication volume information display command).

Example

Two groups (called GRP1) consisting of a source volume and a replica volume are created.

```
# /opt/FJSVswsrp/bin/swsrpsetvol -Xgroup GRP1 /dev/sdb1@SRC /dev/sdb1@TARG-1
swsrpsetvol completed
# /opt/FJSVswsrp/bin/swsrpsetvol -Xgroup GRP1 /dev/sdc1@SRC /dev/sdc1@TARG-1
swsrpsetvol completed
# /opt/FJSVswsrp/bin/swsrpvolinfo -L
Server Original-Volume Size      Replica-Volume    Size      Copy        Op-
Server Group
SRC    /dev/sdb1@SRC   4.0 Gbyte /dev/sdb1@TARG-1 4.0 Gbyte bi-direction
```

```
both       GRP1
SRC     /dev/sdc1@SRC    4.0 Gbyte /dev/sdc1@TARG-1 4.0 Gbyte bi-direction
both       GRP1
#
```

7.5 Operations

Before performing replication, refer to "7.4 Preparations" to set up the environment required for replication.

This section describes the operation by command.

7.5.1 Executing snapshot replication

Use swsrpmake (Replication creation command) to perform snapshot replication.
Refer to "7.1.1 Snapshot replication processing" for an explanation of snapshot replication.

The operation status of a physical copy can be checked by executing swsrpstat (Operation status display command).

QuickOPC type replication

Execute QuickOPC replication by specifying the -T option in swsrpmake (Replication creation command).
If no OPC session exists when swsrpmake (Replication creation command) is executed, the command starts snapshot processing (OPC physical copying), and tracks processing from the source volume to the destination volume.

Figure 7.16 When replication creation command is executed (first time)

To check the execution status of physical copying, use swsrpstat (Operation status display command) in the same way as for an ordinary snapshot replication.
After snapshot processing (OPC physical copy) is complete, only tracking processing is active.
To check the tracking status, use swsrpstat (Operation status display command) with the -L option.

Figure 7.17 When snapshot processing is completed

Entering swsrpmake (Replication creation command) with the -T option specified during tracking processing performs the physical copying of only the data that has been generated since the previous snapshot processing. This means that physical copying can be accomplished in a shorter period of time.

Figure 7.18 When replication creation command is executed (second time)

When you want to perform a restoration while tracking processing is being executed, you need to perform a restoration by OPC (to achieve this, you need to execute swsrpmake (Replication creation command) without the -T option). QuickOPC cannot be executed in the reverse direction while tracking processing is being executed. The replication using QuickOPC is done as follows:

```
[backup]
swsrpmake -T <original volume name> <replica volume name>


[restore]
swsrpmake <replica volume name> <original volume name>
```

Although a restoration is executed with OPC, only the data that has been updated since the previous replication (it can be obtained from the **Update** field of swsrpstat) is copied.

Therefore, in replication using QuickOPC, not only a physical backup but also restoration is completed in a short period of time.

The restore execution status can be checked by executing swsrpstat (Operation status display command) with the -E option specified.

SnapOPC type replication

Execute SnapOPC type replications with the -C option specified in swsrpmake (Replication creation command).

When swsrpmake (Replication creation command) is executed, a SnapOPC session will be set up between the copy source volume and the copy destination volume.

Example

```
# /opt/FJSVswsrp/bin/swsrpmake -C /dev/sda1 /dev/sdc1
FROM=/dev/sda1@SV1,TO=/dev/sdc1@SV1 swsrpmake completed
#
```

Figure 7.19 When the replication creation command is executed

Unlike normal OPCs and QuickOPCs, SnapOPCs do not copy all of the data from the source volume, but instead copy only the data that has been updated on the source or destination since SnapOPC started. This kind of copy processing is referred to as "Copy-on-Write".

Figure 7.20 When the copy source volume is updated

Figure 7.21 When the copy destination volume is updated

Note: The units for host I/O and storage device copies are different (512 bytes for host I/O and 8 kilobytes for storage device copies), and therefore data copies also occur when the copy destination is updated.

The status of SnapOPC sessions can be checked using swsrpstat (Operation status display command).

The following example shows the execution of swsrpstat (Operation status display command) immediately after a SnapOPC snapshot has started. While SnapOPC is being performed, "copy-on-write" is displayed in the Status field, and the amount of data updated since the last copy was created is displayed in the Update field as a percentage.

Example

```
# /opt/FJSVswsrp/bin/swsrpstat -L /dev/sdb1
Server Original-Volume Replica-Volume Direction Status        Execute Trk  Update Rcv  Split Xfer
Snap-Gen
SV1   /dev/sdb1@SV1  /dev/sbc1@SV1 regular   copy-on-write ----   off  0%   ---- ----  ---- ----
#
```

If swsrpmake (Replication creation command) is executed again during SnapOPC processing, the SnapOPC session that has already been set up will be cancelled, and a new session will be set up.

🏮 Note

. . . .

When there is insufficient Snap Data Volume or Snap Data Pool capacity, the SnapOPC+ execution status changes to error suspend status ("failed"), and replication volume cannot be used.

The SnapOPC execution status can be checked in swsrpstat (Operation status display command) output result Status column.

Example

```
# /opt/FJSVswsrp/bin/swsrpstat -L /dev/sdb1
Server Original-Volume Replica-Volume Direction Status Execute Trk  Update Rcv  Split Xfer Snap-Gen
SV1    /dev/sdb1@SV1   /dev/sdc1@SV1  regular   failed ----   off  ----   ---- ----  ---- ----
#
```

When the SnapOPC execution status is error suspend status ("failed"), refer to "8.4.2.3 Troubleshooting when a lack of free space has occurred in the Snap Data Volume or Snap Data Pool".

- -

Perform restorations from Snap Data Volume by running an OPC using swsrpmake (Replication creation command).

```
# /opt/FJSVswsrp/bin/swsrpmake /dev/sdc1 /dev/sdb1
FROM=/dev/sdc1@SV1,TO=/dev/sdb1@SV1 swsrpmake completed
#
```

When restorations are executed, the SnapOPC session from the source volume to the destination volume is maintained as is, and a normal OPC from the replication destination volume to the replication source volume is started. At this point, the time taken to restore the physical copy is reduced, because only data that has been updated since the last copy is restored.

Figure 7.22 When restoration is executed
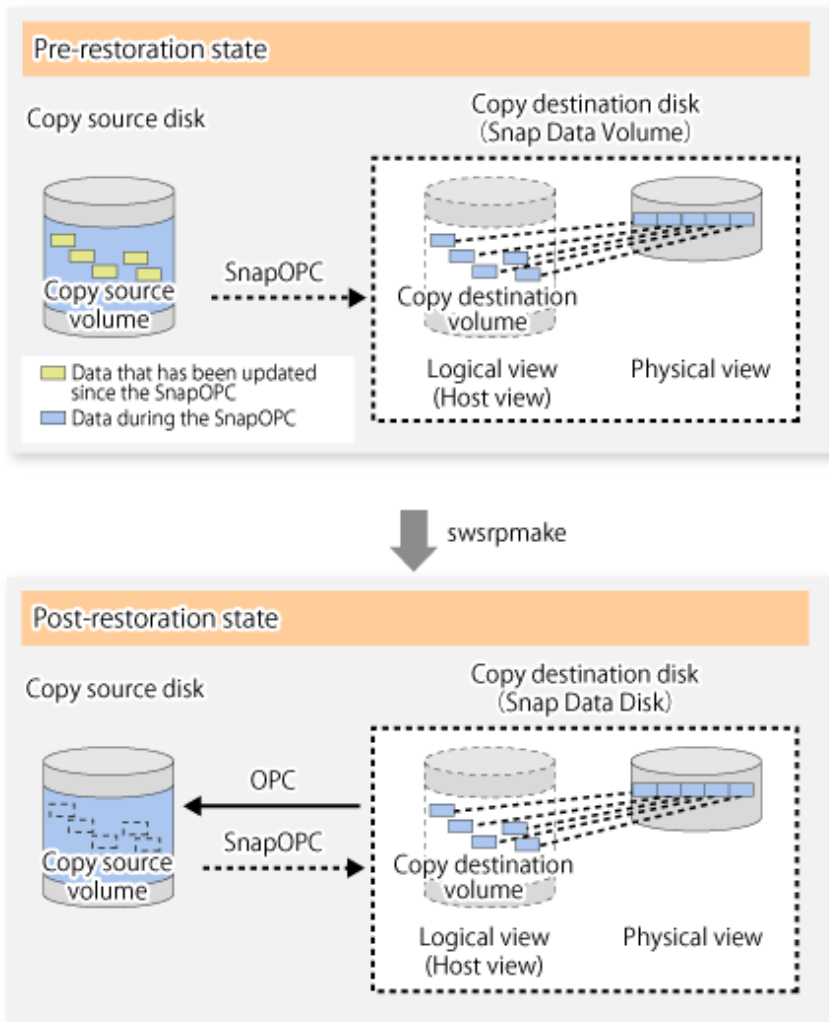
The execution status of restorations can be checked by specifying the -E option with swsrpstat (Operation status display command).

```
# /opt/FJSVswsrp/bin/swsrpstat -E /dev/sda1
Server Original-Volume Replica-Volume Direction Status Execute
```

```
SV1    /dev/sda1@SV1   /dev/sdb1@SV1  reverse   snap   80%
#
```

> **Note**
> ....
> If a SnapOPC is being performed between the source volume and the destination volume, restorations to volumes other than the source volume cannot be executed. To restore to a volume other than the source volume, operating system copy functions (such as the cp command or the copy command) must be used.

Figure 7.23 When restoring to a volume other than the copy source volume

Additionally, if SnapOPCs are being performed to multiple copy destination volumes, restoration cannot be performed.

Figure 7.24 When SnapOPC is performed for multiple copy destination volumes

In this case, restoration using an OPC can be performed by cancelling the other SnapOPCs. However, the backup data on the copy destination volumes whose SnapOPC sessions were cancelled will be lost.

Figure 7.25 When SnapOPC session cancelled to perform restoration

To perform a restoration while still maintaining all SnapOPC sessions, operating system copy functions (such as the cp command or the copy command) must be used for the restoration.

However, if restoration is performed using operating system functions, the amount of updated data on the source volume will increase, and there is a risk that the capacity of the SnapOPC volume will be insufficient.

Figure 7.26 When performing restoration without cancelling SnapOPC session

SnapOPC+ type replication

Execute swsrpmake (Replication creation command) using the -P option to perform SnapOPC+ replication. This sets a SnapOPC+ session between the copy source volume and the copy destination volume. After the session is set, copy-on-write is performed between the copy source volume and the copy destination volume.

An example of executing swsrpmake (Replication creation command) using the -P option is shown below.

Execution example

```
# /opt/FJSVswsrp/bin/swsrpmake -P /dev/sdb1@SV1 /dev/sdc1@SV1
FROM=/dev/sdb1@SV1,TO=/dev/sdc1@SV1 swsrpmake completed
#
```

At this time, the (logically copied) copy destination volume is saved as a snap generation number.

The next time this command is executed with a different copy destination volume for the same copy source volume, the copy-on-write processing being executed between the copy source volume and the previous generation of the copy destination volume is stopped. Then, a SnapOPC+ session is set between the copy source volume and the newly specified copy destination volume, and copy-on-write is performed.

An example of executing swsrpmake (Replication creation command) using the -P option for the newly specified copy destination volume is shown below.

Execution example

```
# /opt/FJSVswsrp/bin/swsrpmake -P /dev/sdb3@SV1 /dev/sdd3@SV1
FROM=/dev/sdb3@SV1,TO=/dev/sdd3@SV1 swsrpmake completed
#
```

This time, the (logically copied) copy destination volume is saved as snap generation number 2.

Similarly, each time there is a new copy destination volume, a snap generation number is assigned.

![Note icon] **Note**

If an earlier snap generation (other than the oldest snap generation) is specified as the copy destination volume when swsrpmake (Replication creation command) is executed, the command terminates with an error. If the oldest snap generation is specified as the copy destination volume, that snap generation is automatically discarded and a replica is created as the newest snap generation. In this case, subsequent snap generations (second, third) are assigned a snap generation number that is one generation prior (second generation => first generation, and third generation => second generation).

Figure 7.27 When the oldest snap generation number is specified as the replication volume

The operation status of SnapOPC+ replication can be checked by executing swsrpstat (Operation status display command) with the -L option.

For the most recent snap generation, "copy-on-write(active)" is displayed in the **Status** field. For past snap generations, "copy-on-write(inactive)" is displayed. In the **Update** field, the amount of data that has finished being updated after replication creation, is displayed as a percentage. In the **Snap-Gen** field, the snap generation number is displayed.

Execution example

```
# /opt/FJSVswsrp/bin/swsrpstat -L /dev/sdb3
Server Original-Volume  Replica-Volume Direction  Status                 Execute Trk Update Rcv
Split Xfer Snap-Gen
```

```
SV1    /dev/sdb3@SV1    /dev/sdc3@SV1   regular    copy-on-write(inactive) ----    off 0%      ----
----  ---- 1
SV1    /dev/sdb3@SV1    /dev/sdd3@SV1   regular    copy-on-write(active)   ----    off 5%      ----
----  ---- 2
#
```

📒 Note

. . . .

When there is insufficient Snap Data Volume or Snap Data Pool capacity, the SnapOPC+ execution status changes to error suspend status ("failed"), and the execution status of SnapOPC+ that was executed before it will also change to error suspend status ("failed"). Replication volume of error suspend status ("failed") cannot be used.

The SnapOPC+ execution status can be checked in swsrpstat (Operation status display command) output result **Status** field.

Execution example

```
# /opt/FJSVswsrp/bin/swsrpstat -L /dev/sdb3
Server Original-Volume   Replica-Volume Direction Status Execute Trk Update Rcv  Split Xfer Snap-Gen
SV1    /dev/sdb3@SV1     /dev/sdc3@SV1   regular    failed ----    off ----   ---- ----  ---- ----
SV1    /dev/sdb3@SV1     /dev/sdd3@SV1   regular    failed ----    off ----   ---- ----  ---- ----
#
```

When the SnapOPC+ execution status is error suspend status ("failed"), refer to "8.4.2.3 Troubleshooting when a lack of free space has occurred in the Snap Data Volume or Snap Data Pool".

. . . .

To restore from a Snap Data Volume, use swsrpmake (Replication creation command) to start OPC.

```
# /opt/FJSVswsrp/bin/swsrpmake /dev/sdd3@SV1 /dev/sdb3@SV1
FROM=/dev/sdd3@SV1,TO=/dev/sdb3@SV1 swsrpmake completed
#
```

The SnapOPC+ session from the replication source volume to the replication destination volume is maintained even if the replication creation command is executed.

Execution of restoration while maintaining the SnapOPC+ session reduces the physical copying time, because physical copying is performed only for data updated after the replica creation.

Figure 7.28 Restoration with SnapOPC+ session maintained

To check the restoration execution status, execute swsrpstat (Operation status display command) with the -E option.

```
# /opt/FJSVswsrp/bin/swsrpstat -E /dev/sdb3
Server Original-Volume Replica-Volume Direction Status Execute
SV1    /dev/sdb3@SV1   /dev/sdc3@SV1  ----      ----   ----
SV1    /dev/sdb3@SV1   /dev/sdd3@SV1  reverse   snap   80%
#
```

> **Note**
> ....
>
> Restoration may cause a Snap Data Pool to run low on free disk space, due to updates issued to the most recent snap data generation where the copy-on-write status is active. Make sure that there is enough free space in the Snap Data Pool usage area before performing restoration.
>
> The most recent snap generation is the data written to the replication source volume by the restoration, updated by the previously existing data. The update amount to the most recent snap generation generated by the restoration is the total of the Copy usage amount for the restoration target snap generation and subsequent snap generations except for the most recent snap generation.

An example of how to calculate the update amount when restoring from snap generation (Snap-Gen) 2 is displayed below.

Use the procedure below to check the update amount for restoration:

1. Use swsrpstat (Operation status display command) to check the device name of the restoration target and subsequent snap generations, except for the most recent snap generation (Snap-Gen 4 data in the example below).

```
# /opt/FJSVswsrp/bin/swsrpstat -L /dev/sdb3
Server Original-Volume Replica-Volume Direction Status                Execute Trk Update Rcv
Split Xfer Snap-Gen
SV1    /dev/sdb3@SV1   /dev/sdc3@SV1   regular   copy-on-write(inactive) ----    off 8%     ----
----  ---- 1
SV1    /dev/sdb3@SV1   /dev/sdd3@SV1   regular   copy-on-write(inactive) ----    off 12%    ----
----  ---- 2
SV1    /dev/sdb3@SV1   /dev/sde3@SV1   regular   copy-on-write(inactive) ----    off 0%     ----
----  ---- 3
SV1    /dev/sdb3@SV1   /dev/sdf3@SV1   regular   copy-on-write(active)   ----    off 3%     ----
----  ---- 4
```

   In this example, /dev/sdd3 and /dev/sde3 are targeted.

2. Use swstsdv (Snap Data Volume operation/reference command) with the "stat" subcommand to find the total copy usage amount for the device in the previous step.

   If a Snap Data Pool is used, also add the Snap Data Pool usage capacity.

   [/dev/sdd3 disk usage]

```
# /opt/FJSVswsts/bin/swstsdv stat /dev/sdd3
BoxID = 00E4000M3#####E450S20A####KD4030639004##
LUN = 110 (0x6E)
Rate Logical(sector) Physical(sector) Used(sector) Copy(sector) Host(sector) Pool(sector)
100% 8388608        1048576          1048576      1048384      192          640
```

   [/dev/sde3 disk usage]

```
# /opt/FJSVswsts/bin/swstsdv stat /dev/sde3
BoxID = 00E4000M3#####E450S20A####KD4030639004##
LUN = 111 (0x6F)
Rate Logical(sector) Physical(sector) Used(sector) Copy(sector) Host(sector) Pool(sector)
4%   8388608        1048576          46928        16           46912        0
```

   In this example, the quantity updated by the restoration is 1049040 (1048384+640+16) sectors.

To check the Snap Data Pool total capacity and the usage area capacity, use swstsdv (Snap Data Volume operation/reference command) with the "poolstat" subcommand.

If the Snap Data Volume is not encrypted, then check the capacity of the usage area and the total capacity where Pool-Type is Normal. Otherwise, check the capacity of the usage area and the total capacity where Pool-Type is encrypted.

```
# /opt/FJSVswsts/bin/swstsdv poolstat -G /dev/sdb3
BoxID = 00E4000M3#####E450S20A####KD4030639004##
Pool-Type Rate Total(sector) Used(sector) Copy(sector) Host(sector) Free(sector)
```

```
Normal     10%  20971520    2097152        0           2097152      18874368
Encrypted  0%   20971520    0              0           0            20971520
```

The disk usage in this example is 15% =~ (2097152 + 1049040) / 20971520 x 100

If the value obtained by adding the size of the Snap Data Pool usage area to the restoration update amount is less than the total capacity, then restoration is possible. However, in order to safely perform restoration, it is recommended to extend the Snap Data Pool if the disk usage after restoration is predicted to exceed 70%.

In addition, if the disk usage is expected to exceed 50%, then consider extending the Snap Data Pool after restoration and increasing the monitoring frequency of the Snap Data Pool.

For details on Snap Data Pool monitoring, refer to "7.2.3.4 Snap Data Volume/Snap Data Pool monitoring".

Note

If SnapOPC+ is being performed between the replication source volume and the replication destination volume, restoration cannot be performed to a volume other than the replication source volume.

Point

As a precaution against hardware malfunctions with SnapOPC+, it is recommended to operate it in conjunction with making full copies using OPC/QuickOPC/EC(REC).
An example of performing QuickOPC on Sundays and SnapOPC+ on Mondays to Saturdays is displayed below.

Figure 7.29 Example of operation using SnapOPC+ and QuickOPC

7.5.1.1 Concurrent OPC functions

Concurrent OPC function is one of ETERNUS Disk Storage system's function which creates snapshots of the multiple logical volumes simultaneously. This function allows backup the replication of the database consisting of multiple volumes in a consistent state. This function is available to use in one of these replications; OPC, QuickOPC, SnapOPC, or SnapOPC+.

The following diagram shows the operation in ETERNUS Disk Storage system.

Figure 7.30 Operations in ETERNUS Disk Storage system

> **Note**
>
> - If you force-quit swsrpcancel (Replication cancellation command) with task manager or Ctrl-C, the session may remain defined concurrent OPC. In this case, execute the command again, or stop the session with swsrpcancel (Replication cancellation command).
>
> - The maximum number of copy pairs that performs concurrent OPC simultaneously is 1024.

7.5.2 Executing synchronized replication

To perform synchronized replication, use the following procedure:

1. Start synchronized processing using swsrpstartsync (Synchronous processing start command). Use swsrpcancel (Replication cancellation command) to cancel synchronized processing that has already started.

2. After making sure that equivalency maintenance state has been established with swsrpstat (Operation status display command), temporarily stop synchronized processing with swsrpmake (Replication creation command) to create a replica of the source volume.

3. To copy the updated (ie, incremental) data, restart synchronized processing with swsrpstartsync (Synchronous processing start command).

7.5.2.1 Intra-box synchronous replication

Intra-box synchronous replication creates a replication from a source volume to a destination volume by using the EC function of ETERNUS Disk storage system.
The EC function operates in a mode in which a copy is made to a destination volume in synchronization through a write to a source volume (ie, in synchronous mode).
It is also possible to reverse the copy direction from suspend status using the EC functionality.

7.5.2.2 Inter-box synchronous replication

Inter-box synchronous replication creates a replication from a source volume to a destination volume by using the REC function of ETERNUS Disk storage system.

The REC function provides three copy operation modes that can be selected for operation:

- Transfer mode

- Recovery mode

- Split mode

In addition, the REC function enables the copy direction to be reversed in suspended status.

Transfer mode

The transfer mode provides the REC data transmission modes described below.

Table 7.2 Transfer mode types

Transfer method	Transfer mode	Explanation
Synchronous transfer method	Synchronous	In this mode, responses to write requests from the server do not return as "complete" until data has been written to the copy source volume and copied to the destination volume.
		Synchronizing data copies with writes to the copy source guarantees the integrity of the data on both the copy source volume and the copy destination volume when the copy completes.
		This mode is suitable to RECs within sites with short delay times, because it has a large impact on the performance of write accesses from the server.
Asynchronous transfer methods	Stack	This mode has only a small impact on the time taken to respond to the server because it only records the position of the block that has been updated before returning "complete" to the server. The data for this block is transferred using an independent transfer engine.
		Copies can be executed even when there is only limited bandwidth, but in this case the amount of data that has not been transferred may build up.
	Consistency	This mode guarantees the sequence of transfers between copy sessions for the copy destination ETERNUS Disk storage system, and is suitable for operations, such as on databases, that perform mirroring using copies that consist of multiple areas.

Transfer method	Transfer mode	Explanation
		This mode uses part of cache memory for transfer buffers. Data is transferred to the copy destination by first "storing" a block of data to be copied in a transmission transfer buffer and then using a reception transfer buffer to "expand" the data that has been transferred.
	Through	This mode is for transferring data that has not been transferred when copies in Stack mode or Consistency mode are stopped or suspended.

To perform a synchronous replication in Stack mode or Consistency mode, use swsrpstartsync (Synchronous processing start command), swsrpmake (Replication creation command), and swsrpchsync (Synchronization mode change command). To check the execution status, use swsrpstat (Operation status display command). The figures below show how synchronous replication is performed in Stack mode or Consistency mode.

Figure 7.31 Flow of synchronous replication (for Stack mode/Consistency mode)

> 🏳 **Note**
> ....
> If the equivalency maintain status REC operating mode is changed, after the operating mode is changed, the user must wait until the equivalency maintain status is attained.
> ....

Recovery mode

Recovery mode consists of two methods of restarting the copying after recovery is performed from an inter-box path error (e.g., halt status).

Table 7.3 Recovery mode types

Mode	Description
Automatic Recovery	In this mode, the REC session automatically switches from halt status to regular status, and copy processing resumes when the inter-box FCRA path is recovered.
Manual Recovery	In this mode, the REC session remains in halt status and copy processing does not resume even if the inter-box FCRA path is recovered. Manually resume copying. This mode is used, for example, when operating a standby database.

Split mode

Split mode includes two modes for performing a write operation to the source volume when REC is used for synchronous transfer mode, and for recovery from an inter-box path error (halt status).

Table 7.4 Split mode types

Mode	Description
Automatic Split	This split mode forcibly executes successful write operations to source volumes even if the inter-box FCRA path is fully blocked and halt status occurs. This mode enables write operations to source volumes even if the inter-box FCRA path is fully blocked, thus this mode does not affect transactions. When the inter-box FCRA path is recovered, copy processing resumes according to the recovery mode configurations.
Manual Split	This split mode rejects write operations to source volumes (it returns an error) if the inter-box FCRA path is fully blocked and halt status occurs. This mode enables source volumes and destination volumes to be fully synchronized even if the FCRA path is fully blocked.

Mode	Description
	When the inter-box FCRA path is recovered, copy processing resumes according to the recovery mode configurations.

Copy direction reverser

The copy direction reverser allows you to smoothly switch between different sites.

The following figures show an example of how to switch the copy direction:

1. Assume that Site A is operating and REC is operating from Site A to Site B.

Figure 7.32 When REC performed from Site A to Site B

2. To switch the sites, execute swsrpmake (Replication creation command) to make a replication to Site B. Then, stop operating Site A.

Figure 7.33 When replication creation command used to create a replica at Site B
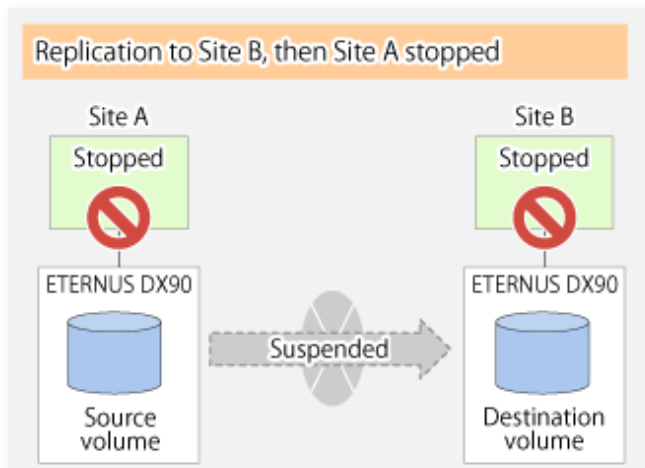
3. Execute swsrprevsync (Synchronization reverse command) to reverse the copy direction.

Figure 7.34 When the synchronization reverse command is executed

4. Bring Site B into operation. At this stage, synchronization is still suspended, thus any update to the volume at Site B is not reflected at Site A.

Figure 7.35 When operation is started at Site B (synchronization suspend state)

5. Start (resume) synchronization from Site B to Site A. Updates made to the volume in Site B while synchronization is suspended are reflected in Site A with differential copies.

Figure 7.36 When operation is started at Site B (synchronization resumed)

7.5.2.2.1 Initial copy skip function

The initial copy skip function is used when the initial copy cannot be executed because of insufficient line capacity.
The following example shows the initial copy being skipped via tape transfer.

1. Suppose that operations at Site A have stopped.

Figure 7.37 If operations at Site A have stopped

2. Next, synchronous processing begins, using the initial copy skip function. At this point, a REC session is set up, and the status is Replication Established. Data is not copied to the copy destination volume.

Figure 7.38 When synchronization is started by the initial copy skip function

3. Next, the data on the copy source volume is backed up to tape.

Figure 7.39 When copy source data is backed up to tape

4. The tape medium is sent to Site B and jobs at Site A restart.

Figure 7.40 When Site A operations restart

5. The data on the tape medium is restored to the copy destination volume. At this point, the data on the copy destination volume is the same as the data on the copy source volume that existed before operations restarted.

Figure 7.41 When data is restored from tape to the copy destination

6. Synchronous processing restarts in Remain mode. Restarting synchronous processing in Remain mode means that only data that has been updated on the copy source volume is reflected on the copy destination volume. If Remain mode is not used, all data on the copy source volume is copied.

Figure 7.42 When synchronization is restarted in Remain mode

7.5.2.3 Concurrent suspend function

The concurrent suspend function simultaneously suspends multiple EC/REC sessions for ETERNUS Disk storage system.
By using this function, a copy with consistency can be easily taken.
For example, this is useful for a database composed of multiple volumes.
The behavior that takes places within the ETERNUS Disk storage system is shown below.

Figure 7.43 Internal operation of ETERNUS Disk storage system

Replication using the concurrent suspend function is performed by specifying the -Xconcur option for swsrpmake (Replication creation command).

Additionally, if concurrent suspend is performed using the Consistency transfer mode, it is no longer necessary to make temporary mode changes during multiple creation processes. (Refer to the figure below). Accordingly, the operating procedure used when concurrent suspend is executed in Consistency mode is the same as the procedure for Through mode and Synchronous mode.

Figure 7.44 When Consistency mode used

7.5.2.3.1 Changing the timeout value for the concurrent suspend function

A timeout may occur before suspend is completed if there is a large quantity of data, if the communication link is poor, or if the REC Disk buffer is used for inter-box synchronous replication (REC function) in Consistency mode.
If the REC Disk buffer is used, refer to "7.5.2.3.2 Concurrent suspend when the REC Disk buffer is used".
The following message is shown if it timeout.

```
swsrp2709 Concurrent suspend failed. Session ID=[-], Reason=[timeout]
```

If this message is shown, change the value of timeout by creating a file which sets the timeout value. Create a file which sets the timeout value for the concurrent suspend function under the following name. The file must be created on the same server that contains the source volume.

In case of normal operation (non-clustered environment)

```
/etc/opt/FJSVswsrp/data/DEFAULT/check.ini
```

In case of operation in clustered environment

```
/etc/opt/FJSVswsrp/<logical node>/data/DEFAULT/check.ini
```

Table 7.5 File format of the check.ini

Section name	Key name	Value
[check]	ConcurCheckTimeout	Set the timeout value for the concurrent suspend function in units of seconds. Setting range: From 1 to 2147483647 sec. If you set a value beyond this range, the default value 60 sec will be applied.

Example

Setting the timeout value for the concurrent suspend function to 120 seconds.

```
[check]
ConcurCheckTimeout=120
```

If the setting file for the concurrent suspend function is not created, the default value (60 seconds) will be applied.

7.5.2.3.2 Concurrent suspend when the REC Disk buffer is used

If the REC Disk buffer is used for inter-box synchronous replication (REC function) in Consistency mode, check that no transferred data exists in the REC Disk buffer before executing the concurrent suspend.
If you execute the concurrent suspend when any data to be transferred exists in the REC Disk buffer, timeout may occur.

You can check the data volume of the REC Disk buffer in the **DiskRate** field when you execute swsrprecbuffstat (REC buffer information display command) with specifying -L option.

7.5.2.4 Destination Access Permission function

The Destination Access Permission function serves to inhibit user access to replication target volumes when copying is in progress and any access would therefore be inappropriate.

It starts synchronization (EC/REC) from the replication source volume to the replication target volume, and when full or differential copying is complete, they are maintained in a state of equivalency.
In the state of maintained equivalency, any upgrades on the replication source volume are successively reflected on the replication target volume, so that any access to the replication target volume would possibly be inappropriate.
The Destination Access Permission function can be used to inhibit access from users to the replication target in such unsuitable moments.

To set the Destination Access Permission function, execute swsrpstartsync (Synchronous processing start command) with specifying the -Xda option. Until replication processing is stopped by swsrpcancel (Replication cancellation command) neither WRITE nor READ access to the replication target volume can be performed.

For changing this setting, once the replication operation has been stopped by swsrpcancel (Replication cancellation command), execute swsrpstartsync (Synchronous processing start command) again.

Whether access is inhibited can be checked by performing swsrpstat (Operation status display command).

7.5.3 Backup and restoration using the replication function

Data can be restored from the replica volume to the original volume if a volume pair has been defined with bi-directional copying by using swsrpsetvol (Replication volume information setting command).

Restoration can be executed according to the following procedures:

1. Execute swsrpcancel (Replication cancellation command) when the target volume for restoration has an EC session.

2. Execute swsrpmake (Replication creation command). The specification of an original volume and the replica volume is reversed at the time that the replication is executed. For the backup operation that uses QuickOPC/SnapOPC/SnapOPC+, execute swsrpmake (Replication creation command) without specifying the -T, -C and -P option.

> 📝 **Note**
> ....
> It is not possible to perform restoration using replication if the ETERNUS Disk storage system does not support SnapOPC+ restoration. If restoration is required, copy the data from the copy destination volume manually.
>
> Figure 7.45 Manual restoration with SnapOPC+
>
>
>
> ....

7.5.4 Execution of synchronization reversal

Using swsrprevsync (Synchronization reverse command) the physical copy amount can be greatly reduced and the AdvancedCopy time shortened for restore when there is a small amount of differential data.
Additionally, it is also possible to execute restore without cancelling other copy jobs in operations that match multiple copy jobs. Therefore, there is no need for overall copy after restore, so operations can be restarted with the load on ETERNUS Disk storage system kept to a minimum.

Figure 7.46 Example of EC reversal operation

7.6 Changing the Configuration Information

When a Managed Server or device required for the replication operation has been changed, the information set in AdvancedCopy Manager must also be changed. This section explains how to change this information.

7.6.1 Changing a source volume or replica volume

To change the attributes (e.g., copy direction, operation servers of server-to-server replication) of the source volume and the replica volume, delete the information with swsrpdelvol (Replication volume information deletion command) and then execute swsrpsetvol (Replication volume information setting command) again.

7.6.2 Deleting a source volume or replica volume

To delete the source volume or the replica volume that have been configured, use swsrpdelvol (Replication volume information deletion command).

7.6.3 Changing the device information of a source or replica volume

When changing the device information and configuration for a source/replica volume, the device configuration is changed after deletion processing of the source/replica volume, and the configuration of the volume is then performed.

Be sure to perform the following steps before changing the size of a source/replica volume or a configuration. If this is not done, the deletion of the source/replica volume after device configuration change may not be possible.

1. Perform deletion process for the device change on the source/replica volume. For details of this, refer to "7.6.2 Deleting a source volume or replica volume".

2. Change the device configuration.

3. Perform any additional steps may be required in the processing of a device. For more information on this, refer to "7.4.4 Fetching device information on a Managed Server".

4. Configure the source/replica volume with the new information.

7.7 Stopping the Operation

To stop replication processing that is in progress, or to change synchronized replication into snapshot replication, use swsrpcancel (Replication cancellation command).

📵 Note

. . . .

To stop replication processing by SnapOPC+, start with the oldest snap generation.
To forcibly stop the second or subsequent snap generations, use the -Xforce option with swsrpcancel (Replication cancellation command). This will stop the specified snap generation and all earlier snap generations.

. . . .

If the -Xforce and -Xgroup options are specified at the same time, the SnapOPC+ sessions in the group are stopped, and all the snap generations and snap generations before it in the group are also stopped.

To stop replication operation, stop the daemon on the Managed Server. In general, the daemon stops when the system stops.
The daemon can be stopped independently if it needs to be stopped for some reason. For more information, refer to "Chapter 2 Starting and Stopping Daemons".

Note
....
When the daemon/service on the Managed Server has stopped, all functions of AdvancedCopy Manager running on the Managed Server will stop.
Before stopping the Management Server daemon/service, make sure that operations on all Managed Servers under management have been stopped.
....

7.8 SDX Object Replication

This section describes the replication operation of SDX objects.
A good understanding of the basic operation of normal volumes is required for this operation.

For details on GDS, refer to the "PRIMECLUSTER(TM) Global Disk Services User's Guide (Linux)".

7.8.1 Operation with Units of the SDX object operation

There are two modes of operation according to the difference in copy units:

1. Replication operation in units of logical volumes (method with GDS Snapshot)

2. Replication operation in units of slices (method without GDS Snapshot)

For details on how to confirm and select the type of operation, refer to the section titled "3.8.1.1 Configuring the unit of SDX object operations". Replication in units of logical volumes (method with GDS Snapshot) is applicable only to replication within a server.

For replication between servers, you must use units of slices.

Note
....
The two types of operation cannot be performed on a single Managed Server.
....

7.8.2 Replication operation in units of logical volumes (method with GDS Snapshot)

SDX objects of GDS can be replicated in units of logical volumes.
Replication is performed in linkage with GDS Snapshot.
You can specify a GDS logical volume name as the target volume. Therefore, you need not consider the physical disk structure of the SDX object. For operation in units of slices, all physical slices that make up the mirror volume must be specified, and the physical structure must be considered in the configuration and operation of replication.

Figure 7.47 Backup operation in units of logical volumes

To specify a GDS logical volume in an AdvancedCopy Manager command, specify the name of the GDS logical volume.

```
/dev/sfdsk/<class-name>/dsk/<volume-name>
```

🛑 Note

Configuration and conditions of unsupported SDX objects

The supported volume types are single volumes and mirror volumes.
AdvancedCopy Manager does not support the following SDX objects (volumes):

- Volume of a shadow class

- Stripe volume

- Switch volume

- Volume in a concatenation group

Stripe groups and concatenation groups, however, can be used as subordinate groups of a mirror group.
If the total number of the slices that make up the original volume and slices that make up all replica volumes associated with the original volume at the start of synchronous replication or snapshot-type replication is 33 or greater, copying cannot be processed.
Also, if the original volume is in the invalid state, copying cannot be processed.

📘 Point

To confirm whether the type of volume being used is either single, mirror, stripe, concatenation, or switch, execute the command "sdxinfo -e long" and check the volume type attribute displayed in the TYPE column for the matching volume label in the OBJ column.

[Volume object type attributes]

- single : Single volume (supported)

- mirror : Mirror volume (supported)

- stripe : Stripe volume (not supported)

- concat : Volume in a concatenation group (not supported)

- switch : Switch volume (not supported)

....

7.8.2.1 Configuring the replication operation

Note the following about performing the replication operation.
For replication with the function in linkage with GDS Snapshot, AdvancedCopy Manager processes binding and separation of volumes by using the GDS Snapshot function.
Therefore, the original and replica volumes must be SDX objects that can be used as a master object and proxy object respectively for a GDS Snapshot.
The copy source volume is the master object, and the copy destination volume is the proxy object.
Specify the GDS volume structure before registering a volume with AdvancedCopy Manager.
Note the compatibility requirements described below for replication in units of physical slices.

For detailed conditions on what constitutes a master/proxy volume, refer to "Preconditions for Proxy Constitution", "Cautions," in the "PRIMECLUSTER(TM) Global Disk Services User's Guide (Linux)".

- The original volume and replica volume must be the same type of SDX objects. Replication volume information based on logical volumes in a combination that includes slices or volumes of another OS cannot be set. When logical volumes are used, the original volume and replica volume must be the same size.

If the volumes are separated or there is no copy session, AdvancedCopy Manager assumes a non-replicated state. The following conditions will result in this state:

- OPC is started by swsrpmake (Replication creation command), and copy processing is completed.

- swsrpstartsync (Synchronous processing start command) is executed with the "software-copy" operand specified, and synchronous processing is suspended by swsrpmake (Replication creation command).

If a pair of separated replication volumes is specified, swsrpstartsync (Synchronous processing start command)swsrpmake (Replication creation command), or swsrpmake (Replication creation command) or swsrpdelvol (Replication volume information deletion command) will automatically clear the separated state.

To clear the separated state without a copy session for operational reasons, delete the replication volumes by using swsrpdelvol (Replication volume information deletion command) or by using the "sdxproxy" command, referring to "PRIMECLUSTER(TM) Global Disk Services User's Guide (Linux)".

7.8.2.2 Replication volumes

To use the advanced copy function, the structures of the original volume and replica volume must satisfy the following conditions.

Structures of copy source volumes

To use the advanced copy function, the structure of the copy source volume must be one of the following:

- Single volume

- Not the root class

- Mirror volume belonging to a mirror group that has no subordinate groups

For a mirror volume in a mirror group connected to a subordinate group, only software-copy can be used.

AdvancedCopy Manager does not support the following SDX objects (volumes):

- Volume of a shadow class

- Stripe volume

- Switch volume

- Volume in a concatenation group
  Stripe groups and concatenation groups, however, can be used as subordinate groups of a mirror group.

If the total number of the slices that make up the original volume plus the slices that make up all replication volumes associated with the original volume at the start of synchronous processing or snapshot-type replication processing is 33 or greater, copying cannot be processed.

Also, if the copy source volume is in the invalid state, copying cannot be processed.

Figure 7.48 Copy source volume structures supported by the Advanced Copy function
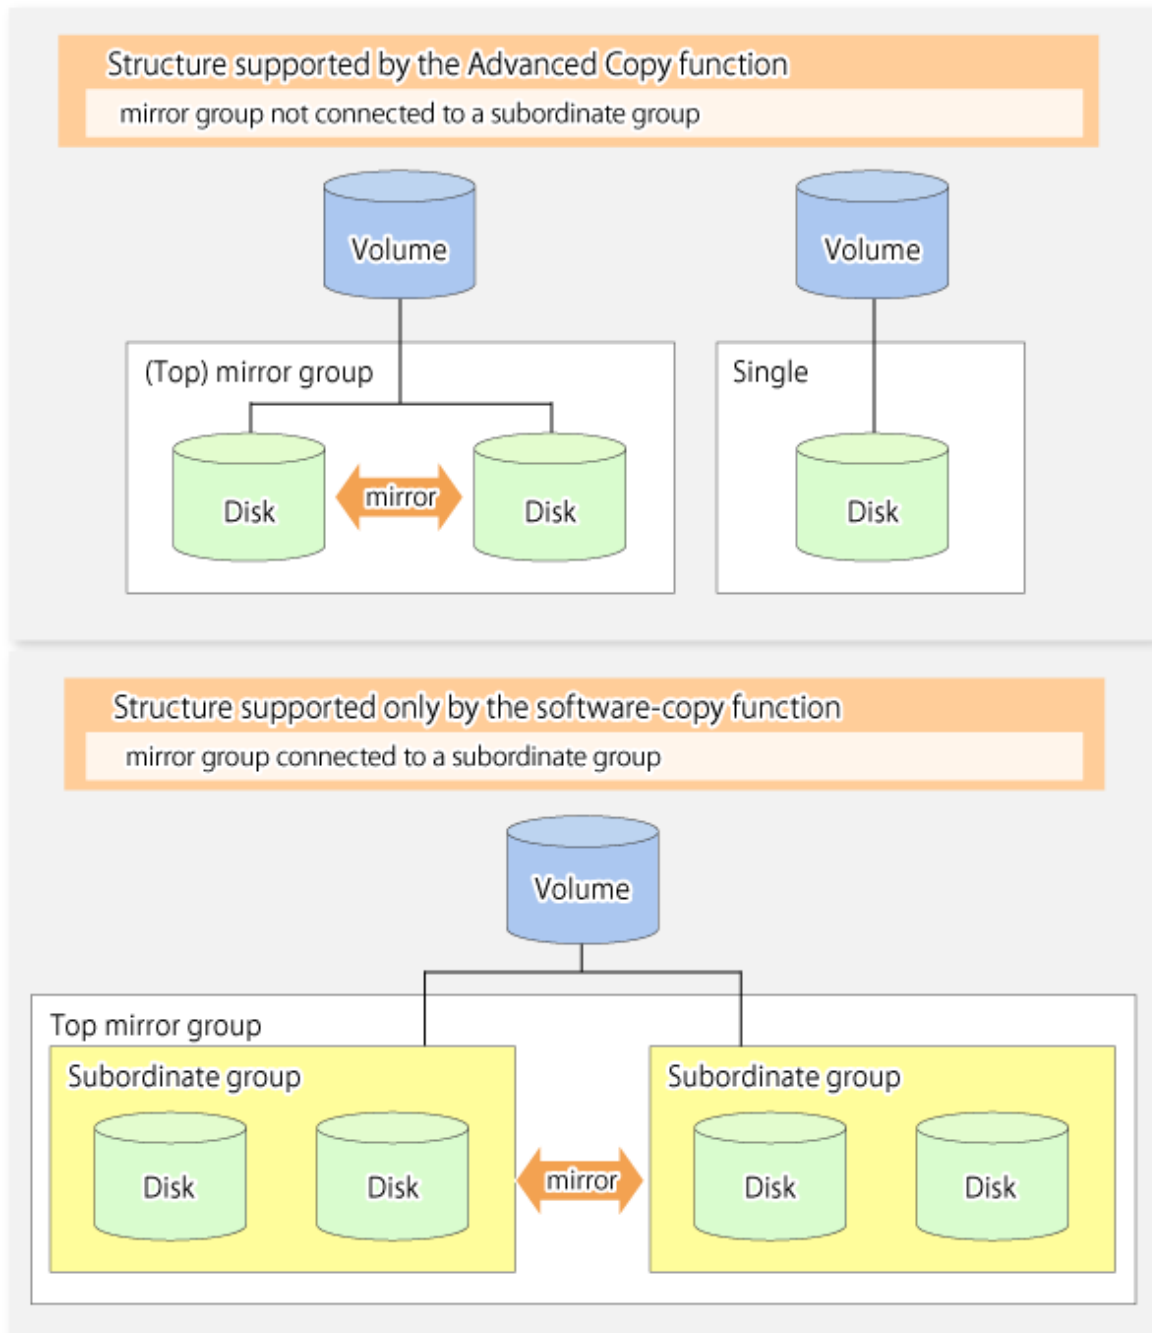
Structures of copy destination volumes

To use the advanced copy function, the structure of the copy destination volume must be one of the following:

- Single structure

  This means that the type attribute of the volume is "single" or the type attribute is "mirror" and one disk makes up the disk group.

- Not the root class

- Not a Stripe or Concatenation type

If the copy destination volume is a mirror volume, only the software-copy function can be used. To use the advanced copy function in both directions, both volumes must be in a single structure. Neither a stripe nor concatenation volume can be used as the copy destination volume.

Figure 7.49 Copy destination volume structures supported by the Advanced Copy function
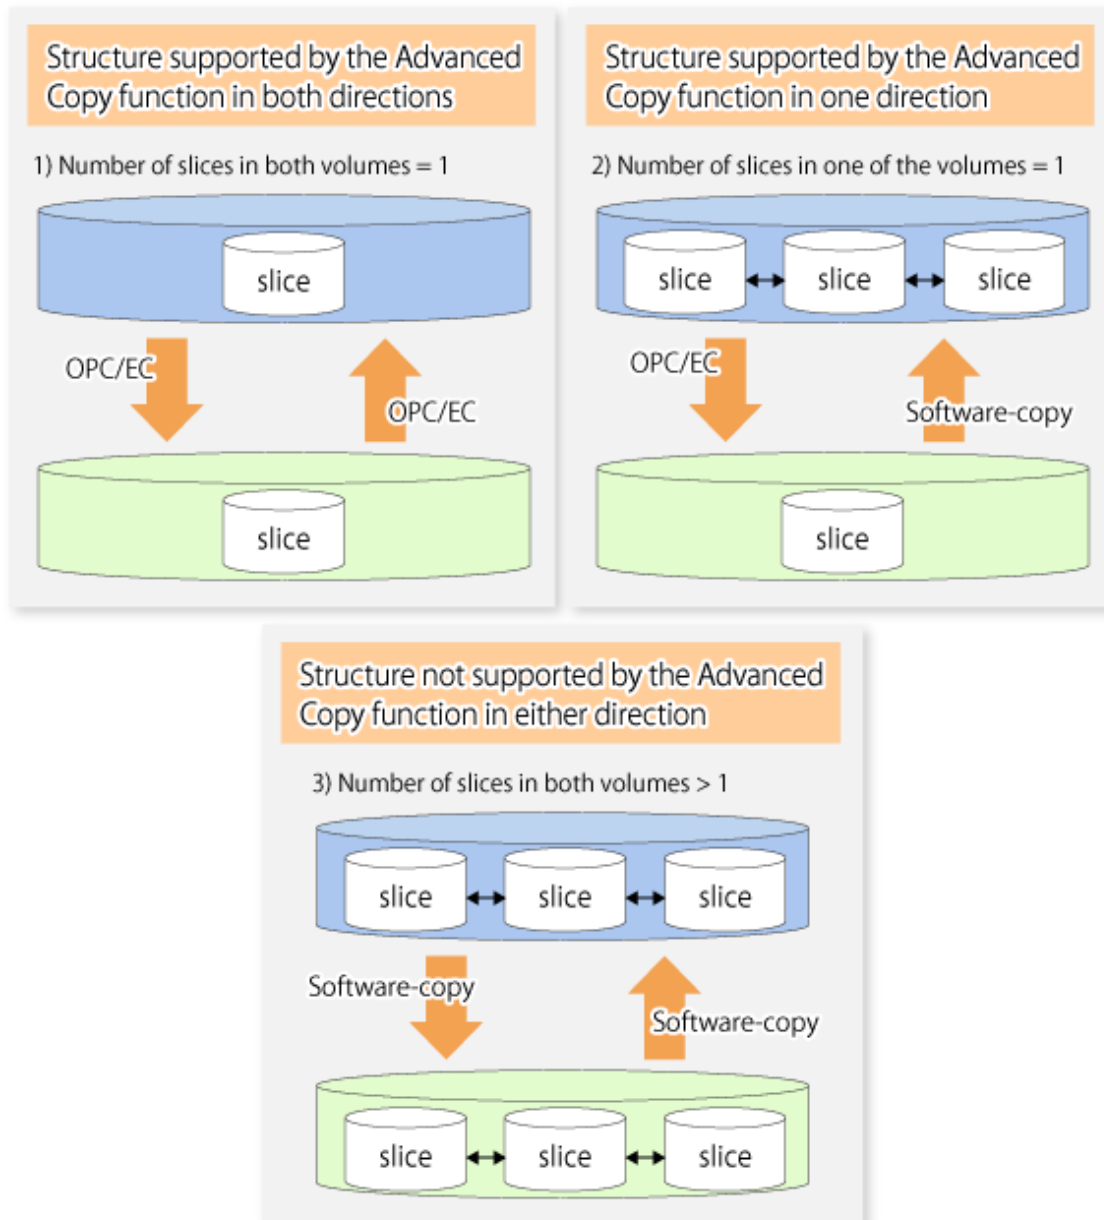
7.8.2.3 Copy functions available for each GDS structure

The two types of GDS Snapshot copy functions are described below.
The availability of the functions depends on the SDX object structure.

- Advanced Copy : Copying with the hardware function of ETERNUS Disk storage system

- Software-copy : Copying with the disk driver of PRIMECLUSTER GDS

The following table shows the relationship between SDX object structures and their available copy functions.

Table 7.6 SDX object structures and available copy functions

Replication structure	Unit	Copy source/ destination	SDX object structure(*)		Available copy function
Within a server	Logical volume	Copy source	Single		Advanced Copy/ Software-copy
			Mirror	No subordinate group exists.	Advanced Copy/ Software-copy
				Subordinate groups exist.	Software-copy
		Copy destination	Single		Advanced Copy/ Software-copy
			Mirror	This consists of one disk.	Advanced Copy/ Software-copy
				This consists of multiple disks.	Software-copy

(*) Stripe/concatenation/switch-types of objects cannot be copied.

7.8.2.4 Fetching device information from a Managed Server

To perform the replication operation with a Managed Server and obtain information relating to a subordinate device of the Managed Server, you must register the server.
For details on how to fetch the device information, refer to "7.4.4 Fetching device information on a Managed Server".

7.8.2.5 Configuring replication volumes

Note the following about configuring the original and replica volumes.

- Confirm that the volume is of a structure as described in "7.8.3.1.1 Supported logical volume structures".

- Confirm that the original/replica volume is not a master volume or proxy volume.

Register GDS logical volumes as the original and replica volumes.

```
# /opt/FJSVswsrp/bin/swsrpsetvol -h SRC-SV /dev/sfdsk/CLS1/dsk/VOL1 /dev/sfdsk/CLS1/dsk/
VOL2
swsrpsetvol completed
#
```

7.8.2.6 Executing replication

Note the following about performing replication of SDX objects.
To start or resume synchronous processing (swsrpstartsync (Synchronous processing start command)) or to start snapshot processing (swsrpmake (Replication creation command)), the logical volume of the copy destination must be inactive.
If the logical volume of the copy destination is active when the command is executed, processing to stop the volume will be performed

within the command. If the copy destination volume is a file system volume, the file system is unmounted during replication preprocessing and the volume is then stopped.

If the volume cannot be stopped (because the volume is accessed by an application), the command will be terminated abnormally.

Example of synchronous replication

```
# /opt/FJSVswsrp/bin/swsrpstartsync -h SRC-SV /dev/sfdsk/CLS1/dsk/VOL1 /dev/sfdsk/CLS1/dsk/VOL2
FROM=/dev/sfdsk/CLS1/dsk/VOL1, TO=/dev/sfdsk/CLS1/dsk/VOL2 swsrpstartsync completed
Equivalency holding state
# swsrpmake -h SRC-SV /dev/sfdsk/CLS1/dsk/VOL1 /dev/sfdsk/CLS1/dsk/VOL2
FROM=/dev/sfdsk/CLS1/dsk/VOL1, TO=/dev/sfdsk/CLS1/dsk/VOL2 swsrpmake completed
```

Example of snapshot replication

```
# /opt/FJSVswsrp/bin/swsrpmake -h SRC-SV /dev/sfdsk/CLS1/dsk/VOL1 /dev/sfdsk/CLS1/dsk/VOL2
FROM=/dev/sfdsk/CLS1/dsk/VOL1, TO=/dev/sfdsk/CLS1/dsk/VOL2 swsrpmake completed
#
```

7.8.2.7 Operational notes

Do not use the "sdxproxy" command directly to change the state of a volume registered for the replication function.
Otherwise, the "sdxproxy" command may create an inconsistency between the management information of the replication function and the actual volume state. In this case, use swsrpcancel (Replication cancellation command) to cancel the replication operation.

7.8.3 Replication operation in units of slices (method without GDS Snapshot)

SDX objects of GDS can be replicated in units of slices where replication is performed without GDS Snapshot.
This operation however requires that the physical disks that make up the GDS logical volume be specified.

For details on GDS, refer to the "PRIMECLUSTER(TM) Global Disk Services User's Guide (Linux)".

Figure 7.50 Replication operation in units of slices

> 🔰 **Note**
> ....
>
> - If mirroring between boxes is enabled and copying to the original volume by OPC is required in the event of a box failure, both mirrored systems must be replicated. In such cases, the required capacity of the replica volume is not the size of the logical volume but the size of the physical volume.
>
> - The SDX objects that can be used for operations in slice units are as follows:
>
>   - Single volumes
>
>   - Non-layered mirror volumes
>
> - You can determine if the volume being used is the single, mirror, stripe, concatenation, or switch volume type by executing "sdxinfo -e long". (Check the value in the TYPE column in the line displayed for the volume in the OBJ column.)
>
>   [Volume object type attribute]
>
>   - single: Single volume (operation possible)
>
>   - mirror: Mirror volume (operation possible)
>
>   - stripe: Stripe volume (operation not possible)
>
>   - concat: Volume in a concatenation group (operation not possible)
>
>   - switch: Switch volume (operation not possible)
>
> ....

When specifying a name in the AdvancedCopy Manager command, use the following name format that combines a logical volume name and an AdvancedCopy Manager device name. Refer to "PRIMECLUSTER GDS SDX objects" for details.

```
/dev/sfdsk/<class-name>/dsk/<volume-name>:<device name where volume is constructed>
```

If, as a countermeasure for a missing device name, the udev name is to be used instead of the sdxinfo DEVNAM value, and if the operation is an inter-server operation, both servers must run AdvancedCopy Manager Version 13.3 or later.

7.8.3.1 Configuring the replication operation

For notes on configuring the replication operation, refer to "10.1.5.2 Notes on SDX object operations (replication management)".

7.8.3.1.1 Supported logical volume structures

Replication with the copy destination volume of a mirror slice is not supported.
Consequently, if either the original volume or replica volume is a mirror slice, you can perform replication only from the mirror slice to a physical slice.
To copy data to a mirror slice, use the "dd" command. Copying must be executed for all mirror slices that make up the logical volume.

7.8.3.2 Fetching device information from a Managed Server

To perform the replication operation with a Managed Server and obtain information about a subordinate device of the Managed Server, you must register the server.
For details on how to fetch the device information, refer to "7.4.4 Fetching device information on a Managed Server".

7.8.3.3 Configuring replication volumes

Register the mirror slices that make up a mirror volume of GDS as replication volumes.

```
# /opt/FJSVswsrp/bin/swsrpsetvol /dev/sfdsk/CLS1/dsk/VOL1:sda /dev/sfdsk/CLS1/dsk/
VOL2:sdb
swsrpsetvol completed
#
```

7.8.3.4 Executing replication

Example of snapshot replication

```
# /opt/FJSVswsrp/bin/swsrpmake /dev/sfdsk/CLS1/dsk/VOL1:sda /dev/sfdsk/CLS1/dsk/VOL2:sdb
FROM=/dev/sfdsk/CLS1/dsk/VOL1:sda, TO=/dev/sfdsk/CLS1/dsk/VOL2:sdb swsrpmake completed
#
```

States of the objects in replication creation

Replication can be performed when the SDX object that makes up the relevant logical volume is in one of the states listed below. If the object is in any other state, replication cannot be performed. (AdvancedCopy Manager checks the state of the SDX object by using the sdxinfo command of GDS.)

1. The state of the mirror volume is ACTIVE or STOP.

2. The state of the physical disk is ENABLE.

3. The state of the mirror slice is ACTIVE or TEMP.

Replication pre-processing and post-processing

Pre-processing and post-processing scripts are not executed when a slice of the SDX object is in the TEMP state.

7.9 VxVM Volume Operation

This section describes the replication operation for volumes under VERITAS Volume Manager (VxVM).

7.9.1 Replication operation in units of physical disks

If a VxVM volume is the replication target, execute replication in the physical disk units that comprise the VxVM volume.

When replication operation is performed in units of physical disks, disk group integrity must be maintained and, therefore, all the physical disks in the disk group must be replicated synchronously.

Point

For the physical volume which should be synchronized, confirm it by using either Web Console or stgxfwcmdispdev (Device information display command).

Note

- Before starting this operation, be sure to understand the basic operation of ordinary volumes.

- Snapshot backup is recommended for backing up in units of physical disks. In the case of synchronized backup, commands that access destination volumes, such as VxVM commands, cannot be executed during full copy or differential copy.

- For operation in a cluster configuration, the device name (/dev/sd#) of the physical disks that comprise the disk group must be the same at all servers that comprise the cluster, and the ETERNUS Disk storage system's disk indicated by the device name must be the same.

7.9.1.1 Configuration

When configuring the disk groups of the original and replica volumes, you should note the following:

- The number, sizes, and types of VM disks must be the same.

- The configurations of logical volumes must be the same.

Figure 7.51 Copy source volume and copy destination volume

7.9.1.2 Preparations

7.9.1.2.1 Confirming the disk group configuration information file

The disk group must be reconfigured in replication post-processing. Confirm that a volume group configuration information file has been saved.

```
/etc/vx/cbr/bk/<disk group name>.<disk group ID>
```

7.9.1.2.2 Configuring the original and replica volumes

When setting the original and replica volumes, specify all disks in the volume group.

Example of setting original and replica volumes

```
# /opt/FJSVswsrp/bin/swsrpsetvol /dev/vx/dmp/sdg /dev/vx/dmp/sdx
swsrpsetvol completed
# /opt/FJSVswsrp/bin/swsrpsetvol /dev/vx/dmp/sdh /dev/vx/dmp/sdy
swsrpsetvol completed
#
```

7.9.1.3 Replication

Perform Replication by synchronizing all physical disks in the disk group.

Perform the required pre-processing or post-processing work for each volume group respectively before and after the replication operation. Disable pre-processing and post-processing when replicating individual physical disks.

Example of snapshot backup

```
(Perform pre-processing for the source and target volumes.)

# /opt/FJSVswsrp/bin/swsrpmake -f -t /dev/vx/dmp/sdg /dev/vx/dmp/sdx
FROM=/dev/vx/dmp/sdg@SV1, TO=/dev/vx/dmp/sdx@SV1 swsrpmake completed
# /opt/FJSVswsrp/bin/swsrpmake -f -t /dev/vx/dmp/sdh /dev/vx/dmp/sdy
FROM=/dev/vx/dmp/sdh@SV1, TO=/dev/vx/dmp/sdy@SV1 swsrpmake completed
#

(Perform post-processing for the source and target volumes.)
```

Example of synchronous replication

```
(Perform pre-processing for the target volume.)

# /opt/FJSVswsrp/bin/swsrpstartsync -t /dev/vx/dmp/sdg /dev/vx/dmp/sdx
FROM=/dev/vx/dmp/sdg@SV1, TO=/dev/vx/dmp/sdx@SV1 swsrpstartsync completed
# /opt/FJSVswsrp/bin/swsrpstartsync -t /dev/vx/dmp/sdh /dev/vx/dmp/sd
FROM=/dev/vx/dmp/sdh@SV1, TO=/dev/vx/dmp/sdy@SV1 swsrpstartsync completed
#

(After state of equivalency upkeep)
(Perform pre-processing for the source volume.)

# /opt/FJSVswsrp/bin/swsrpmake -f -t /dev/vx/dmp/sdg /dev/vx/dmp/sdx
FROM=/dev/vx/dmp/sdg@SV1, TO=/dev/vx/dmp/sdx@SV1 swsrpmake completed
# /opt/FJSVswsrp/bin/swsrpmake -f -t /dev/vx/dmp/sdh /dev/vx/dmp/sdy
FROM=/dev/vx/dmp/sdh@SV1, TO=/dev/vx/dmp/sdy@SV1 swsrpmake completed
#

(Perform post-processing for the source and target volumes.)
```

The table below summarizes the pre-processing and post-processing work to be performed before and after replication.

Table 7.7 Replication pre-processing and post-processing

Type	Pre-processing	Post-processing
Source volume	1. Secure data integrity by stopping access to all logical volumes in the volume group. 2. If file systems are included, unmount all file systems in the volume group. 3. Import the disk group, when the disk group is not imported for replication within the server.	If file systems are included, remount the volumes that were unmounted by pre-processing.
Target volume	1. Stop access to all logical volumes in the volume group. 2. If file systems are included, unmount all file systems in the volume group. 3. Deport the disk group. 4. A disk group subordinate's physical disk is set to offline.	1. The physical disk set to offline with pre-processing is set to online. 2. Reconfigure the volume group. 3. If file systems are included, remount the volumes that were unmounted by pre-processing.

Reconfiguring the disk group (for replication within the server)

Reconfigure the disk group as follows:

1. Pre-commit analysis for restoration

```
# /etc/vx/bin/vxconfigrestore -p dstdg
Diskgroup dstdg configuration restoration started ......

Installing volume manager disk header for sdx ...
Installing volume manager disk header for sdy ...
-
dstdg's diskgroup configuration is restored (in precommit state).
Diskgroup can be accessed in read only and can be examined using
vxprint in this state.

Run:
  vxconfigrestore -c dstdg ==> to commit the restoration.
  vxconfigrestore -d dstdg ==> to abort the restoration.
 #
```

2. Commit the change required for restoring the configuration of the copy destination disk group.

```
# /etc/vx/bin/vxconfigrestore -c dstdg
Committing configuration restoration for diskgroup dstdg ....

dstdg's diskgroup configuration restoration is committed.
#
```

Reconfiguring the disk group (for replication between servers)

1. Import the disk group "srcdg" on the master server as disk group "dstdg" on the target server.

```
# /usr/sbin/vxdg -C -n dstdg import srcdg
#
```

🛈 Note
. . . .

When the disk group name is the same in the server of the source and destination, the -n option is not specified.
. . . .

2. Execute recovery processing for the volume in the disk group "dstdg" on the target server.

```
# vxrecover -g dstdg -sb
#
```

3. Because udid of the copy destination disk is not an adjustment, this is restored.

```
# vxdisk updateudid sdx sdy
#
```

🛈 Note
. . . .

- In case of a clustered system, when a disk group resource has been defined, use the online/offline process instead of the disk group import/deport command.
  If a mount point is defined as a clustered system resource, use the mount resource online/offline process instead of the file system mount/unmount process.

- When performing a system disk exchange, there are cases when multiple conflicting configuration information backups may exist for one disk group.
  In such cases, the disk group ID displayed after executing the above command needs to be set and executed instead of the disk group.

- When the volumes within a disk group are required to be run in synchronous mode in the background and depending on the volume configuration it may take some time for synchronous processing to complete after this operation.
  It is also possible to use the volumes during this time.

Chapter 8 Operation Maintenance

This chapter provides information on how to maintain a database, perform troubleshooting, and change the operation information.

8.1 Maintenance of Database

AdvancedCopy Manager holds various repositories of management information on the usage of its functions. They are named generically and are referred to as a database in this chapter.

A database is important information. You are recommended to periodically back up a database that is frequently updated, when the space available for repositories becomes insufficient after long-time usage of the database that is different from originally planned, as a safeguard against running out of space. (Note that it differs from other types of databases such as ORACLE etc.)

The following table lists instances when backup of the database is needed.

Table 8.1 Databases on a Management Server

Type	When a backup to be created	Backup method (references)	Size
Database	A Managed Server has been added. The configuration information has changed as a result of loading of new device information or for some other reason. The information on backup management or the replication management operation has changed.	"8.1.1.3.1 Saving a database" or "8.1.2 Making a package backup of databases"	Refer to "8.1.1.3.1 Saving a database".
Registry	DB setup has been performed on a Management Server (a Management Server transaction has been created in a clustered system), or the storage management information (server name, IP address or port number) has changed.	"8.1.1.4.1 Backing up a registry" or "8.1.2 Making a package backup of databases"	Refer to "8.1.1.4.1 Backing up a registry".

Table 8.2 Databases on a Managed Server

Type	When backup to be made	Backup method (reference)	Size
Registry	A Managed Server has been added or Management Server information (server name, IP address or port number) has changed.	"8.1.1.4.1 Backing up a registry" or "8.1.2 Making a package backup of databases"	Refer to "8.1.1.4.1 Backing up a registry".
Backup management list	A backup transaction has been completed.	"8.1.1.1.1 Backing up a backup management list and recovery control file" or "8.1.2 Making a package backup of databases"	Refer to "8.1.1.1.1 Backing up a backup management list and recovery control file".
Pre-processing and post-processing scripts for backup management	A pre-processing or post-processing script has changed.	Saving a database by using the copy command or "8.1.2 Making a package backup of databases"	Check the size of the pre-processing and post-processing script file for backup management.

Type	When backup to be made	Backup method (reference)	Size
Replication management list	A replication transaction has been completed.	"8.1.1.2.1 Backing up a replication management list" or "8.1.2 Making a package backup of databases"	Refer to "8.1.1.2.1 Backing up a replication management list".
Pre-processing and post-processing scripts for replication management	A pre-processing or post-processing script has changed.	Saving a database by using the copy command or "8.1.2 Making a package backup of databases"	Check the size of the pre-processing and post-processing script file for replication management.
Policy file for management information package backup	The policy file has changed.	Saving a database by using the copy command or "8.1.2 Making a package backup of databases"	Check the size of the policy file for management information package backup.

ⓟ Point

. . . .

You can create backups of databases either individually or as a whole.

If a Managed Server also operates as a Management Server, the Managed Server contains a database for the Management Server as well as for the Managed Server.

. . . .

8.1.1 Making individual backups of databases

8.1.1.1 Maintaining a backup management list

This section provides information on the maintenance of a backup management list. A backup management list is maintained on the server that performs the backup.

8.1.1.1.1 Backing up a backup management list and recovery control file

A backup management list and a recovery control file created at the time of backing up the Symfoware database spaces are backed up. For details of this, refer to swstresback (Resource backup command).

Backup management list backup data size

Check the size of the file that the subdirectory of the following directory contains.

```
/etc/opt/FJSVswsts/data
```

In the cluster environment, check the size of the file that the subdirectory of the following directory contains.

```
/etc/opt/FJSVswsts/<logical node name>/data
```

For backing up Symfoware databases, check the file sizes in the following directories:

```
Target directory for output of the recovery control file
```

8.1.1.1.2 Restoring a backup management list and recovery control file

A backup management list and a recovery control file created at the time of backing up the Symfoware database spaces are restored. For details of this, refer to swstresrst (Resource restore command).

8.1.1.2 Maintaining a replication management list

This section describes the maintenance of a replication management list.

8.1.1.2.1 Backing up a replication management list

To back up a replication management list, back up the repository of the Management Server. For information about backing up a repository, refer to "Saving a database" in the "ETERNUS SF AdvancedCopy Manager Operation Guide" relevant to the OS of the Management Server.

8.1.1.2.2 Restoring a replication management list

Perform the following steps to restore a replication management list:

1. Restore the repository to the Management Server. For information about restoring a repository, refer to "Recovering a database" in the "ETERNUS SF AdvancedCopy Manager Operation Guide" relevant to the OS of the Management Server.

    - If the Management Server is Windows
      Select "RDB dictionary and DSI recovery (LOAD Mode)" as the recovery method.

    - If the Management Server is Solaris or Linux
      Specify 0 for stgdbrcv -m option (recovery mode) as the recovery method.

2. On the Managed Server, execute swsrprecoverres (Resource adjustment command) with the -r option.

8.1.1.3 Maintaining a database

This section describes the maintenance of a database used by a Management Server.

8.1.1.3.1 Saving a database

Execute stgdbdmp (Database save command) to save a database space.

Execute this command on the Management Server.

After saving a database, refer to "8.7 Maintaining a database of ETERNUS SF Manager" and save a database of ETERNUS SF Manager as well.

P Point

．．

- Use this save method only when the configuration information is changed due to the addition of a Managed Server, or fetching of new device information. It is recommended to save a database regularly (for example, once a day).

- To safeguard against an input-output error in an external file while saved data is being acquired, it is recommended to keep at least two generations of saved data.

．．

Database backup data size

Calculate the size using the following formula:

60 + (number of devices to be managed x 0.001) megabytes

8.1.1.3.2 Finding an error in a database

Refer to the following file, and beginning from the last line, search for any message with "rdb: ERROR: qdgXXXXX-" in its contents (also check the messages displayed at the same time). Then, obtain the value of "qdgXXXXX" to determine the recovery mode.

For non-cluster operation

```
/var/opt/FJSVswstf/log/RDBSWSTF.log
```

For cluster operation

```
/var/opt/FJSVswstf/<logical-node-name>/log/RDBSWSTF.log
```

Table 8.3 Relationship between qdg messages and recovery modes

qdg message	Description of error	Target of checking	Recovery mode
qdg13039u	I/O error or detected damage	*<Directory for storing data of a DB space in a repository>*/hdrdb_rep1	1
qdg12079e	File delete		
qdg03400u	Rollback disabled	-	
qdg03401u			
qdg03121u	I/O error	*<Directory for the RDB dictionary>*/hdrdic	2
qdg12078u	File delete		
qdg03420u	I/O error, detected damage, etc.	*<Directory for DB file>*/SWSTFDB/rdbdir/ DIR_FILE1	
qdg12150u	File open error		

qdg message	Description of error	Target of checking	Recovery mode
qdg12151u		*<Directory for DB file>*/SWSTFDB/rdbdir/ DIR_FILE2	
qdg12251e	I/O error	*<Directory for DB file>*/SWSTFDB/rdbdir/ rdblogmanage	3
qdg12443u			
qdg12434u	File error		
qdg02290u	File delete		
qdg12111u	I/O error	*<Directory for RDB log files>*/hdrlog	4
qdg12440u			
qdg12441u			
qdg12442u			
qdg02296u	File delete		
qdg03815u	System call error	*<Directory for DB file>*/SWSTFDB/rdbdir/archive1	5
qdg12275W		*<Directory for DB file>*/SWSTFDB/rdbdir/archive2	
qdg12250u	Damage detection		
qdg12483u	File read error		

Note

- If multiple errors occur simultaneously, recovery using recovery modes 1 to 5 may not be possible.

  To restore a database to its state when it was last saved, re-create the database and execute "8.1.1.3.1 Saving a database" in recovery mode 0.
  Data updated or inserted after the last time the database was saved cannot be recovered.

- The swstsrsemtch (Resource match command) with -x option will need to be executed on all Managed Servers on which the backup is performed, after which swsrprecoverres (Resource adjustment command) with -r option will need to be executed on the Managed Servers on which the replication is performed.

8.1.1.3.3 Recovering a database

Execute the database recovery command for database recovery if an I/O error occurs in a database space.
Execute this command on the Management Server.

For details on this command, refer to "9.3.19 stgdbrcv (Database recovery command)".

8.1.1.4 Maintaining a registry

This section describes the maintenance of a registry. Perform this procedure on a Management Server or a Managed Server.

8.1.1.4.1 Backing up a registry

In a command prompt window, back up the following files using the copy command, etc.

File names

In the case of a non-cluster

```
/etc/opt/swstorage/swnode.ini
```

In the case of a cluster

```
/etc/opt/swstorage/<Logical-node-name>/swnode.ini
```

Example

Save it in the /backup directory.

```
# cp /etc/opt/swstorage/swnode.ini /backup
```

Registry backup data size

Check the size of the file to be backed up.

8.1.1.4.2 Restoring a registry

In a command prompt window, use the copy command or another method to restore the following files.

File names

In the case of a non-cluster

```
/etc/opt/swstorage/swnode.ini
```

In the case of a cluster

```
/etc/opt/swstorage/<Logical-node-name>/swnode.ini
```

Example:

Restore the following file saved in the /backup directory:

```
# cp /backup/swnode.ini /etc/opt/swstorage
```

8.1.2 Making a package backup of databases

This section describes the function of making a backup of management information (e.g., databases) from just one Managed Server or Management Server.

With a package backup of management information, you can make a backup of management information for all Managed Servers and Management Server from just one Managed Server or Management Server. Executing a package backup of management information runs backup commands (making backups of backup management resources and databases) for the databases of other Managed Servers and Management Server, and collects the backup data and management information created by these backup commands.

The server on which you can execute a package backup of management information is called the "package-backup-of-the-management-information execution server". The servers whose data is to be backed up are called the "package-backup-of-the-management-information target servers".

Figure 8.1 Management Server and Managed Server

Example

Assume that the package backup of management information function has been used to back up the management information on a Management Server, Managed Server A, and Managed Server B onto Managed Server B in the diagram above. In this case, Managed Server B is the backup execution server. The Management Server, Managed Server A, and Managed Server B are the backup target servers. That is, Managed Server B operates as both the backup execution server and backup target server.

The following table shows the types of management information that is handled in a package backup of management information.

Table 8.4 Types of management information backed up by package backup of management information

Function	Type of data
Database	- Database
Registry	- Registry
Backup management	- Backup management list, or backup management list and recovery control file - Pre-processing and post-processing scripts - Transaction volume locking specification file for backups - Backup volume locking specification file for backups - Backup-restored volume locking specification file
Replication management	- Pre-processing and post-processing scripts - Copy-source volume locking specification file - Copy-destination volume locking specification file
Package backup of management information	- Policy file

 Note

Managed Server management information created by Softek AdvancedCopy Manager Version 10.3 (V10.0L40) or earlier cannot be handled by the package backup of management information function. To create a backup, back up each management information component individually as described in "Maintaining a database" in the "ETERNUS SF AdvancedCopy Manager Operation Guide" applicable to the Managed Server.

8.1.2.1 Preliminaries

Before creating a package backup of management information, you must perform the following:

1. Determine the management information package backup target servers and the types of management information to be backed up.

2. Determine the management information package backup execution server.

3. Create a directory where the management information package backup execution server will store the backup data. The size of the directory must not be less than the total of the following sizes:

   - The required size for a database is the size of the database backup data multiplied by 1.2.
     For the backup data size, refer to "Maintaining a database" in the "ETERNUS SF AdvancedCopy Manager Operation Guide" for the operating system running on the Management Server.

   - The required size for the registries is the total of the data sizes of the registry backup data on the management information package backup target servers.
     For the backup data size, refer to "Maintaining a database" in the "ETERNUS SF AdvancedCopy Manager Operation Guide" for the operating system running on the management information package backup target server.

   - The required size for backup management is the total of the sizes below for each management information package backup target server:
     For the size of each component of backup data, refer to "Maintaining a database" in the "ETERNUS SF AdvancedCopy Manager Operation Guide" for the operating system running on the management information package backup target server.

     - Size of the backup data for the backup management lists (size of the backup data for the backup management lists plus the size of the backup data for the recovery control files) multiplied by 1.2

     - Size of the pre-processing and post-processing scripts for backup management

     - Size of the transaction volume locking specification file for backups for backup management

     - Size of the backup volume locking specification file for backups for backup management

     - Size of the backup-restored volume locking specification file for backup management

   - The required size for replication management is the total of the following sizes for each management information package backup target server:

     - Size of the pre-processing and post-processing scripts for replication management

     - Size of the copy-source volume locking specification file for replication management

     - Size of the copy-destination volume locking specification file for replication management

4. Check the size of the working directory for the management information package backup target servers (or the shared disk for shared data for AdvancedCopy Manager in a clustered system). The sizes described below are needed only when a package backup of management information is being created. Use the largest value of the sizes of the databases to be backed up from the management information package backup target servers. The required sizes are as follows:

   - The required size for the databases is the size of the database backup data multiplied by 2.2.
     For the backup data size, refer to "Maintaining a database" in the "ETERNUS SF AdvancedCopy Manager Operation Guide" for the operating system running on the Management Server.

   - The required size for backup management is the size of the backup data for the backup management lists (size of the backup data for the backup management lists plus the size of the backup data for the recovery control files) multiplied by 2.2.
     For the backup data size, refer to "Maintaining a database" in the "ETERNUS SF AdvancedCopy Manager Operation Guide" for the operating system running on the management information package backup target server.

   - The working directory does not require any space for registry, replication management, and policy file backups.

5. Execute the stgmgrinfoset command as described in stgmgrinfoset (Management information package backup server registration/deletion command), on the management information package backup execution server (determined in step 2).

6. Create a policy file on the management information package backup execution server.

Table 8.5 Management information package backup execution server and policy file paths

Management information package backup execution server		Policy file path
OS	Operation type	
Windows	Non-cluster	*<Environment directory>*\etc\stgmgrinfo_policy.ini
	Cluster	*<AdvancedCopy-Manager-shared-data>*:\etc\opt\swstorage\etc\stgmgrinfo_policy.ini
Solaris, Linux, HP-UX, AIX	Non-cluster	/etc/opt/FJSVswstf/stgmgrinfo_policy.ini
	Cluster	/etc/opt/FJSVswstf/*<logical-node-name>*/stgmgrinfo_policy.ini

📷 Note

. . . .

- Do not specify more than one Managed Server as a management information package backup execution server.

- When this server is used for executing transactions on the Managed Server in a clustered system, create a directory for storing backup data on a shared disk.

- For databases and backup management, backup commands are executed on the corresponding management information package backup target servers. In addition, for database backups and backup management (ie, backup management lists) backups, a temporary capacity that is approximately 2.2 times greater than the total size of the backups is required to save the backup data in an archive file.

. . . .

Table 8.6 Policy file format

Section name	Key name	Description (**Note**)	Whether specifiable on	
			Management Server	Managed Server
Server name	DBData	Specifies whether to make a backup of the database. Yes: Make a backup of the database. No: Do not make a backup of the database.	A	C
	BkupData	Specifies whether to make a backup of the backup management. Yes: Make a backup of the backup management. No: Do not make a backup of the backup management.	B	B
	ReplData	Specifies whether to make a backup of the replication management. Yes: Make a backup of the replication management. No: Do not make a backup of the replication management.	B	B

A: Specifiable
B: Specifiable only when the applicable function is operating on the server
C: Unspecifiable

Note: If this key is not specified or if "Yes" is not in the form "Y" (uppercase) + "es" (lowercase), then "No" is set.

📘 Example

. . . .

Set up a policy file as follows:

- The Management Server named ACMMGR creates backups for the database and backup management.

- The Managed Server named ACMAGT1 creates backups for backup management.

- The Managed Server named ACMAGT2 creates backups for replication management.

```
[ACMMGR]
DBData=Yes
BkupData=Yes
[ACMAGT1]
BkupData=Yes
[ACMAGT2]
ReplData=Yes
```

![Note icon] **Note**

- Create a policy file only if it will be used.

- The policy file cannot include any settings for backups of registries and of the policy file for management information package backup.

8.1.2.2 Backup operations

Executing stgmgrinfobkup (Management information package backup command) on the management information package backup execution server stores the backup data in the directory under the following structure:

Table 8.7 Directory structure for execution of management information package backup and data types

Directory			Data type
*1	Server name	DBData	- Database
		RegiData	- Registry
		BkupData	FJSVswsts	- Backup management list, or backup management list and recovery control file
			Scripts	- Pre-processing and post-processing scripts for backup management
			data	- Transaction volume locking specification file for backups, for backup management - Backup volume locking specification file for backups, for backup management - Backup-restored volume locking specification file for backups, for backup management
		ReplData	Scripts	- Pre-processing and post-processing scripts for replication management
			data	- Copy-source volume locking specification file for replication management - Copy-destination volume locking specification file for replication management
		MgrInfoData	- Policy file for management information package backup

*1 is the backup data storage directory specified by stgmgrinfoset (Management information package backup server registration/deletion command).

Example

When *1 is /ACMBkupData and the server name is ACMAGT1, the pre-processing and post-processing scripts for backup management are put in the following directory:

```
/ACMBkupData/ACMAGT1/BkupData/Scripts
```

Initial operations

The policy file backup operations are described below. Backup of registries which are not included in the backup targets when a policy file is used to perform backup are also detailed. The initial operations are as follows:

1. Execute stgmgrinfobkup (Management information package backup command) on the management information package backup execution server to make backups of the registries on all of the Managed Server.

2. Execute stgmgrinfobkup (Management information package backup command) on the management information package backup execution server to make a backup of the policy file.

Normal-use operations

Execute stgmgrinfobkup (Management information package backup command) on the management information package backup execution server when you need to make backups of the target data.
For the timing of the backups, refer to "Maintaining a database" in the "ETERNUS SF AdvancedCopy Manager Operation Guide" for the applicable operating system.

Operations when a server is added or deleted

Operations when server added

1. On the management information package backup execution server, set up a policy for the added Managed Server in the policy file.

2. Execute stgmgrinfobkup (Management information package backup command) on the management information package backup execution server to make a backup of the registry of the added Managed Server.

3. Execute stgmgrinfobkup (Management information package backup command) on the management information package backup execution server to make a backup of the policy file.

Operations when server deleted

1. On the management information package backup execution server, delete the information about the Managed Server that was deleted from the policy file.

2. In the backup data directory on the management information package backup execution server, delete the directory containing the name of the deleted Managed Server.

3. Execute stgmgrinfobkup (Management information package backup command) on the management information package backup execution server to make a backup of the policy file.

Operations when a server name is changed

1. On the management information package backup execution server, delete the information for the Managed Server from the policy file.

2. In the backup data directory on the management information package backup execution server, delete the directory containing the name of the Managed Server.

3. On the management information package backup execution server, set up a policy for the renamed Managed Server in the policy file.

4. Execute stgmgrinfobkup (Management information package backup command) on the management information package backup execution server to make a backup of the registry of the renamed Managed Server.

5. Execute stgmgrinfobkup (Management information package backup command) on the management information package backup execution server to make a backup of the policy file.

8.1.2.3 Restoration operations

This section describes how to use the backup data created by stgmgrinfobkup (Management information package backup command) to execute restoration.
Perform the operations below as a root user or with an account that belongs to the Administrators group.

Restoring a registry

1. Execute stgmgrinfoset (Management information package backup server registration/deletion command) on the management information package backup execution server to determine which directory contains the backup data.

2. On the management information package backup execution server, copy the registry backup to the following directory:

   Windows version:

   ```
   <Directory-determined-in-step1>\<server-whose-registry-to-be-restored>\RegiData
   ```

   Solaris, Linux, HP-UX or AIX version:

   ```
   <Directory-determined-in-step1>/<server-whose-registry-to-be-restored>/RegiData
   ```

   1. Transfer the registry backup that was copied in step 2 to the target server in binary mode.

   2. Restore the registry on the server. For details on restoring a registry, refer to "Restoring a registry" in the "ETERNUS SF AdvancedCopy Manager Operation Guide" for the operating system running on the target server.

Restoring a database

1. Execute stgmgrinfoset (Management information package backup server registration/deletion command) on the management information package backup execution server to determine which directory contains the backup data.

2. On the management information package backup execution server, copy the backup of the database to the following directory:

   Windows version:

   ```
   <Directory-determined-in-step1>\<management-server-name>\RegiData
   ```

   Solaris, Linux, HP-UX or AIX version:

   ```
   <Directory-determined-in-step1>/<management-server-name>/RegiData
   ```

3. Transfer the backup of the database that was copied in step 2 to the Management Server in binary mode.

4. On the Management Server, expand the database backup transferred in step 3.

   - On the Windows version Management Server:

     Expand the database backup using the stgarcext (Backup data expansion command) in the "ETERNUS SF AdvancedCopy Manager Operation Guide (for Windows)".

   - On a Solaris/Linux version Management Server:

     Expand the database backup by using the tar command supported by the operating system.

5. On the Management Server, restore the database.

   For details on restoring a database, refer to "Restoring a database" in the "ETERNUS SF AdvancedCopy Manager Operation Guide" for the operating system running on the Management Server.

   - The backup data storage directory specified for database restoration is the directory expanded in step 4.

   - Specify "RDB Dictionary, DSI(LOAD Mode)" for the Windows version as a recovery mode. Specify "0" for other versions.

Restoring a backup management list, or backup management list and recovery control file

1. Execute stgmgrinfoset (Management information package backup server registration/deletion command) on the management information package backup execution server to determine which directory contains the backup data.

2. On this server, copy the backup of the backup management list or the backups of the backup management list and the recovery control file to the following directory:

    Windows version:

    ```
    <Directory-determined-in-step1>\<server-whose-list/file-to-be-restored>\BkupData\FJSVswsts
    ```

    Solaris, Linux, HP-UX or AIX version:

    ```
    <Directory-determined-in-step1>/<server-whose-list/file-to-be-restored>/BkupData/FJSVswsts
    ```

3. Transfer the backup of the backup management list or the backups of the backup management list and the recovery control file copied in step 2 to the target server in binary mode.

4. Expand the backup of the backup management list or the backups of the backup management list and recovery control file transferred in step 3 on the target server.

    - On a Windows version server:

      Expand the backup of the backup management list using the stgarcext (Backup data expansion command) in the "ETERNUS SF AdvancedCopy Manager Operation Guide (for Windows)".

    - On a Solaris/Linux/HP-UX/AIX version server:

      Expand the backup of the backup management list or the backups of the backup management list and recovery control file by using the tar command supported by the operating system.

5. Restore the backup management list, or the backup management list and recovery control file on the target server. For details on restoring, refer to "Restoring a backup management list" or "Restoring a backup management list and recovery control file" in the "ETERNUS SF AdvancedCopy Manager Operation Guide" for the operating system running on the target server. The Backup-Directory specified in the resource restoration command is the directory expanded in step 4.

Restoring the backup management pre-processing/post-processing scripts

1. Execute stgmgrinfoset (Management information package backup server registration/deletion command) on the management information package backup execution server to determine which directory contains the backup data.

2. On this server, copy the backups of the backup management pre-processing/post-processing scripts for backup management to the following directory:

    Windows version:

    ```
    <Directory-determined-in-step1>\<server-whose-scripts-to-be-restored>\BkupData\Scripts
    ```

    Solaris, Linux, HP-UX or AIX version:

    ```
    <Directory-determined-in-step1>/<server-whose-scripts-to-be-restored>/BkupData/Scripts
    ```

3. Transfer the backups of the backup management pre-processing/post-processing scripts copied in step 2 to the target server in binary mode.

4. Transfer the backups of the backup management pre-processing/post-processing scripts copied in step 3 to the directory in which the scripts will be executed on the target server in binary mode. For the directory in which the pre-processing and post-processing scripts for backup management are executed (that is, the directory to which the scripts are to be copied), refer to "Pre-processing and post-processing for backup/restoration" in the "ETERNUS SF AdvancedCopy Manager Operation Guide" for the operating system running on the target server.

5. On a Solaris/Linux/UP-UX/AIX-version server, change the access permission assigned to the pre-processing and post-processing scripts for backup management copied in step 4 on the target server by using the chmod command supported by the operating system. Set execution permission for the root user, which is the owner.

Restoring locking specification files for backup management

1. Execute stgmgrinfoset (Management information package backup server registration/deletion command) on the management information package backup execution server to determine which directory contains the backup data.

2. On this server, copy the backups of the locking specification files for backup management to the following directory:

   Windows version:

   ```
   <Directory-determined-in-step1>\<server-whose-files-to-be-restored>\BkupData\data
   ```

   Solaris, Linux, HP-UX or AIX version:

   ```
   <Directory-determined-in-step1>/<server-whose-files-to-be-restored>/BkupData/data
   ```

3. Transfer the backups of the locking specification files for backup management copied in step 2 to the target server in binary mode.

4. Copy the backups of the locking specification files for backup management transferred in step 3 to the target server. For the directory in which the locking specification files for backup management are used (that is, the directory to which the files are to be copied), refer to "Appendix A Pre-processing and Post-processing of Backup and Restoration" in the "ETERNUS SF AdvancedCopy Manager Operation Guide (for Windows)".

Restoring pre-processing and post-processing scripts for replication management

1. Execute stgmgrinfoset (Management information package backup server registration/deletion command) on the management information package backup execution server to determine which directory contains the backup data.

2. On this server, copy the backups of the pre-processing and post-processing scripts for replication management to the following directory:

   Windows version:

   ```
   <Directory-determined-in-step1>\<server-whose-scripts-to-be-restored>\ReplData\Scripts
   ```

   Solaris, Linux, HP-UX or AIX version:

   ```
   <Directory-determined-in-step1>/<server-whose-scripts-to-be-restored>/ReplData/Scripts
   ```

3. Transfer the backups of the pre-processing and post-processing scripts for replication management copied in step 2 to the target server in binary mode.

4. Copy the backups of the pre-processing and post-processing scripts for replication management transferred in step 3 to the directory in which the scripts will be executed on the target server. For the directory in which the pre-processing and post-processing scripts for replication management are executed (that is, the directory to which the scripts are to be copied), refer to "Pre-processing and post-processing of replication" in the "ETERNUS SF AdvancedCopy Manager Operation Guide" for the operating system running on the target server.

5. On a Solaris/Linux/UP-UX/AIX-version server, change the access permission assigned to the pre-processing and post-processing scripts copied in step 4 on the target server by using the chmod command supported by the operating system. Set execution permission for the root user, which is the owner.

Restoring locking specification files for replication management

1. Execute stgmgrinfoset (Management information package backup server registration/deletion command) on the management information package backup execution server to determine which directory contains the backup data.

2. On this server, copy the backups of the locking specification files for replication management to the following directory:

Windows version:

```
<Directory-determined-in-step1>\<server-whose-files-to-be-restored>\ReplData\data
```

Solaris, Linux, HP-UX or AIX version:

```
<Directory-determined-in-step1>/<server-whose-files-to-be-restored>/ReplData/data
```

3. Transfer the backups of the locking specification files for replication management copied in step 2 to the target server in binary mode.

4. Copy the backups of the locking specification files for replication management transferred in step 3 to the target server. For the directory in which the locking specification files for replication management are used (that is, the directory to which the files are to be copied), refer to "Pre-processing and Post-processing of Replication" in the "ETERNUS SF AdvancedCopy Manager Operation Guide (for Windows)".

Restoring a policy file for management information package backup

1. Execute stgmgrinfoset (Management information package backup server registration/deletion command) on the management information package backup execution server to determine which directory contains the backup data.

2. On this server, copy the policy file located in the following directory to the policy file for the management information package backup command:

Windows version:

```
<Directory-determined-in-step1>\<server-whose-file-to-be-restored>\MgrInfoData
```

Solaris, Linux, HP-UX or AIX version:

```
<Directory-determined-in-step1>/<server-whose-file-to-be-restored>/MgrInfoData
```

For information on the directory for the policy file for the management information package backup command, refer to step 6 in the "8.1.2.1 Preliminaries".

8.2 Troubleshooting

8.2.1 Troubleshooting: Insufficient free space in a repository

This section describes the actions to take if there is insufficient free space in a repository used by a Management Server.
Perform this procedure on a Management Server.

Point

You can use this troubleshooting method when there is insufficient free space in a repository, or if the operation status is significantly different from the status at the initial installation, in order to ensure that there is sufficient free space in the repository in the future.

Increase the amount of repository space as follows:

1. Back up the registry of a Management Server (refer to "8.1.1.4.1 Backing up a registry")

2. Perform stored data backup processing (refer to "8.2.1.1 Database backup")

3. Release a repository (refer to "Cancel of repository settings" in the "ETERNUS SF Installation and Setup Guide")

4. Re-estimate the database size (refer to "Estimation of database size" in the "ETERNUS SF Installation and Setup Guide")

5. Re-create partitions (refer to "Setup of ETERNUS SF Manager" in the "ETERNUS SF Installation and Setup Guide")

6. Create a repository (refer to "Setup of ETERNUS SF Manager" in the "ETERNUS SF Installation and Setup Guide")

7. Starting of a RDB daemon (refer to "9.5.3 Starting and stopping the RDB daemon")

8. Perform saved data restoration processing (refer to "8.2.1.2 Database restoration")

9. Perform restoration of the registry of a Management Server (refer to "8.1.1.4.2 Restoring a registry")

8.2.1.1 Database backup

This section describes the backing up of a database. You can save data stored in a database to the specified directory. This method is mainly used for copying data for reuse for other purposes. For ordinary backup purposes, use the method described in "8.1.1.3.1 Saving a database".

📝 Note

- -

- A directory in which data will be saved must be created in advance.

- To start the database backup, perform as a root user.

- Use the location specified at Management Server installation.

- -

Back up a database as follows:

1. Start database backup.
   Enter the following command to start database backup:

   ```
    # /opt/swstorage/bin/stgdbbackup
   ```

2. Confirm to continue processing.
   Enter "y" to continue processing or "n" to stop processing.

   ```
   ////////////////////////////////////////////////////////////////
         AdvancedCopy Manager Storage Management Server ( kiku )
                           Database backup
   ////////////////////////////////////////////////////////////////


      +-------------------------------------------------------------+
      |                                                             |
      |    All of the stored data will be saved in an external file. |
      |                                                             |
      |   [Notes]                                                   |
      |      Prepare a directory for storing the data to be saved.  |
      |                                                             |
      +-------------------------------------------------------------+


            Do you want to continue with processing? [y/n] ==> y
   ```

3. Enter the name of the data storage directory the data will be saved in.
   Using an absolute pathname, enter the name of the directory in which to store save data.

📝 Note

- -

If the data has already been saved in the directory the processing is not performed.
Delete the save data in the directory or specify another directory, and then retry processing.

- -

```
          //////////////////////////////////////////////////////////////
                  AdvancedCopy Manager Storage Management Server ( kiku )
                                      Database backup
          //////////////////////////////////////////////////////////////

                  Enter the absolute path name of the directory
                  for storing the data to be saved.

                              Enter ==> /storage/backdir
```

4. To continue with database backup processing, enter a final confirmation.
   Enter "y" to continue with database backup processing or "n" to stop it.
   If "y" is entered, the window display returns that of step 2.

```
          //////////////////////////////////////////////////////////////
                  AdvancedCopy Manager Storage Management Server ( kiku )
                                      Database backup
          //////////////////////////////////////////////////////////////

                  Directory containing the saved data : /storage/backdir

                  Do you want to continue with processing? [y/n] ==> y
```

5. Database backup copies are created.
   Database backup copies are created, and execution results are displayed on the window as necessary.
   An example of such execution results is shown below.

```
     ////// AdvancedCopy Manager database backup //////

     ================================================================
     DB data backup will start.
     ================================================================
     ....
     Unloading of data is complete.(base)
     ...
     Unloading of data is complete.(backup)
     .
     Unloading of data is complete.(repl)


     ================================================================
     Backup of DB data is complete.
     ================================================================

     Database backup ended normally.

     #
```

8.2.1.2 Database restoration

This section describes the restoration of a database. You can restore data that has been saved as described in "8.2.1.1 Database backup" to a database as shown below.

Note

. . . .

- Data currently stored in a database will be lost in this processing because it will be overwritten by saved data.

- To start the database restoration, perform as a root user.

- Use the location specified at Management Server installation.

- When the database restoration has completed normally, the database should then be backed up to save this copy of the data. For more information on this, refer to "8.1.1.3.1 Saving a database".

Restore a database as follows:

1. Start database restore processing.
   Enter the following command to start database restore processing:

   ```
   # /opt/swstorage/bin/stgdbrestore
   ```

2. Confirm to continue processing.
   Enter "y" to continue processing or "n" to stop processing.

   ```
   ////////////////////////////////////////////////////////////////////
         AdvancedCopy Manager Storage Management Server ( kiku )
                           Database restore
   ////////////////////////////////////////////////////////////////////


   +-----------------------------------------------------------------+
   |                                                                 |
   |    All of the saved data will be restored to database.          |
   |                                                                 |
   |   [Notes]                                                       |
   |      Check the directory containing the saved data.             |
   |                                                                 |
   +-----------------------------------------------------------------+


             Do you want to continue with processing? [y/n] ==> y
   ```

3. Enter the name of the data storage directory where the data has been saved.
   Using an absolute pathname, enter the name of the directory in which the data is saved.

   Note

   If any part of saved data in the directory is missing, processing is not performed because data consistency may not be maintained.

   ```
   ////////////////////////////////////////////////////////////////////
         AdvancedCopy Manager Storage Management Server ( kiku )
                           Database restore
   ////////////////////////////////////////////////////////////////////


             Enter the absolute path name of the directory
             in which the saved data is stored.

                   Enter ==> /storage/backdir
   ```

4. To continue with database restore processing, enter the final confirmation.
   Enter "y" to continue with database restore processing or "n" to stop it.
   If "n" is entered, the window display returns to that of step 2.

   ```
   ////////////////////////////////////////////////////////////////////
         AdvancedCopy Manager Storage Management Server ( kiku )
                           Database restore
   ////////////////////////////////////////////////////////////////////


             Directory containing the saved data : /storage/backdir
   ```

```
            Do you want to continue with processing? [y/n] ==> y
```

5. Database restore processing is performed.
   Database restore processing is performed, and execution results are displayed on the window as necessary. Sample results are shown below.

```
////// AdvancedCopy Manager database restore //////


================================================================
DB data will be restored.
================================================================
....
Loading of data is complete.(base)
....
Updating of the optimal information is complete.(base)
...
Loading of data is complete.(backup)
...
Updating of the optimal information is complete.(backup)
.
Loading of data is complete.(repl)
.
Updating of the optimal information is complete.(repl)


================================================================
DB data has been completely restored.
================================================================

Database restore processing ended normally.

#
```

8.2.2 Collecting maintenance information

8.2.2.1 Collecting error survey information

Collecting error survey information on the Management Server

The collecting error survey information on the Management Server can be performed using either the command or Web Console.

When using the command, login to the Management Server and execute esfsnap (AdvancedCopy Manager's manager troubleshooting data collection command).
When using Web Console, refer to "Collecting troubleshooting data" in the "ETERNUS SF Web Console Guide".

Send the collected error survey information to a Fujitsu system engineer.

Note

If an error occurs during installing ETERNUS SF, esfsnap (AdvancedCopy Manager's manager troubleshooting data collection command) may not exist. In this case, collect the following data.

- System log

- The files under the following directories.

    - /var/opt/FJSVswtrc/trc

    - /var/opt/FJSVswstf

- /etc/opt/swstorage

Collecting error survey information on the Managed Server

Login to the Managed Server and execute acmcapture (AdvancedCopy Manager's agent troubleshooting data collection command) on each server.

Please send the collected error survey information to a Fujitsu system engineer.

8.2.3 Troubleshooting: Failure in accessing a repository database

This section describes how to determine the cause of a failure in accessing a repository database and the actions to perform when this occurs. Execute the steps below until the problem is resolved and no further steps are required.
Perform this procedure on the Management Server.

1. Check whether the RDB daemons are still running.

   Check method

   Execute the following command to check for any process related to the RDB daemons.

   ```
   # ps -ef | grep RDBSWSTF | grep -v grep
   ```

   Required action

   If no process related to the three RDB daemons exists, retry processing after restarting (start after stopping) the daemons. For an explanation on restarting the daemons, refer to "9.5.3 Starting and stopping the RDB daemon".

2. Check that there is sufficient free space as required to update a repository database.

   Check method

   Execute the following command to check the usage ratio of space required to update a repository.

   ```
   # /opt/swstorage/bin/stgdbloginf
   ```

   Required action

   If the usage ratio is 100%, save the database space according to "8.1.1.3.1 Saving a database". (Consequently, there should be sufficient free space required to update a repository.)
   After using "8.1.1.3.1 Saving a database", restart (start after stopping) the RMI daemon. For an explanation on restarting the daemon, refer to "9.5.2 Starting and stopping the RMI daemon".
   Then, perform this step again.

3. Check whether the "8.1.1.3.1 Saving a database" process is in progress.

   Check method

   Execute the following command to check for any process related to "8.1.1.3.1 Saving a database".

   ```
   # ps -ef | grep stgdbdmp | grep -v grep
   ```

   Required action

   If any process related to "8.1.1.3.1 Saving a database" exists, retry processing after waiting for the end of "8.1.1.3.1 Saving a database".

4. Check whether the capacity of repository is insufficient.

Check method

In the log file a /var/opt/FJSVswstf/log/RDBSWSTF.log file (in cluster operation, they are /var/opt/FJSVswstf/*<logical-node-name>*/log/RDBSWSTF.log), search for the strings "JYP5019E" or "JYP5045E" in any message taking the form "rdb: ERROR: qdgXXXXXX-"

Required action

When a character sequence exists, please extend repository space with reference to "8.2.1 Troubleshooting: Insufficient free space in a repository".
Then, Retry processing.

5. Check whether an I/O error or any other error has occurred in the database space.

Check method

Check the /var/opt/FJSVswstf/log/RDBSWSTF.log file, and beginning from the last line, search for any message with "rdb: ERROR: qdgXXXXX-" in its contents (also check the messages displayed at the same time). Then, search for a string taking the form of "qdgXXXXX," and check for it in "Relationships between the qdg message and recovery mode".

Required action

If "qdgXXXXX" is in the table, examine the target according to the error description corresponding to the qdg messages in "Relationships between the qdg message and recovery mode", and determine if there is a problem. If a problem is found, take corrective action as appropriate to solve the problem, and specify the appropriate settings. Then, recover the database space by using "8.1.1.3.3 Recovering a database".
After the database recovery command ends normally, acquire the latest save data by using "8.1.1.3.1 Saving a database".
Then, restart (start after stopping) the AdvancedCopy Manager daemons. For an explanation on restarting the daemons, refer to "9.5.4 Starting and stopping AdvancedCopy Manager daemons".
Retry processing.

📝 Note

. . . .

- If the database recovery command is executed with the value "0" specified for -m option, perform the following tasks after the command ends normally:

  - Perform the resource consistency command specified with -x option (refer to swstsrsemtch (Resource match command)) on all Managed Servers in a backup operation.

  - Perform the resources adjustment command (refer to swsrprecoverres (Resource adjustment command)) on the Managed Servers that are performing replication using the resources adjustment command.

- If an error such as an I/O error occurs in a database space, a file required for analyzing the cause of the error may be created in the *<DB file directory>*/SWSTFDB/core. If the corresponding recovery ends normally, delete the file in this directory.

. . . .

6. Collect the files in /var/opt/FJSVswstf/log, and contact a Fujitsu system engineer.

8.3 Troubleshooting: Hardware error during backup

This section describes the actions to perform if a hardware error, etc. occurs during a backup using AdvancedCopy Manager.

8.3.1 Overview

If a hardware error occurs during backup, the following symptoms may be displayed:

- swstbackstat (Backup execution status display command) displays "failed" in the Status column.

- swstsyncstat (Backup synchronous processing progress display command) displays "failed" in the Status column.

- swstreststat (Restore execution status display command) displays "failed" in the Status column.

- swsthistdisp (History information display command) displays "failed" in the Status column.

- swsttrkstat (Tracking status display command) displays "failed" in the Status column.

- The message swst0634, swst0719 or swst0740 is displayed during command execution.

If the execution status display command and history information display command display "failed" in the Status column, a hardware error has occurred during physical copying, and the copy processing has failed.
If the message swst0634, swst0719 or swst0740 is displayed during command execution and the same error occurs even if the command is re-executed, a hardware error has occurred, causing Advanced Copy to fail.

If any of these messages, statuses or symptoms is generated, a hardware error may have occurred in the transaction volume or backup volume. Check the hardware status, and if a hardware error is found, fix the cause of the error, then re-execute the processing.

8.3.2 Troubleshooting

The following figure shows the flow of troubleshooting tasks if a hardware or similar fault occurs.

Figure 8.2 Troubleshooting flow (for faults during backup)
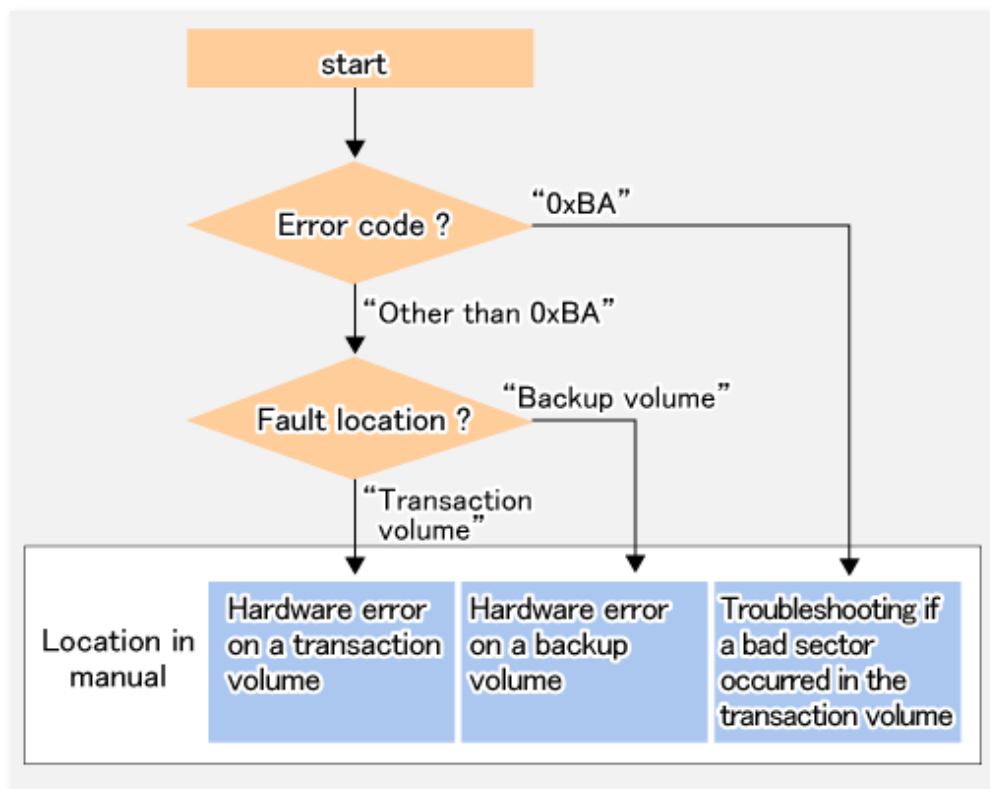
Note

- Refer to "8.3.1 Overview" for details of the "status column" and "fault location".

- Use ETERNUS Web GUI to check the copy process error codes. Use the following method to check the error codes:

    - On the [Display status] menu, click [Advanced Copy status display] in the status display.

    - At "Session status", click the "Number of active sessions" link for the relevant copy type.

    - Refer to the value in the "Error code" column of the relevant copy process.

    The following table shows the meanings of the error codes.

Table 8.8 Meanings of error codes

Error code	Meaning

0xBA	If a) or b) below applies, a bad sector was created in the transaction volume.
	
	  a. QuickOPC has not yet performed physical copying and tracking is in progress
	
	  b. EC/REC is in the suspend status (replication established status)
	
	Note:
	If a bad sector is created in a transaction volume when a) or b) applies, the ETERNUS Disk storage system automatically changes the copy processing to the error suspend state. This prevents a restart of QuickOPC or EC/REC resume and prevents the backup volume from being overwritten with invalid transaction volume data.

Other than 0xBA	An error other than the above occurred.

Error code 0xBA is returned only for the following ETERNUS Disk storage system:

- ETERNUS4000 and ETERNUS8000 (firmware version V11L30-0000 or later)

For ETERNUS Disk storage system other than the above, the event indicated by error code 0xBA does not occur. (When the above a) or b) apply, the copy status does not change even if a bad sector occurs in the copy source volume.)

....

8.3.2.1 Hardware error on a transaction volume

When a hardware fault occurs in the transaction volume, perform the repair work according to the following procedures.

🛈 **Note**

....

- The backup history information for which a copy failed cannot be used to perform a restoration.

- If there is no normal (that is, Status is "succeeded") backup history information, restoration is not possible.

....

1. If an error occurred during physical copying, cancel the processing in which the error occurred, with one of the following commands:

    - If an error occurred during backup (OPC) physical copy: swsthistdel (History information deletion command)

    - If an error occurred during synchronous processing (EC): swstcancelsync (Backup synchronous processing cancel command)

    - If an error occurred during restore (OPC) physical copy: swstcancelrest (Restore cancellation command)

    - If an error occurred during tracking processing (QuickOPC): swstcanceltrk (Tracking cancel command)

2. Execute the execution status display commands (swstbackstat (Backup execution status display command), swstsyncstat (Backup synchronous processing progress display command), swstreststat (Restore execution status display command), swsttrkstat (Tracking status display command) and swsthistdisp (History information display command)) to verify that no other errors have occurred.

3. Fix the hardware error in the transaction volume.

4. Execute swstrestore (Restore execution command) to restore the transaction volume.

8.3.2.2 Hardware error on a backup volume

When a hardware fault occurs in backup volume, perform repair work according to the following procedures.

1. Cancel the processing in which the error occurred, with one of the following commands:

    - If an error occurred during backup (OPC) physical copy: swsthistdel (History information deletion command)

    - If an error occurred during synchronous processing (EC): swstcancelsync (Backup synchronous processing cancel command)

- If an error occurred during restoration (OPC) physical copy: swstcancelrest (Restore cancellation command)

- If an error occurred during tracking processing (QuickOPC): swstcanceltrk (Tracking cancel command)

If the processing cannot be cancelled using one of the above commands, use ETERNUS Web GUI to cancel it.

2. Execute swstsrsemtch (Resource match command).

3. Cancel the processing in which the error occurred. At this point, the execution status display commands (swstbackstat (Backup execution status display command), swstreststat (Restore execution status display command) and swsthistdisp (History information display command)) display "succeeded" in the Status column. However, because the copy processing is suspended, ensure that the commands are cancelled. If an error occurred in the synchronous processing (EC), the command has already been cancelled, so no response is required.

- If an error occurred during backup (OPC) physical copy: swsthistdel (History information deletion command)

- If an error occurred during restoration (OPC) physical copy: swstcancelrest (Restore cancellation command)

4. Execute the execution status display commands (swstbackstat (Backup execution status display command), swstsyncstat (Backup synchronous processing progress display command), swstreststat (Restore execution status display command), swsttrkstat (Tracking status display command) and swsthistdisp (History information display command)) to verify that no other errors have occurred.

5. Use swstdevinfoset (Device information setting command) to delete the backup volume in which the error occurred.

6. Use swstdevinfoset (Device information setting command) to register a new backup volume. If the backup volume in which the error occurred is repaired and reused, execute the option [Collect or reflect the information for a specific device] from the Web Console, and store the information again in the backup volume.

7. Re-execute the processing in which the error occurred.

8.3.2.3 Troubleshooting if a bad sector occurred in the transaction volume

If a bad sector occurred in the transaction volume, use the following procedure to restore the transaction volume:

1. Cancel the copy processing for which the bad sector occurred.

- If the bad sector occurred during the suspend state of synchronous processing (EC): swstcancelsync (Backup synchronous processing cancel command)

- If the bad sector occurred during QuickOPC tracking: swstcanceltrk (Tracking cancel command)

2. Use swstbackstat (Backup execution status display command), swstsyncstat (Backup synchronous processing progress display command), swstreststat (Restore execution status display command), swsttrkstat (Tracking status display command) and swsthistdisp (History information display command) to check for other errors.

3. Restoration is performed by overwriting the area containing the bad sector. Select the appropriate method, in accordance with the usage or use status of the transaction volume, from the methods below.

- Restoration method 1:

If the area can be reconstructed from high-level software (file system, DBMS, or similar), reconstruct the area.

- Restoration method 2:

If the area containing the bad sector is an area that is not being used, such as an unused area or a temporary area, use a system command (for example, the UNIX dd command or the Windows format command) to write to the area.

- Restoration method 3:

Use swstrestore (Restore execution command) to restore the transaction volume from a backup volume. (If the backup volume for the copy processing for which the bad sector occurred contains backup history information, restoration is also possible from that backup volume.)

 Note

The transaction volume cannot be restored by this method, if there is no backup history information.

8.4 Hardware or Other Error during Replication

This section describes the corrective action to take when a hardware or other type of error occurs during AdvancedCopy Manager replication.

8.4.1 Overview

If a hardware error occurs during the replication, the following symptoms are displayed:

- swsrpstat (Operation status display command) displays "failed" or "?????" in the Status column.

- The message swsrp2606 is output during command execution.

- The message swsrp0634, swsrp0719 or swsrp0740 is displayed during command execution.

If the operation status display command displays "failed" in the Status column and swsrp2606 is displayed during command execution, a hardware error has occurred during physical copying, and the copy processing has failed.
If the message swsrp0634, swsrp0719 or swsrp0740 is displayed during command execution, and the same error occurs even if the command is re-executed, a hardware error has occurred, causing Advanced Copy to fail.

If a hardware error occurs during remote copying, the following symptoms may be displayed:

- swsrpstat (Operation status display command) displays "halt" in the Status column.

- The message swsrp2648 is displayed during command execution.

If the operation status display command displays "halt" in the Status column, and the message swsrp2648 is displayed during command execution, a hardware error has occurred during physical copying, and the remote copy processing may have failed.
If any of these messages, statuses or symptoms is generated, a hardware error or path close may have occurred in the source or replica volume. Check the hardware status, and if a hardware error is found, fix the cause of the error and re-execute the command.

8.4.2 Troubleshooting

The following figure shows the troubleshooting flow when a hardware or similar fault occurs.

Figure 8.3 Troubleshooting flow (when a fault occurs during replication)
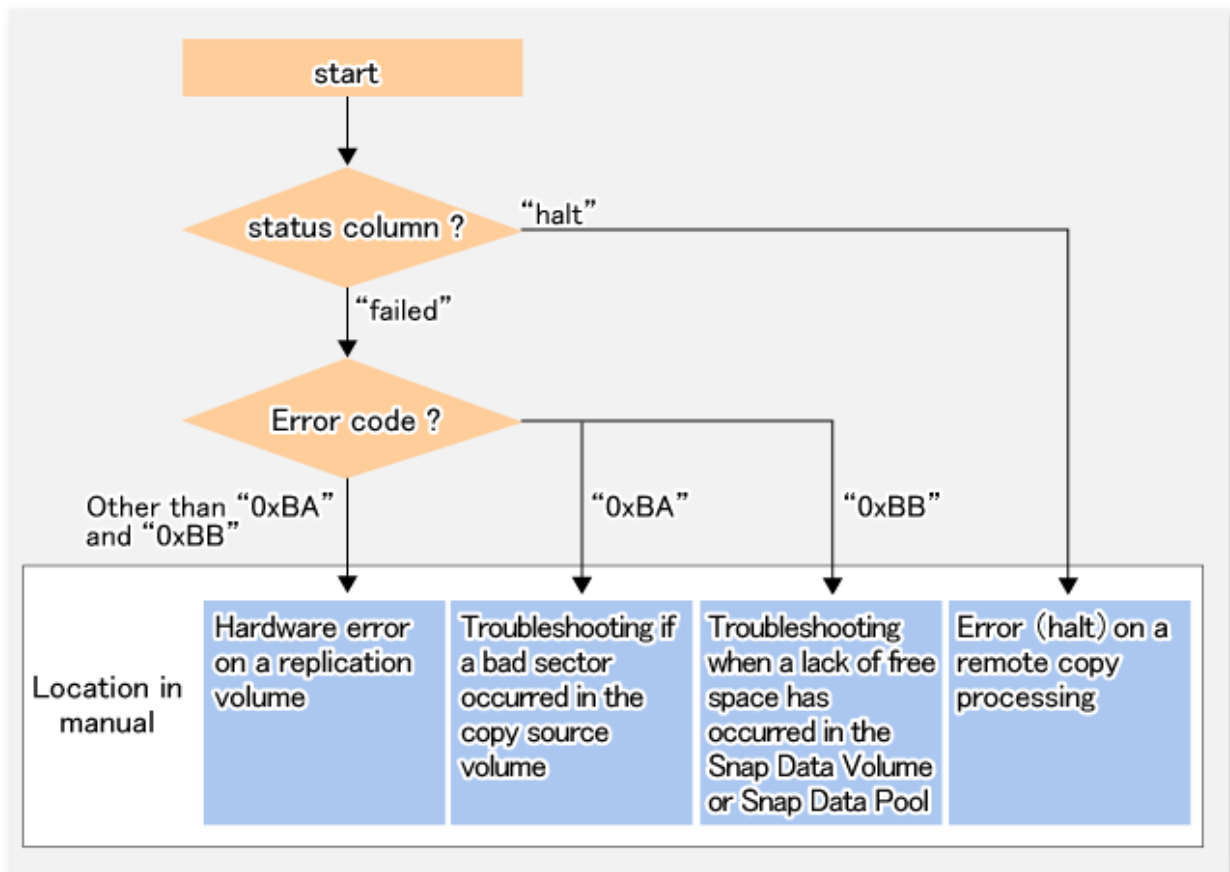
Note

- Refer to "Overview" for details of the "status column".

  If the status column is "?????", check if ETERNUS Web GUI copy processing is in the error suspend status ("failed") or the hardware suspend status ("halt").

  If the copy processing is in either of these states, take the action indicated in the above troubleshooting flow.

  In other cases where another cause is likely (such as the Managed Server or switch etc) contact your Fujitsu SE.

- Use ETERNUS Web GUI to check the error codes. Use the following two methods to check.

  - Checking with swsrpstat (Operation status display command)

    Execute the command with the -O option.

  - Checking with ETERNUS Web GUI

    1. On the [Display status] menu, click [Advanced Copy status display] in the status display.

    2. At "Session status", click the "Number of active sessions" link for the relevant copy type.

    3. Refer to the value in the "Error code" column of the relevant copy process.

  The following table shows the meanings of the error codes.

Table 8.9 Meanings of error codes

Error code	Meaning
0xBA	If a) or b) below applies, a bad sector was created in the transaction volume.
	a. QuickOPC has not yet performed physical copying and tracking is in progress
	b. EC/REC is in the suspend status (replication established status)

Error code	Meaning
	Note: If a bad sector is created in a transaction volume when a) or b) applies, the ETERNUS Disk storage system automatically changes the copy processing to the error suspend state. This prevents a restart of QuickOPC or EC/REC resume and prevents the copy destination volume from being overwritten with invalid copy source volume data.
0xBB	A lack of free space has occurred in the Snap Data Volume or Snap Data Pool
Other than 0xBA and 0xBB	An error other than the above occurred.

Error codes 0xBA and 0xBB are returned only for the following ETERNUS Disk storage system:

- ETERNUS4000 and ETERNUS8000 (firmware version V11L30-0000 or later)

For ETERNUS Disk storage system other than the above, the events indicated by error code 0xBA and 0xBB are identified by the following methods:

Table 8.10 Error code events

Event	Identification method
Events indicated by 0xBA	These events do not occur. In cases a) and b) above, the copy status does not change even if a bad sector occurs at the copy source volume.
Events indicated by 0xBB	Use ETERNUS Web GUI to check the capacity already used on the Snap Data Volume in order to determine whether or not a lack of free space has occurred. - On the [Display status] menu, click [Volume list] in the status display. - Click the link to Snap Data Volume in the "Volume type" column of the relevant volume. - Refer to the value shown in the "Capacity already used" column. If this event applies, refer to "8.4.2.3 Troubleshooting when a lack of free space has occurred in the Snap Data Volume or Snap Data Pool".

8.4.2.1 Hardware error on a replication volume

When a hardware error occurs in a duplicate volume, perform the repair work on the error according to the following procedures.

1. Use swsrpcancel (Replication cancellation command) to cancel the processing in which the error occurred. If the processing cannot be cancelled from the operation server when inter-server replication is performed, cancel it from a non-operational server.

2. If the processing cannot be cancelled by using the command, use ETERNUS Web GUI to cancel it.

3. Execute swsrprecoverres (Resource adjustment command)

4. Execute swsrpstat (Operation status display command) to verify that no other errors have occurred.

5. Use swsrpdelvol (Replication volume information deletion command) to delete the replication volume in which the error occurred.

6. Use swsrpsetvol (Replication volume information setting command) to register a new replication volume. If the replication volume in which the error occurred is repaired and reused, execute the option [Collect or reflect the information for a specific device] from the Web Console and store the information again in the replication volume.

7. Re-execute the processing in which the error occurred.

8.4.2.2 Troubleshooting if a bad sector occurred in the copy source volume

If a bad sector occurred in the copy source volume, use the following procedure to restore the copy source volume:

1. Use swsrpcancel (Replication cancellation command) to cancel processing for which the error occurred.
   If inter-server replication was being performed and cancellation is not possible from the active server, cancel processing from the inactive server.
   If processing cannot be cancelled using commands, use ETERNUS Web GUI to cancel it.

2. Execute swsrpstat (Operation status display command) to check for other errors.

3. Restoration is performed by overwriting the area containing the bad sector. Select the appropriate method, in accordance with the usage or use status of the copy source volume, from the methods below.

   - Restoration method 1
     If the area can be reconstructed from high-level software (file system, DBMS, or similar), reconstruct the area.

   - Restoration method 2
     If the area containing the bad sector is an area that is not being used, such as an unused area or a temporary area, use a special-purpose tool (for example, the UNIX dd command) to write to the area.

   - Restoration method 3
     Use swsrpmake (Replication creation command) to restore the data from the copy destination volume. (Restoration is also possible from the copy destination volume of the copy process for which the bad sector occurred.)

8.4.2.3 Troubleshooting when a lack of free space has occurred in the Snap Data Volume or Snap Data Pool

A Snap Data Volume lack of free space occurs when the Snap Data Pool is not being used, whereas a Snap Data Pool lack of free space occurs when the Snap Data Pool is being used.
The following section provides information on recovery when a lack of free space has occurred in either the Snap Data Volume or Snap Data Pool.

Recovery of insufficient free space in Snap Data Volume

When a lack of free space has occurred in the Snap Data Volume, follow these steps to undertake recovery:

- Cancel the processing in which the error occurred with swsrpcancel (Replication cancellation command) command.

- If inter-server replication was being performed and cancellation is not possible from the active server, cancel processing from the inactive server.

- If processing cannot be cancelled using commands, use ETERNUS Web GUI to cancel it.

The likely causes of a lack of free space in the Snap Data Volume are as follows:

a. The estimate of the physical size of the Snap Data Volume is not accurate.

b. The estimate of the physical size of the Snap Data Volume is accurate but, as a result of a large volume being updated in the Snap Data Volume when a SnapOPC/SnapOPC+ session does not exist, the physical capacity of the Snap Data Volume is being used up.

The usage status of the Snap Data Volume can be checked by specifying "stat" in swstsdv (Snap Data Volume operation/reference command) subcommand.

If "a." applies, re-estimate the physical size of the Snap Data Volume, and recreate the Snap Data Volume.
If "b." applies, use ETERNUS Web GUI or, specify "init" in swstsdv (Snap Data Volume operation/reference command) subcommand, and then initialize the Snap Data Volume.
Recreation of the partition (slice) is required after recreation/initialization of the Snap Data Volume.

Recovery of insufficient free space in Snap Data Pool

When a lack of free space has occurred in the Snap Data Pool, follow these steps to undertake recovery:

- Cancel the processing in which the error occurred with swsrpcancel (Replication cancellation command).

- If inter-server replication was being performed and cancellation is not possible from the active server, cancel processing from the inactive server.

- If processing cannot be cancelled using commands, use ETERNUS Web GUI to cancel it.

The following are likely causes of a lack of free space in the Snap Data Pool:

    a. The estimate of the size of the Snap Data Pool is not accurate.

    b. The estimate of the size of the Snap Data Pool is accurate but, as a result of a large volume being updated in the Snap Data Volume when a SnapOPC/SnapOPC+ session does not exist, the capacity of the Snap Data Pool is being used up.

The use status of the Snap Data Pool can be checked by specifying "poolstat" in swstsdv (Snap Data Volume operation/reference command) subcommand.

If "a." applies, re-estimate the size of the Snap Data Pool, and after increasing the size of the Snap Data Pool, recreate the Snap Data Volume.

If "b." applies, use ETERNUS Web GUI or, specify "init" in swstsdv (Snap Data Volume operation/reference command) subcommand, then initialize the Snap Data Volume.

Recreation of the partition (slice) is required after recreation/initialization of the Snap Data Pool.

8.4.2.4 Error (halt) on a remote copy processing

The REC restart (Resume) method varies, depending on the halt status.

Execute swsrpstat (Operation status display command) with the -H option specified to check the halt status, and then implement the relevant countermeasure.

- For "halt(use-disk-buffer)" or "halt(use-buffer)"

This status means that data is saved to the REC Disk buffer or REC buffer because data cannot be transferred due to a path closure (halt). In order to restart REC, perform path recovery before a space shortage occurs for the REC Disk buffer or REC buffer.
After recovery, the ETERNUS Disk storage system restarts REC automatically.

If a space shortage has already occurred for the REC Disk buffer or REC buffer, the "halt(sync) or halt (equivalent)" status shown below occurs. Implement the countermeasures for that status.

- For "halt(sync) or halt(equivalent)"

This status means that data transfer processing was discontinued due to a path closure (halt).

The REC restart method differs for different REC Recovery modes.

For the Automatic Recovery mode

    1. Remove the cause that made all paths close (halt).

    2. ETERNUS Disk storage system automatically restarts (Resume) REC.

For the Manual Recovery mode

    1. Remove the cause that made all paths close (halt).

    2. Use swsrpmake (Replication creation command) to forcibly suspend the REC that is in the halt status.

```
[For volume units]
swsrpmake -j < replication source volume name > <replication destination
volume name >

[For group units]
swsrpmake -j -Xgroup <group name>
```

    3. Use swsrpstartsync (Synchronous processing start command) to restart (Resume) the REC. The -t option must be specified if REC is being restarted after a forcible suspend.

```
[For volume units]
swsrpstartsync -t <replication source volume name > <replication destination
volume name >

[For group units]
swsrpstartsync -t -Xgroup <group name>
```

8.5 Changing the Operating Environment

Server information batch change operations preserve the consistency of server information in systems managed by AdvancedCopy Manager. All Managed Server services and transactions in the systems that include a Management Server and Management Server transactions must be running when these changes are made.

Note

- - If the server information change command (stgxfwcmmodsrv) is used to change IP addresses, or similar.

  Do not execute backup management commands or replication management commands at the Management Server that executes stgxfwcmmodsrv (Server information change command), nor at any of the Managed Servers managed by Management Server transactions, either before or during execution of stgxfwcmmodsrv (Server information change command).

  - - If executed, this command may terminate abnormally. If stgxfwcmmodsrv (Server information change command) terminates abnormally, implement recovery in accordance with the action indicated in the error message. If backup management or replication management commands are accidentally executed either before or during execution of the server information change commands.

  - - If the backup management and replication management commands terminate abnormally, wait until stgxfwcmmodsrv (Server information change command) terminates normally, then re-execute the backup management and replication management commands.

8.5.1 Changing the IP address of Management Server

Change the IP address of a Management Server as follows:
If only a physical IP address needs to be changed and the Management Server uses a logical IP address, then perform steps 6 only. Otherwise, perform all steps.

1. If replication management function was used at the Management Server to define it as a copy source or copy destination volume, delete the copy source or copy destination volume settings by referring to "7.6.2 Deleting a source volume or replica volume".

2. On the Management Server, change the system IP address.

3. On the Management Server, restart the AdvancedCopy Manager daemon. For information on restarting a daemons, refer to "Chapter 2 Starting and Stopping Daemons".
   Note that his step is not required if the system was restarted in Step 1.

4. On the Management Server, create a server information change instruction file, and then execute stgxfwcmmodsrv (Server information change command) with the -f option specified to change the IP address.

5. Restart all AdvancedCopy Manager daemons on the Management Server. Refer to "Chapter 2 Starting and Stopping Daemons" for details.

6. If the replication management function is used, reconstruct the replication management environment (perform the "7.4.5 Configuring the source volume and replica volume" task and subsequent tasks).

8.5.2 Changing the IP address of Managed Server

Change the IP address of a Managed Server as follows:

1. If replication management function was used at the target Managed Server to define it as a copy source or copy destination volume, delete the copy source or copy destination volume settings by referring to "7.6.2 Deleting a source volume or replica volume".

2. On the target Managed Server, stop the AdvancedCopy Manager daemons. For more information on stopping daemons, refer to "Chapter 2 Starting and Stopping Daemons".

3. On the target Managed Server, change the system IP address.

4. On the target Managed Server, restart the AdvancedCopy Manager daemons. For more information on starting a daemons, refer to "Chapter 2 Starting and Stopping Daemons".

5. On the Management Server, create a server information change instruction file, and then execute stgxfwcmmodsrv (Server information change command) with the -f option specified to change the IP address.

6. If the replication management function is used, reconstruct the replication management environment (perform the "7.4.5 Configuring the source volume and replica volume" task and subsequent tasks).

8.5.3 Changing the port number of Management Server or Managed Server

Change the port number of a Management Server or Managed Server as follows:

1. If replication management function was used at the target Management Server or Managed Server to define it as a copy source or copy destination volume, delete the copy source or copy destination volume settings by referring to "7.6.2 Deleting a source volume or replica volume".

2. On the target Management Server or Managed Server, change the port number of the communication daemon (stgxfws) of AdvancedCopy Manager indicated by /etc/services.

3. On the target Management Server or Managed Server, restart the AdvancedCopy Manager daemons. For information on restarting daemons, refer to "Chapter 2 Starting and Stopping Daemons".

4. On the Management Server, create a server information change instruction file, and then execute stgxfwcmmodsrv (Server information change command) with the -f option specified to change the port number.

5. If the replication management function is used, reconstruct the replication management environment (perform the "7.4.5 Configuring the source volume and replica volume" task and subsequent tasks).

8.5.4 Changing the server name of Management Server or Managed Server

Change the server name of a Management Server or Managed Server as follows:

1. If replication management function was used at the target Management Server or Managed Server to define it as a copy source or copy destination volume, delete the copy source or copy destination volume settings by referring to "7.6.2 Deleting a source volume or replica volume".

2. Change the system server name at the Management Server or Managed Server that is to be changed.
   If the server name to be changed is only that of a Management Server or Managed Server managed by AdvancedCopy Manager, this step is not required.

3. On the Management Server, create a server information change instruction file, and then execute stgxfwcmmodsrv (Server information change command) with the -f option specified to change the server name.

4. If the replication management function is used, reconstruct the replication management environment (perform the "7.4.5 Configuring the source volume and replica volume" task and subsequent tasks).

See
. . . .
For information about performing the above change in a cluster operation, refer to "8.5.7.4 Changing the server name of Management Server transaction or Managed Server transaction".
. . . .

8.5.5 Changing the device configuration

Changing the device configuration refers to any of the following:

- Mount name change

- Change from File System operation to RawDevice operation

- Capacity change of partitions (mainly for re-creation of LUN mapping)

- Start position for change of partitions (mainly for re-creation of LUN mapping)

Note
. . . .
For correct backup management, the backup volume and the transaction volume must be the same size. For more information on this, refer to "Size of a Logical Disk" in "10.1.1 General notes".
. . . .

8.5.5.1 Backup management

Change the device configuration by using the following procedure. This processing is executed for all volumes on the disk containing the volume whose device configuration is to be changed.

1. Check whether there is a transaction volume on the target disk that needs to be backed up. If there is, delete the transaction volume. For information on this, refer to "3.6.2.1 Deleting a transaction volume".

2. Check whether a backup volume exists in the target disk. If it does, delete the backup volume. For information on this, refer to "3.6.2.2 Deleting a backup volume".

3. Change the device configuration. If the logical drive in the primary partition or extended partition has been deleted, restart the system.

4. Because the post-change device configuration needs to be reflected in the AdvancedCopy Manager data set, execute device information change processing. Refer to "3.4.4 Fetching device information on a Managed Server", for details on changing the device information.

5. Reset the deleted transaction or backup volume in the backup. Refer to "3.4.6 Setting the operation type for a device", for details on adding the transaction volume or backup volume.

P Point

If the device configuration is changed (in step 3 above) before the transaction and backup volumes are deleted (in steps 1 and 2), history information cannot be deleted, and synchronized backup processing cannot be cancelled. This could possibly prevent the deletion of the transaction and backup volumes.
In this case, set the emergency operational mode, and delete the history information or cancel synchronized backup processing.
For details of the emergency operational mode, refer to the section describing swsthistdel (History information deletion command) and swstcancelsync (Backup synchronous processing cancel command).

8.5.5.2 Replication management

Change the device configuration as shown below. This processing applies to all the volumes on a disk whose device configuration is to be changed.

1. Check whether there are any replication volumes on the target disk. If there are, delete them. For information about this, refer to "7.6.2 Deleting a source volume or replica volume".

2. Change the device configuration. If a logical drive in the primary partition or an extended partition is deleted, restart the system.

3. Because the post-change device configuration needs to be reflected in the AdvancedCopy Manager data set, execute device information change processing for AdvancedCopy Manager. Refer to "7.4.4 Fetching device information on a Managed Server", for details about changing the device information.

4. Reset a deleted replica volume in the replication operation. Refer to "7.4.5 Configuring the source volume and replica volume", for details about adding the replica volume.

P Point

If the device configuration is changed (in step 2 above) before the replication volumes are deleted (in step 2), replication processing cannot be cancelled. This could possibly prevent the deletion of the replication volumes.
In this case, set the emergency operational mode, and cancel replication processing.
For details on the emergency operational mode, refer to the section describing swsrpcancel (Replication cancellation command).

8.5.6 Change of password or deletion of the user who was specified as the repository access user

When a user who is specified as the "repository access user" by the environmental setup at the time of the installation of the AdvancedCopy Manager, or when the password of the specified user name is changed by the operating system, the information of a "repository access user" must be reconfigured using stguserset (Repository access user change command).

See

For information about performing the above change in a cluster operation, refer to "8.5.7.5 Deleting a repository accessible user or changing a repository accessible user's password".

8.5.7 Changing the configuration in cluster operation

This section describes how to make a change to the operational configuration in cluster operation.

8.5.7.1 Changing the IP address of Management Server transaction

Refer to "Changing the IP address of Management Server transaction" under "Changing the Operation Environment for Solaris version, Linux version" in the "ETERNUS SF Cluster Environment Setup Guide".

8.5.7.2 Changing the IP address of Managed Server transaction

Change the IP address of a Managed Server transaction by performing the following steps:

1. If replication management function was used at the target Managed Server transaction to define it as a copy source or copy destination volume, delete the copy source or copy destination volume settings by referring to "7.6.2 Deleting a source volume or replica volume".

2. Stop the Managed Server transaction.
   For details on how to stop it, refer to the relevant cluster software manual.

3. Change the IP address resource.
   For details on how to change an IP address resource, refer to the relevant cluster software manual.

4. Start the Managed Server transaction.
   For details on how to start it up, refer to the relevant cluster software manual.

5. On the Management Server that manages the Managed Server transactions, create a server information change instruction file, and then execute stgxfwcmmodsrv (Server information change command) with the -f option specified to change the IP address.

6. If the replication management function is used, reconstruct the replication management environment (perform the "7.4.5 Configuring the source volume and replica volume" task and subsequent tasks).

8.5.7.3 Changing the port number of Management Server transaction or Managed Server transaction

For changing the port number of Management Server transaction, refer to "Changing the port number of Management Server transaction" under "Changing the Operation Environment for Solaris version, Linux version" in the "ETERNUS SF Cluster Environment Setup Guide".

Change the port number of a Managed Server transaction by performing the following steps:

1. If replication management function was used by the target Managed Server transaction to define it as a copy source or copy destination volume, delete the copy source or copy destination volume settings by referring to "7.6.2 Deleting a source volume or replica volume".

2. Stop the Management Server transaction or Managed Server transaction.
   For details on how to stop it, refer to the relevant cluster software manual.

3. On the primary node and secondary node of the target Managed Server transaction, change the port number of the AdvancedCopy Manager communication daemon (stgxfws_logical-node-name) specified in /etc/services.

4. Start the targeted Managed Server transaction.
   For details on how to start it up, refer to the relevant cluster software manual.

5. Create a server information change instruction file at the Management Server that manages the Managed Server. Then, execute stgxfwcmmodsrv (Server information change command) with the -f option specified to change the port number.

6. If the replication management function is used, reconstruct the replication management environment (perform the "7.4.5 Configuring the source volume and replica volume" task and subsequent tasks).

8.5.7.4 Changing the server name of Management Server transaction or Managed Server transaction

For changing the server name of Management Server transaction, refer to "Changing the server name of Management Server transaction" under "Changing the Operation Environment for Solaris version, Linux version" in the "ETERNUS SF Cluster Environment Setup Guide".

Change the server name of a Managed Server transaction by performing the steps below.

1. If replication management function was used by the target Managed Server transaction to define it as a copy source or copy destination volume, delete the copy source or copy destination volume settings by referring to "7.6.2 Deleting a source volume or replica volume".

2. Create a server information change instruction file at the Management Server that manages the Managed Server. Then, execute stgxfwcmmodsrv (Server information change command) with the -f option specified to change the port number.

3. If the replication management function is used, reconstruct the replication management environment (perform the "7.4.5 Configuring the source volume and replica volume" task and subsequent tasks).

8.5.7.5 Deleting a repository accessible user or changing a repository accessible user's password

For changing the deleting a repository accessible user or changing a repository accessible user's password of Management Server transaction, refer to "Deleting a repository accessible user or changing a repository accessible user's password" under "Changing the Operation Environment for Solaris version, Linux version" in the "ETERNUS SF Cluster Environment Setup Guide".

8.5.7.6 Changing the logical node name

For changing the logical node name of Management Server transaction, refer to "Changing the logical node name" under "Changing the Operation Environment for Solaris version, Linux version" in the "ETERNUS SF Cluster Environment Setup Guide".

Change the logical node name of a Managed Server transaction by performing the following steps:

1. Check whether the target Managed Server transaction is operating.
   If the transaction is not operating, start the Managed Server transaction.
   For details on how to start it up, refer to the relevant cluster software manual.

2. Delete the configuration of the Managed Server transaction.
   When using the backup management function, delete the transaction volume and backup volume by referring to "3.6.2 Deleting a device".
   When it is defined as a copy source or copy destination volume by using the replication management function, delete the copy source or copy destination volume settings by referring to "7.6.2 Deleting a source volume or replica volume".

3. Save the following AdvancedCopy Manager data.

   - Pre-processing/post-processing scripts
     Save the pre-processing and post-processing scripts for backup management and replication management.

   For a "package-backup-of-the-management-information execution server" of AdvancedCopy Manager, save the following data:

   - Policy file for package backup of the management information
     For details about the policy file path, refer to "8.1.2.1 Preliminaries".

4. Delete /etc/opt/swstorage/<*logical-node-name*>/swnode.ini on the node operating the transaction.

5. Delete the cluster environment of the target transaction.

   For details about how to delete, refer to "Deletion on Solaris/Linux/HP-UX version clustered systems" in the "ETERNUS SF Cluster Environment Setup Guide".

Note

Do not perform the following step in "Deletion on Solaris/Linux/HP-UX version clustered systems" of the "ETERNUS SF Cluster Environment Setup Guide".

- Step 1 described in "Deletion Procedure"

6. Create the cluster environment of the target transaction.

For details on how to create, refer to "Customization for Solaris/Linux/HP-UX version clustered system" in the "ETERNUS SF Cluster Environment Setup Guide".

Note

When you perform the cluster set up, specify the name of the new logical node.

7. Start the target transaction.
For details on how to start it up, refer to the relevant cluster software manual.

8. Restore the following AdvancedCopy Manager data.

- Pre-processing/post-processing scripts
Restore the pre-processing and post-processing scripts for backup management and replication management.

For a "package-backup-of-the-management-information execution server" of AdvancedCopy Manager, restore the following data:

- Policy file for package backup of the management information
For details about the policy file path, refer to "8.1.2.1 Preliminaries".

9. Change the server name.

Perform the procedure described in "8.5.7.4 Changing the server name of Management Server transaction or Managed Server transaction".

10. To use the backup management or replication management function, rebuild the backup management or replication management environment. (For backup management, perform the "3.4.5 Configuring the environment for a backup operation server" task and subsequent tasks. For replication management, perform the "7.4.5 Configuring the source volume and replica volume" task and subsequent tasks.)

11. For a "package-backup-of-the-management-information execution server" of AdvancedCopy Manager, execute stgmgrinfoset (Management information package backup server registration/deletion command) to register as its server.

8.5.7.7 Changing the transaction name

Change the name of a Managed Server transaction by performing the following steps:

1. Check whether the target Managed Server transaction is operating.
If the transaction is not operating, start the Managed Server transaction.
For details about how to start it up, refer to the relevant cluster software manual.

2. Save the following AdvancedCopy Manager data.

- AdvancedCopy Manager management file

On the node operating the transaction, save /etc/opt/swstorage/<*logical-node-name*>/swnode.ini.

- Backup management list
Make a backup of the backup management list.
For details about how to do this, refer to "8.1.1.1.1 Backing up a backup management list and recovery control file".

- Pre-processing/post-processing scripts
Save the pre-processing and post-processing scripts for backup management and replication management.

For a "package-backup-of-the-management-information execution server" of AdvancedCopy Manager, save the following data:

- Policy file for package backup of the management information
  For details about the policy file path, refer to "8.1.2.1 Preliminaries".

3. Delete /etc/opt/swstorage/<*logical-node-name*>/swnode.ini on the node operating the transaction.

4. Delete the cluster environment of the target transaction.

   For details on how to do this, refer to "Deletion on Solaris/Linux/HP-UX version clustered systems" in the "ETERNUS SF Cluster Environment Setup Guide".

   ## Note
   ....

   Do not perform the following step of "Deletion on Solaris/Linux/HP-UX version clustered systems" in the "ETERNUS SF Cluster Environment Setup Guide".

   - Step 1 described in "Deletion Procedure"
   ....

5. On the clustered system, change the name of cluster transaction.
   For details on how to do this, refer to the relevant cluster software manual.

6. Create the cluster environment of the target transaction.

   For details on how to do this, refer to "Customization for Solaris/Linux/HP-UX version clustered system" in the "ETERNUS SF Cluster Environment Setup Guide".

   ## Note
   ....

   When you perform the cluster set up, do not change any information other than the transaction name.
   ....

7. Start the target transaction.
   For details on how to start it up, refer to the relevant cluster software manual.

8. Restore the following AdvancedCopy Manager data.

   - AdvancedCopy Manager management file

     On the node operating the transaction, restore /etc/opt/swstorage/<*logical-node-name*>/swnode.ini.

   - Backup management list
     Restore the backup management list.
     For details about how to do this, refer to "8.1.1.1.2 Restoring a backup management list and recovery control file".

   - Pre-processing/post-processing scripts
     Restore the pre-processing and post-processing scripts for backup management and replication management.

   - When using the replication management function, restore the environment by using swsrprecoverres (Resource adjustment command) with the -r option.

   For a "package-backup-of-the-management-information execution server" of AdvancedCopy Manager, restore the following data:

   - Policy file for package backup of the management information
     For details about the policy file path, refer to "8.1.2.1 Preliminaries".

9. For a "package-backup-of-the-management-information execution server" of AdvancedCopy Manager, execute stgmgrinfoset (Management information package backup server registration/deletion command) to register as its server.

8.5.7.8 Changing a shared disk for shared data

Change a shared disk that is used for Managed Server transaction by performing the following steps:

1. Check whether the target Managed Server transaction is operating.
   If the transaction is not operating, start the Managed Server transaction.
   For details about how to start it up, refer to the relevant cluster software manual.

2. Save the following AdvancedCopy Manager data.

   - AdvancedCopy Manager management file

     On the node operating the transaction, save /etc/opt/swstorage/<*logical-node-name*>/swnode.ini.

   - Backup management list
     Make a backup of the backup management list.
     For details about how to do this, refer to "8.1.1.1.1 Backing up a backup management list and recovery control file".

   - Pre-processing/post-processing scripts
     Save the pre-processing and post-processing scripts for backup management and replication management.

   For a "package-backup-of-the-management-information execution server" of AdvancedCopy Manager, save the following data:

   - Policy file for package backup of the management information
     For details about the policy file path, refer to "8.1.2.1 Preliminaries".

3. Delete /etc/opt/swstorage/<*logical-node-name*>/swnode.ini on the node operating the transaction.

4. Delete the cluster environment of the target transaction.

   For details about how to do this, refer to "Deletion on Solaris/Linux/HP-UX version clustered systems" in the "ETERNUS SF Cluster Environment Setup Guide".

Note

· ·

Do not perform the following step of "Deletion on Solaris/Linux/HP-UX version clustered systems" in the "ETERNUS SF Cluster Environment Setup Guide".

   - Step 1 described in "Deletion Procedure"

· ·

5. Create the cluster environment of the target transaction.

   For details about how to do this, refer to "Customization for Solaris/Linux/HP-UX version clustered system" in the "ETERNUS SF Cluster Environment Setup Guide".

Note

· ·

   - To change the name of a mount point, specify the new mount point name when you perform the cluster set up.

   - To change a physical disk, specify the new physical disk when you mount it.

   - Change the configuration of a shared disk in the clustered system.

· ·

6. Start the target transaction.
   For details about how to start it up, refer to the relevant cluster software manual.

7. Restore the following AdvancedCopy Manager data.

   - AdvancedCopy Manager management file

     On the node operating the transaction, restore /etc/opt/swstorage/<*logical-node-name*>/swnode.ini.

   - Backup management list
     Restore the backup management list.
     For details about how to do this, refer to "8.1.1.1.2 Restoring a backup management list and recovery control file".

   - Pre-processing/post-processing scripts
     Restore the pre-processing and post-processing scripts for backup management and replication management.

   - When using the replication management function, restore the environment by using swsrprecoverres (Resource adjustment command) with the -r option.

   For a "package-backup-of-the-management-information execution server" of AdvancedCopy Manager, restore the following data:

- Policy file for package backup of the management information
    For details about the policy file path, refer to "8.1.2.1 Preliminaries".

8. For a "package-backup-of-the-management-information execution server" of AdvancedCopy Manager, execute stgmgrinfoset (Management information package backup server registration/deletion command) to register as its server.

8.5.8 Operation on Model upgrade for ETERNUS Disk storage system

Perform the procedure of "Operation on Model upgrade for ETERNUS Disk storage system" in the "ETERNUS SF Web Console Guide."

8.6 Action when Managed Server information cannot be deleted

If it becomes necessary to delete Managed Server information under the circumstances specified below, this section provides a way to delete the Managed Server information forcibly from the Management Server.

- The operating system on the Managed Server cannot be started because of a hardware failure or another similar reason. From that point, the server will not be used for operations.

- The Managed Server has been physically removed without performing the Managed Server deletion procedure ("3.6.4 Deleting the Managed Server").

- The Managed Server is managed by AdvancedCopy Manager and has been physically removed without being deleted in AdvancedCopy Manager.

For cases other than these, refer to "3.6.4 Deleting the Managed Server" to delete the Managed Server.

📛 Note

・・

- When this procedure is performed, the specified Managed Server is forcibly deleted regardless of the backup or replication status. Therefore, this procedure should be used only for Managed Server that are not planned to be used afterwards.

- Please do not execute a backup or replication while this procedure is performed.

- The components and management information of AdvancedCopy Manager on the forcibly deleted Managed Server are not deleted. When the forcibly deleted server is to be used again, uninstall the AdvancedCopy Manager from the deleted server before it is reconnected to the network. If its server is reconnected to the networks without uninstalling first AdvancedCopy Manager, the AdvancedCopy Manager functions of other servers on the network may not work correctly.

・・

The procedure to delete forcibly a Managed Server is described below:
Confirm that the Managed Server will not be used afterwards or that it will be removed before starting the procedure.

1. Check the status of AdvancedCopy Manager on the Managed Server to be deleted forcibly.
    If AdvancedCopy Manager is operable, use the normal deletion procedure for deleting the Managed Server instead of the forcible deletion procedure. For details of the normal deletion procedure, refer to "3.6.4 Deleting the Managed Server".

    How to check the status of Managed Server

       On the Management Server, execute stgxfwcmsetdev (Device information collection/reflection command) with -n option. Specify the target Managed Server with an operand of -n option.
       After executing the command, check the status of Managed Server as follows.
       Go to step 2, only when the Managed Server cannot be started.

         - When the Managed Server cannot be started:

            The following message is displayed:
            "swstf5405 The acquisition of device information caused a failure.
            Reason=swstf2111 Connection to Storage Server failed."

         - When the Managed Server can be started:

            The following message is displayed:
            "swstf5408 There is no change in information currently being managed."

Alternatively the command may be waiting for input of the device information confirmation. If the command is waiting for input, enter 'q' or 'n' to prevent updating the device information and end the stgxfwcmsetdev command execution.

2. Stop the Advanced Copy sessions.
   Even if the Managed Server information has been deleted, the Advanced Copy sessions for the connected disks (Logical Units) may not be stopped.
   For the unnecessary Advanced Copy sessions, stop its sessions by using ETERNUS Web GUI.

3. Back up the database.
   For the procedures, refer to "8.2.1.1 Database backup".

4. Delete the Managed Server information forcibly.
   Execute stgforcedelsrv (Managed Server forcibly deletion command).

Note

If the stgforcedevsrv command terminated abnormally, please do not proceed from step 5 onwards.
If the steps after step 5 are executed an incorrect database is restored and normal system operation may not be performed adequately afterwards.

5. Restore the database.
   Restore the database from the directory that the database is stored. Its directory must be a directory of "after-dir" specified with -a option in stgforcedelsrv (Managed Server forcibly deletion command) at step 4. For the procedures, refer to "8.2.1.2 Database restoration".

6. Perform the following operation according to the status of the forcibly deleted Managed Server.

   - If a replication operation has been done:

     - If the forcibly deleted Managed Server was part of a server-to-server replication:

       On the Management Server, execute swsrprecoverres (Resource adjustment command) with the -h and -r options. For the operand of -h option, specify the name of the server used for the replication with the forcibly deleted Managed Server. If several servers where used for replication, the swsrprecoverres command has to be executed for each of those servers.

     - If the forcibly deleted Managed Server was not part of a server-to-server replication:

       No operation is required.

   - If a backup operation has been done:

     No operation is required.

8.7 Maintaining a database of ETERNUS SF Manager

This section explains the maintenance for database of ETERNUS SF Manager.

8.7.1 Saving a Storage Cruiser's manager environment

To save a Storage Cruiser's manager environment, refer to "Backup of Manager" in "ETERNUS SF Storage Cruiser Operation Guide".

8.7.2 Maintaining a common control repository

This section explains the maintenance of common control repository.

8.7.2.1 Backup of common control repository

Save the common control repository.

Information

About the notation in procedure

Directory name	Explanation
$BAK_DIR	Backup destination directory.

1. Stop the ETERNUS SF Manager Web service.
   Execute the following shell script to stop the service.

   ```
   /opt/FJSVesfcm/bin/stop-webservice.sh
   ```

2. Create a backup data.

   1. Create the following backup directories:

      ```
      mkdir -p $BAK_DIR/Common
      ```

   2. Execute the following command to create a backup data with OS Administrator privilege.

      If you did not change the port number, specify 15432 for <*port number*>. If you did change the port number, specify the changed port number for <*port number*>.
      The file name of the created backup data is <*file_name*>.

      ```
      /opt/FJSVesfcm/postgres/bin/pg_dump -U esfpostgres -C --attribute-inserts -p <port number>
      -f $BAK_DIR/Common/<file_name> esfdb
      ```

   3. The prompt is displayed, and waiting for input. Input the following strings.

      ```
      master01!
      ```

3. Start the ETERNUS SF Manager Web service.

   Execute the following shell script to start the service.

   ```
   /opt/FJSVesfcm/bin/start-webservice.sh
   ```

8.7.2.2 Restore of common control repository

Restore the common control repository.

Information

About the notation in procedure

Directory name	Explanation
$BAK_DIR	Backup destination directory.

1. Stop the ETERNUS SF Manager Web service.

   Execute the following shell script to stop the service.

   ```
   /opt/FJSVesfcm/bin/stop-webservice.sh
   ```

2. Restore a backup data. The following procedure must be implemented by a user with administrator privileges for the operating system.

   1. Execute the following command.

      If you did not change the port number, specify 15432 for <*port number*>. If you did change the port number, specify the changed port number for <*port number*>.

      ```
      /opt/FJSVesfcm/postgres/bin/psql -U esfpostgres -p <port number> postgres
      ```

2. The prompt is displayed, and waiting for input. Input the following strings in order.

```
master01!
drop database esfdb;
\q
```

3. Execute the following command to restore the backup data.
   If you did not change the port number, specify 15432 for <*port number*>. If you did change the port number, specify the changed port number for <*port number*>.
   The <*file_name*> is the file name specified when the backup data was created.

```
/opt/FJSVesfcm/postgres/bin/psql -U esfpostgres -p <port number> -f $BAK_DIR/Common/
<file_name> postgres
```

4. The prompt is displayed, and waiting for input. Input the following strings.

```
master01!
```

3. Start the ETERNUS SF Manager Web service.

   Execute the following shell script to start the service.

```
/opt/FJSVesfcm/bin/start-webservice.sh
```

Note

After restoring the shared control repository, it takes approximately 2 minutes before the Web Console becomes available.

8.7.3 Maintaining a license management database

This section explains the maintenance of license management database.

8.7.3.1 Backup of license management database

Save the license management database.

Information

About the notation in procedure

Directory name	Explanation
$BAK_DIR	Backup destination directory.

1. Create the following backup directories:

```
mkdir -p $BAK_DIR/LM
```

2. Execute the following command to create a backup data with OS Administrator privilege.
   The file name of the created backup data is <*file_name*>.

```
/opt/FJSVesflm/bin/esflm backup $BAK_DIR/LM/<file_name>
```

8.7.3.2 Restore of license management database

Restore the license management database.

About the notation in procedure

Directory name	Explanation
$BAK_DIR	Backup destination directory.

Execute the following command with administrator privileges for the operating system.
The <*file_name*> is the file name specified when the backup data was created.

```
/opt/FJSVesflm/bin/esflm restore $BAK_DIR/LM/<file_name>
```

8.7.4 Saving an AdvancedCopy Manager CCM environment settings files

If AdvancedCopy Manager CCM is used by the server where ETERNUS SF Manager is installed, save an AdvancedCopy Manager CCM environment settings files.

For details, refer to "How to back up environment settings files" in the "ETERNUS SF Operation Guide for Copy Control Module".

If AdvancedCopy Manager CCM environment settings files are already saved, it is unnecessary to perform this task.

Chapter 9 Commands

This chapter describes using commands.

Note

Set the following paths in the PATH environment variable to execute the commands.

- /usr/sbin

- /usr/bin

9.1 Command list

The commands are classified as follows:

Backup management commands

Environment definition commands

Table 9.1 List of environment definition commands for backup management

Function	Command name	Explanation
Managed Server configuration information setting command	swstsvrset	Sets the configuration information of a Managed Server.
Managed Server configuration information display command	swstsvrdisp	Displays the configuration information of a Managed Server.
Device information setting command	swstdevinfoset	Sets the usage type of a device.
Device use status display command	swstdevdisp	Displays the information of a device.
Backup policy setting command	swstbkpolset	Sets a backup policy.
Backup policy display command	swstbkpoldisp	Displays a backup policy that has been set.
Backup policy deletion command	swstbkpoldel	Deletes a backup policy that has been set.

Operation commands

Table 9.2 List of operation commands for backup management

Function	Command name	Explanation
Backup execution command	swstbackup	Performs the synchronous high-speed backup or snapshot fast backup.
Backup execution status display command	swstbackstat	Displays the backup progress status.
Restore execution command	swstrestore	Performs restoration.
Restore cancellation command	swstcancelrest	Cancels the execution of restoration.

Function	Command name	Explanation
Restore execution status display command	swstreststat	Displays the execution status of restoration.
History information display command	swsthistdisp	Displays the backup history.
History information deletion command	swsthistdel	Deletes the backup history.
Backup synchronous processing start command	swststartsync	Starts backup synchronous processing.
Backup synchronous processing cancel command	swstcancelsync	Aborts backup synchronous processing.
Backup synchronous processing progress display command	swstsyncstat	Displays the progress of the backup synchronous processing.
Execution status display command	swstexecstat	Displays the execution status of a command executed on a transaction volume.
Tracking cancel command	swstcanceltrk	Stops tracking processing.
Tracking execution status display command	swsttrkstat	Displays the status of tracking processing.

Maintenance commands

Table 9.3 List of maintenance commands for backup management

Function	Command name	Explanation
Resource adjustment command	swstsrsemtch	Recovers the consistency of information in a backup management file.
Resource backup command	swstresback	Backup backup management files
Resource restore command	swstresrst	Restores backup management files

Configuration management commands

Table 9.4 List of configuration management commands

Function	Command name	Explanation
Management Server information change command	stgcmmodnode	Changes the server information of the Management Server
Repository access user change command	stguserset	Use this command to change the username or password used to access the repository.
Server information addition command	stgxfwcmaddsrv	This command adds information of a Managed Server newly managed by AdvancedCopy Manager.
Server information change command	stgxfwcmmodsrv	This command changes the information of a Managed Server managed by AdvancedCopy Manager.
Server information deletion command	stgxfwcmdelsrv	This command deletes the information of a Managed Server currently managed by AdvancedCopy Manager.
Device information fetch/ reflect command	stgxfwcmsetdev	This command collects the information of the device most recently detected by the server and displays the differences compared with the information managed by AdvancedCopy Manager.
Device information deletion command	stgxfwcmdeldev	This command deletes the information of a device currently managed by AdvancedCopy Manager.

Function	Command name	Explanation
Server information display command	stgxfwcmdispsrv	This command displays the information of a server managed by AdvancedCopy Manager.
Device information display command	stgxfwcmdispdev	This command displays the information of a device managed by AdvancedCopy Manager.
Partition information display command	stgxfwcmdisppat	This command displays the information of the partitions that comprise a device managed by AdvancedCopy Manager.
Environment information display command	stgenvdisp	Displays the environment values set for AdvancedCopy Manager
Management information batch backup server registration/deletion command	stgmgrinfoset	Registers servers for execution of batch backup of management information
Management information batch backup command	stgmgrinfobkup	Backs up Managed Server and Management Server management information
Management information batch backup status display command	stgmgrinfodisp	Displays the execution status of the management information batch backup command
SDX object management unit configuration command	stgxfwcmsetsdx	Sets the management unit for SDX objects handled by AdvancedCopy Manager
Information collection mode setup command	stgxfwcmsetmode	Sets the information collection mode. This command is executed from the Managed Server.
Repository update area busy rate command	stgdbloginf	Checks the busy rate at the area required when a repository is updated
Database save command	stgdbdmp	Saves the database to the specified directory
Database recovery command	stgdbrcv	Recovers a saved database
Database backup command	stgdbbackup	Backs up the data stored in the database to the specified directory as a batch
Database restore command	stgdbrestore	Restores the data saved as database backup to the database as a batch
Database integrated setup command	stgdbset	Starts database integrated setup and creates database environments and databases
Database integrated unsetup command	stgdbunset	Starts database integrated unsetup, deletes databases and database environments, and cancels repository settings
Repository data migration command	stgrepocnv	Recovers the targeted repository data in a new version repository
Managed Server forcibly deletion command	stgforcedelsrv	Forcibly deletes the information of a Managed Server

Replication management commands

Environment definition commands

Table 9.5 List of environment definition commands for replication management

Function	Command name	Explanation
Replication volume information setting command	swsrpsetvol	Sets the information for a replication volume.
Replication volume information display command	swsrpvolinfo	Displays the information for a replication volume.

Function	Command name	Explanation
Replication volume information deletion command	swsrpdelvol	Deletes replication volume information that has been set.
REC buffer information display command	swsrprecbuffstat	Displays information for the REC buffer.
REC buffer settings change command	swsrprecbuffset	Changes the settings for the REC buffer.

Operation commands

Table 9.6 List of operation commands for replication management

Function	Command name	Explanation
Replication start command	swsrpstartsync	Starts replication when the synchronous type replication function is performed.
Replica creation command	swsrpmake	Creates replication when the snapshot type and the synchronous type replication functions are performed.
Use status display command	swsrpstat	Displays the status of replication operation
Replication cancel command	swsrpcancel	Cancels the replication processing when the snapshot type and the synchronous type replication functions are performed.
Synchronous processing mode change command	swsrpchsync	Changes the operation mode of inter-box synchronization.
Synchronous processing reverse command	swsrprevsync	Reverses the copying direction of inter-box synchronization which is in suspended status.

Maintenance commands

Table 9.7 List of maintenance commands for replication management

Function	Command name	Explanation
Resource adjustment command	swsrprecoverres	Recovers the consistency of the information in the replication management file

Daemon start and stop commands

Table 9.8 Daemon start and stop command list

Function	Command name	Explanation
Communication daemon start and stop	stgfwcom	Starts and stops communication daemons
RMI daemon start and stop	stgrmicom	Starts and stops RMI daemons
RDB daemon start and stop	stgdbcom	Starts and stops RDB daemons
AdvancedCopy Manager daemon start and stop	startacm stopacm	Starts and stops AdvancedCopy Manager daemons

SnapOPC/SnapOPC+ support commands

Table 9.9 List of SnapOPC/SnapOPC+ support commands

Function	Command name	Explanation
Update size measurement command	swstestupdate	Uses the pseudo-session setting function to measure the physical size of a Snap Data Volume

Function	Command name	Explanation
Snap Data Volume operation/ reference command	swstsdv	By specifying a subcommand, Snap Data Volume initialization/status display and Snap Data Pool status display are possible.

Troubleshooting data collection commands

Table 9.10 List of Troubleshooting data collection commands

Function	Command name	Explanation
AdvancedCopy Manager's manager troubleshooting data collection command	esfsnap	This command collects troubleshooting data for AdvancedCopy Manager's manager.
AdvancedCopy Manager's agent troubleshooting data collection command	acmcapture	This command collects troubleshooting data for AdvancedCopy Manager's agent.

9.2 Backup Management Commands

This section describes operations using Backup Management Commands

🔔 **Note**

. . . .

- Only a root user can execute all the commands of backup management. When command execution is carried out by other users, a message "swst0601 This Command cannot be run. This is because the user is not a root user." is displayed and then, an error end.

- Backup operations in Logical Units (disk units) are not supported. An error occurs if swstdevinfoset (Device information setting command) is used to specify a Logical Unit (disk).

. . . .

🔔 **Note**

. . . .

Operating environment for command execution

The table below lists the operating environment requirements for different commands.

Table 9.11 Operating environment for executing backup management commands

Command	Operating environment of AdvancedCopy Manager

swstbackstat	The daemons and services of the Management Server and target Managed Server must be operating in order to execute this command, if the Management Server has either of the following parameter designations:
	
	  - -h option specified
	
	  - No device name specified

swstbackup	To execute this command on a Management Server with the -h option specified, the daemons and services of the Management Server and target Managed Server must be operating. When the Managed Server runs on Windows and is in EC operation, the daemons and services of the Managed Server must be operating.
swstbkpoldel	To execute this command on a Managed Server, the daemons and services of the Management Server must be operating. To execute this command on a Management Server with the -h option specified, the daemons and services of the Management Server and target Managed Server must be operating.

Command	Operating environment of AdvancedCopy Manager
swstbkpoldisp	The daemons and services of the Management Server and target Managed Server must be operating in order to execute this command, if the Management Server has either of the following parameter designations: - -h option specified - No device name specified
swstbkpolset	To execute this command on a Managed Server, the daemons and services of the Management Server must be operating. To execute this command on a Management Server with the -h option specified, the daemons and services of the Management Server and target Managed Server must be operating.
swstcancelrest	To execute this command on a Management Server with the -h option specified, the daemons and services of the Management Server and target Managed Server must be operating.
swstcancelsync	To execute this command on a Management Server with the -h option specified, the daemons and services of the Management Server and target Managed Server must be operating. When the Managed Server runs on Windows and is in EC operation, the daemons and services of the Managed Server must be operating.
swstcanceltrk	When this command is executed with the -h option specified on the Management Server, the daemons/services of the Management Server and target Managed Server must be active.
swstdevdisp	The daemons and services of the Management Server and target Managed Server must be operating in order to execute this command, if the Management Server has either of the following parameter designations: - -h option specified - No device name specified To execute this command with the -o option specified, the daemons and services of the Management Server must be operating.
swstdevinfoset	To execute this command on a Managed Server, the daemons and services of the Management Server must be operating. To execute this command on a Management Server with the -h option specified, the daemons and services of the Management Server and target Managed Server must be operating.
swstexecstat	The daemons and services of the Management Server and target Managed Server must be operating in order to execute this command on a Management Server with either of the following parameter designations: - -h option specified - No device name specified
swsthistdel	To execute this command on a Management Server with the -h option specified, the daemons and services of the Management Server and target Managed Server must be operating.
swsthistdisp	The daemons and services of the Management Server and target Managed Server must be operating in order to execute this command on a Management Server with either of the following parameter designations: - -h option specified - No device name specified
swstresback	-
swstresrst	To execute this command on a Management Server with the -h option specified, the daemons and services of the Management Server and target Managed Server must be operating. When the Managed Server runs on Windows, the daemons and services of the Managed Server must be operating.

Command	Operating environment of AdvancedCopy Manager
	To execute this command with the -x option specified, the daemons and services of the Management Server must be operating.
swstrestore	To execute this command on a Management Server with the -h option specified, the daemons and services of the Management Server and target Managed Server must be operating.
swstreststat	The daemons and services of the Management Server and target Managed Server must be operating in order to execute this command on a Management Server with either of the following parameter designations: - -h option specified - No device name specified
swstsrsemtch	The daemons and services of the Management Server must be operating. To execute this command on a Management Server with the -h option specified, the daemons and services of the target Managed Server must be operating. When the Managed Server runs on Windows, the daemons and services of the Managed Server must be operating.
swststartsync	To execute this command on a Management Server with the -h option specified, the daemons and services of the Management Server and target Managed Server must be operating. When the Managed Server runs on Windows, the daemons and services of the Managed Server must be operating.
swstsvrdisp	To execute this command on a Management Server with the -h option specified, the daemons and services of the Management Server and target Managed Server must be operating.
swstsvrset	To execute this command on a Managed Server, the daemons and services of the Management Server must be operating. To execute this command on a Management Server with the -h option specified, the daemons and services of the Management Server and target Managed Server must be operating.
swstsyncstat	The daemons and services of the Management Server and target Managed Server must be operating in order to execute this command on a Management Server with either of the following parameter designations: - -h option specified - No device name specified
swsttrkstat	When this command is executed in either of the following formats on the Management Server, the daemons/services of the Management Server and target Managed Server must be active: - -h option is specified - No device name is specified

----

Note

In cluster operation

In cluster operation, execute a command by performing the following steps:

1. If necessary, specify the logical node name to the transaction in the environment variable SWSTGNODE. Specify the environment variable as follows:

   ```
   # SWSTGNODE=logical-node-name
   # export SWSTGNODE
   ```

2. Execute the command.

The table below lists the commands available in cluster operation, notes whether environment variables must be configured, and specifies the nodes on which commands are executable. For notes and cautions on command execution, refer to the Remarks column.

Table 9.12 Commands when cluster operation is used

Command	Environment variable	Command execution node	Remarks
swstsvrset	Required	Transaction operating node	-
swstsvrdisp	Required	Transaction operating node	-
swstdevinfoset	Required	Transaction operating node	-
swstdevdisp	Required	Transaction operating node	-
swstbkpolset	Required	Transaction operating node	-
swstbkpoldisp	Required	Transaction operating node	-
swstbkpoldel	Required	Transaction operating node	-
swstbackup	Required	Transaction operating node	-
swstbackstat	Required	Transaction operating node	-
swstrestore	Required	Transaction operating node	-
swstcancelrest	Required	Transaction operating node	-
swstreststat	Required	Transaction operating node	-
swsthistdisp	Required	Transaction operating node	-
swsthistdel	Required	Transaction operating node	-
swststartsync	Required	Transaction operating node	-
swstcancelsync	Required	Transaction operating node	-
swstsyncstat	Required	Transaction operating node	-
swstexecstat	Required	Transaction operating node	-
swstcanceltrk	Required	Transaction operating node	-
swsttrkstat	Required	Transaction operating node	-
swstsrsemtch	Required	Transaction operating node	-
swstresback	Required	Transaction operating node	-
swstresrst	Required	Transaction operation node	-

Note

In a Management Server, when a display-type command that has 'disp' or 'stat' at the end of the command name is executed, the display contents will depend on the existence of the -h option and device name. For details, refer to the following table.

Table 9.13 Display contents of the display-type commands in Management Server

-h option	Device name specified	Display contents
Does not exist	No	Displays all Managed Server information registered using the Backup management functionality. If an error occurs on any of the Managed Servers during display processing, the processing will be terminated.
	Yes	Displays only the specified Management Server device information.
Exists	No	Displays all Managed Server information specified using the -h option.

-h option	Device name specified	Display contents
	Yes	Displays only the specified device information for the Managed Server specified using the -h option.

9.2.1 Environment definition commands

This section describes the environment definition commands for backup management.

9.2.1.1 swstsvrset (Managed Server configuration information setting command)

NAME

swstsvrset - sets configuration information of a Managed Server

SYNOPSIS

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstsvrset [-c RecCtrlFile-Directory] [-w Work-Directory]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstsvrset [-h Server-Name] [-c RecCtrlFile-Directory] [-w Work-Directory]
```

DESCRIPTION

This command sets configuration information of a Managed Server. Execution of this command declares that the Managed Server on which the command is executed is a server that performs backup operations.

For backup and recovery of Symfoware databases, register the storage location of the recovery control file created during backup (target directory for output of the recovery control file) as well as the working directory that is used for the recovery of the databases during restore.

For modifying the target directory for output of the recovery control file or the working directory, specify the directory name to be changed and run the command again.

After modifying the target directory for output of the recovery control file, transfer the recovery control file from the previous directory to the new one. If the transfer fails due to insufficient space, the target directory for output of the recovery control file will not be changed. In that case, the recovery control file will be kept in the directory where it was before the modification.

OPTIONS

-h *Server-Name*

Specifies the name of a Managed Server to be subject to the backup operation.
This option can be specified only on a Management Server.
If this option is omitted, the Managed Server on which the command is entered will be subject to the processing.

-c *RecCtrlFile-Directory*

Specifies the name of the target directory for output of the recovery control file. As the output target, specify an empty directory.
If this option is omitted at the time of first registration, the following directories will be used.

Non-cluster operation: /etc/opt/FJSVswsts/SymfoWARE/Recovery
Cluster operation: /etc/opt/FJSVswsts/<*logical-node-name*>/SymfoWARE/Recovery

If this option is omitted during an update, the previously set values will be carried over.
Do not use this option when Symfoware is not installed.

-w *Work-Directory*

Specifies the name of the working directory.
If this option is omitted at the time of first registration, the following directories will be used.

Non-cluster operation: /var/opt/FJSVswsts/SymfoWARE
Cluster operation: /var/opt/FJSVswsts/<*logical-node-name*>/SymfoWARE

If this option is omitted during an update, the previously set values will be carried over.
Do not use this option when Symfoware is not installed.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Set the configuration information of a Managed Server:

```
# /opt/FJSVswsts/bin/swstsvrset
swstsvrset completed
#
```

- Set the configuration information for the Managed Server on which Symfoware is installed:

```
# /opt/FJSVswsts/bin/swstsvrset -c /etc/opt/FJSVswsts/SymfoWARE/Recovery -w /var/opt/FJSVswsts/
SymfoWARE
swstsvrset completed
#
```

- Set the configuration information of a Managed Server named job1 from a Management Server:

```
# /opt/FJSVswsts/bin/swstsvrset -h job1
swstsvrset completed
#
```

NOTES

- In cluster operation, as the target directory for output of the recovery control file specified by the -c option, you must specify a directory that is located on a shared disk.

- The configuration information for the Managed Server cannot be set if:

    - The -w or the -c option was specified although Symfoware is not installed.

    - The directory specified in the -w option does not exist.

    - The directory specified in the -c option does not exist.

    - The directory name specified in the -w or -c option contains double-byte characters (as used in Asian languages).

- This command can only be executed if none of the following commands are running:

    - swstsvrset (Managed Server configuration information setting command)

    - swstbkpolset (Backup policy setting command)

    - swstbkpoldel (Backup policy deletion command)

    - swstbackup (Backup execution command)

- swstrestore (Restore execution command)

- swstcancelrest (Restore cancellation command)

- swststartsync (Backup synchronous processing start command)

- swstcancelsync (Backup synchronous processing cancel command)

- swstdevinfoset (Device information setting command)

- swsthistdel (History information deletion command)

- swstcanceltrk (Tracking cancel command)

- swstresback (Resource backup command)

- swstresrst (Resource restore command)

- swstsrsemtch (Resource match command)

9.2.1.2 swstsvrdisp (Managed Server configuration information display command)

NAME

swstsvrdisp - displays configuration information of a Managed Server

SYNOPSIS

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstsvrdisp
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstsvrdisp [-h Server-Name]
```

DESCRIPTION

This command displays configuration information of a Managed Server that has been set using swstsvrset (Managed Server configuration information setting command).

The following information will be displayed:

Title	Description
Storage-Server	Displays a Managed Server name.
RecCtrlFile-Dir	Displays the name of the target directory for output of the recovery control file.
Work-Dir	Displays the name of the working directory.

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on a Management Server.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Display configuration information of a Managed Server named Work1:

```
# /opt/FJSVswsts/bin/swstsvrdisp
Storage-Server  = Work1
#
```

- Display the configuration information for the Managed Server named Work2 on which Symfoware is installed:

```
# /opt/FJSVswsts/bin/swstsvrdisp
Storage-Server  = Work2
RecCtrlFile-Dir = /etc/opt/FJSVswsts/SymfoWARE/Recovery/
Work-Dir        = /var/opt/FJSVswsts/SymfoWARE/
#
```

NOTES

- The directory name for outputting recovery control file and the working directory name are displayed only when Symfoware is already installed.

9.2.1.3 swstdevinfoset (Device information setting command)

NAME

swstdevinfoset - sets the device information

SYNOPSIS

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstdevinfoset -t | -b | -o Device-Name
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstdevinfoset [-h Server-Name] -t | -b | -o Device-Name
```

DESCRIPTION

Set a volume that has been allocated to the ETERNUS Disk storage system, as a transaction or backup volume.

A transaction volume refers to a volume that stores transaction data to be backed up. A backup volume refers to a volume used to store a backup of this transaction data.

If you want to perform the backup operation on a regular transaction volume, refer to "3.2.2 Selecting a backup source" for information on a transaction volume, and "3.2.4 Preparing a backup volume" for information on a backup volume.

If you want to perform the backup operation for Symfoware, refer to "4.1.2.3 Selecting a backup source" for information on a transaction volume, and "4.1.2.5 Preparing a backup volume" for information on a backup volume.

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on a Management Server.

-t

Register the specified device as a transaction volume.
Cannot be specified at the same time as the -b and -o options.

-b

Register the specified device as a backup volume.
Cannot be specified at the same time as the -t and -o options.

-o

Deregister the specified device if it has been registered as a transaction or backup volume.
Cannot be specified at the same time as the -t and -b options.

OPERANDS

Device-Name

Specify an AdvancedCopy Manager device name.
Refer to "1.5 Managing a Device on AdvancedCopy Manager" for the devices to be specified.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Register a device as a transaction volume:

```
# /opt/FJSVswsts/bin/swstdevinfoset -t /dev/sdb6
swstdevinfoset completed
#
```

- Register a device as a backup volume:

```
# /opt/FJSVswsts/bin/swstdevinfoset -b /dev/dsk/sdb6
swstdevinfoset completed
#
```

- Deregister a device that has been registered as a transaction volume:

```
# /opt/FJSVswsts/bin/swstdevinfoset -o /dev/dsk/sdb6
swstdevinfoset completed
#
```

NOTES

- A device cannot be registered as a transaction volume if:

    - Another device with the same box information (box identifier, OLU number, EXTENT starting position, and EXTENT size) has already been registered as a transaction volume.

    - The space of the device to be registered overlaps the space of an already registered working or backup volume.

    - The device to be set up contains a Symfoware database space in its configuration.

    - The device is a Logical Unit (disk).

    - The license for the ETERNUS Disk storage system that contains the specified transaction volume is one of the following:

        - The license is not registered in the license management function.

- The license is expired.

- A device cannot be registered as a backup volume if:

    - Another device with the same box information (box identifier, OLU number, EXTENT starting position, and EXTENT size) has already been registered as a backup volume.

    - The space of the device to be registered overlaps the space of an already registered work or backup volume.

    - The database space of Symfoware is built on the device to be configured.

    - The device to be configured is already mounted.

    - The device to be set up is registered as a replication destination volume in replication management.

    - The device to be set up is registered as a replication source volume of the duplicate volume information in which bidirectional copy is possible in replication management.

    - The device is a Logical Unit (disk).

    - The device is an SDX object.

- You must delete related backup history information before:

    - Changing a backup volume to a transaction volume

    - Cancelling the setting of a backup volume

- You must delete a backup policy specified on a transaction volume before:

    - Changing a transaction volume to a backup volume

    - Cancelling the setting of a transaction volume

- The following changes are not possible unless the related synchronous processing (in an EC session) is cancelled:

    - Changing a setting in a suspended transaction volume

    - Changing a setting in a suspended backup volume

- The related tracking processing (in an OPC session) must be cancelled in advance to perform the following changes:

    - Changing the settings of a transaction volume in tracking processing

    - Changing the settings of a backup volume in tracking processing

- For details on how to use SDX objects of GDS, refer to "Configuration and conditions of unsupported SDX objects" and "10.1.5.1 Notes on SDX object operations (backup management)".

- This command can be executed only while none of the following commands are running:

    - swstsvrset (Managed Server configuration information setting command)

    - swstdevinfoset (Device information setting command)

    - swstbkpolset (Backup policy setting command)

    - swstbkpoldel (Backup policy deletion command)

    - swstbackup (Backup execution command)

    - swstrestore (Restore execution command)

    - swstcancelrest (Restore cancellation command)

    - swststartsync (Backup synchronous processing start command)

    - swstcancelsync (Backup synchronous processing cancel command)

    - swsthistdel (History information deletion command)

    - swstcanceltrk (Tracking cancel command)

    - swstsrsemtch (Resource match command)

9.2.1.4 swstdevdisp (Device usage status display command)

NAME

swstdevdisp - displays the device information

SYNOPSIS

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstdevdisp [-t] [-b [-u]] [-o] [Device-Name]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstdevdisp [-h Server-Name] [-t] [-b [-u]] [-o] [Device-Name]
```


DESCRIPTION

This command displays information about a transaction or backup volume or any other definable device, which has been set using swstdevinfoset (Device information setting command).

The following information will be displayed:

Title	Description
Server	Displays a Managed Server name.
Device	Displays the AdvancedCopy Manager device name.
Size	Displays the size of a partition allocated to a device.
Device-Mode	Displays a character string by which a device can be identified: "Transaction" : Represents a transaction volume. "Backup" : Represents a backup volume. "Other" : Represents a device that can be registered as a working or backup volume. A backup volume is indicated as "Backup (used)" if it is being used or "Backup (free)" if it is not used. During the backup process with -suspend or -T option specified, the backup volume remains in use after history information is deleted. When the -u option is specified, "Backup (used-S)" is displayed for a backup volume in the suspend status for which history information was deleted, or "Backup (used-T)" is displayed for a backup volume in the tracking status for which history information was deleted.
Mount-Point (Method)	Displays the mount point of a device. The file system type of the mount point is displayed inside the parentheses. If the device is a transaction volume used for Symfoware, the names of Symfoware database spaces and log groups as well as the RDB system names are displayed, followed by "SymfoWARE" in parentheses.
Backup-Engine	Displays the operation type: "AdvancedCopy" : Backup operation using AdvancedCopy Manager

- For transaction volumes used for Symfoware, the display in the "Mount-Point (Method)" column has the following format:

```
DB-Name  DBS-Name  /  LogGroup  /  RDB-Name  (SymfoWARE)
(1)      (2)          (3)          (4)        (5)


(1) Displays the database name.
(2) Displays the database space name.
```

(3) Displays the log group name.
(4) Displays the RDB system name. If there is no RDB system name, the display shows "-".
(5) Displays that the volume is a Symfoware volume.

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on a Management Server.

-t

Displays information on a device that has been set as a transaction volume.
Displays device information on all the transaction volumes if the device name in an operand is omitted.

-b

Displays a device that has been set as a backup volume.
Displays device information on all the backup volumes if the device name in an operand is omitted.

-u

Changes the display format of the **Device-Mode** field.

-o

Displays all the devices other than those set as a transaction or backup volume.
Displays all the devices other than those set as a transaction or backup volume if the device name in an operand is omitted.

Information on a device that has been set as a transaction volume will be displayed if all of the -t, -b and -o options are omitted.

OPERANDS

Device-Name

Specify an AdvancedCopy Manager device name.
Refer to "1.5 Managing a Device on AdvancedCopy Manager" for the devices to be specified.

Displays device information on all the transaction volumes if this operand and the associated options are omitted.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Display the usage statuses of the registered devices:

```
# /opt/FJSVswsts/bin/swstdevdisp -tb
Server Device      Size        Device-Mode    Mount-Point (Method) Backup-Engine
Work2  /dev/sdb6   1.0 Gbyte Transaction    /mnt/tran1 (ext2)    AdvancedCopy
Work2  /dev/sdd1   2.0 Gbyte Transaction    DB1.DBS6/LOG01/RDB01 (SymfoWARE) AdvancedCopy
Work2  /dev/ddd2   2.0 Gbyte Transaction    DB1.DBS2/LOG02/RDB01 (SymfoWARE) AdvancedCopy
Work2  /dev/sdb5   1.0 Gbyte Transaction    ---- (----)          AdvancedCopy
Work2  /dev/sdc5   1.0 Gbyte Backup (used) ---- (----)          AdvancedCopy
Work2  /dev/sdc6   1.0 Gbyte Backup (free) ---- (----)          AdvancedCopy
Work2  /dev/sdc7   2.0 Gbyte Backup (free) ---- (----)          AdvancedCopy
Work2  /dev/sdc8   2.0 Gbyte Backup (free) ---- (----)          AdvancedCopy
:
#
```

9.2.1.5 swstbkpolset (Backup policy setting command)

NAME

swstbkpolset - sets a backup policy

SYNOPSIS

Setting a backup policy by specifying the transaction volume

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstbkpolset [-i Interval-Days] [-s Save-Number] [Device-Name]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstbkpolset [-h Server-Name] [-i Interval-Days] [-s Save-Number] [Device-
Name]
```

Setting a backup policy by specifying a log group

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstbkpolset -n [-i Interval-Days] [-s Save-Number] Log-Group-Name[/RDB-Name]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstbkpolset -n [-h Server-Name][-i Interval-Days][-s Save-Number] Log-Group-
Name[/RDB-Name]
```

DESCRIPTION

This command sets a backup policy (the number of preservation generations and the number of interval days).
The number of preservation generations is the number of generations of backup data that should be kept at one time. The number of interval days means the number of days from the execution of one backup until the next backup should be performed.

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on a Management Server.

-i *Interval-Days*

Specifies the number of interval days. You can specify a value between 1 and 366. If this option is omitted during initial registration, 30 will be set as the default. If this option is omitted during update, the existing value will be inherited.

-s *Save-Number*

Specifies the number of preservation generations. You can specify a value between 1 and 31. If this option is omitted during initial registration, 2 will be set as the default. If this option is omitted during update, the existing value will be inherited.

-n

Specifies that the setting is being made for a log group.
In the operand, specify the log group name and the RDB system name. If there is no RDB system name, you can omit it.

OPERANDS

Device-Name

>Specify an AdvancedCopy Manager device name.
>Refer to "1.5 Managing a Device on AdvancedCopy Manager" for the devices to be specified.
>If the device name is omitted, the backup policy will be set for all the transaction volumes.

Log-Group-Name

>Specify the name of the log group.

RDB-Name

>Specify the RDB system name after the log group name, delimited by a slash ("/").
>If there is no RDB system name, you can omit it.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- To set a backup policy with 3 backup generations for log group (LOG01/RDB1):

```
# /opt/FJSVswsts/bin/swstbkpolset -n -s 3 LOG01/RDB1
/dev/sdd1 swstbkpolset completed
/dev/sdd2 swstbkpolset completed
    :
LOG01/RDB1 swstbkpolset completed
#
```

- To set 3 as the save number of a transaction volume as a backup policy:

```
# /opt/FJSVswsts/bin/swstbkpolset -s 3
/dev/sdg3 swstbkpolset completed
/dev/sdg4 swstbkpolset completed
#
```

NOTES

- Even if the number of interval days is specified, AdvancedCopy Manager does not automatically create a backup.

- When you specify a backup policy, there must be as many backup volumes in operation according to the specified backup policy. To view information on the number of backup volumes required to perform the backup operation of a regular transaction volume, refer to "3.2.4 Preparing a backup volume".

- You can change the number of preservation generations as long as the "history information count at the time" is equal to or less than the "number of preservation generations to be specified".

- If you want to perform the backup operation for Symfoware, refer to "4.1.2.5 Preparing a backup volume".
  If you want to perform the backup operation for GDS, refer to "3.8.2 Backup operation in units of logical volumes (method with GDS Snapshot)".

- Ordinary transaction volumes must be backed up to ordinary backup volumes.
  If the transaction volume is a GDS logical volume, the backup destination must be a backup volume of a GDS logical volume.

- If GDS SDX objects are to be used, refer to information regarding "Configuration and conditions of unsupported SDX objects".

- If you execute this command with a log group specified, the backup policy will be set for all transaction volumes in that log group.
  You cannot set a backup policy specifying individual transaction volumes that are included in a log group.

- No backup policy can be set if:

    - There are not as many backup volumes as the specified number of preservation generations.

    - An option for more than the maximum number of preservation generations is specified.

    - An option for more than the maximum interval days is specified.

    - You have specified individual transaction volumes that are included in a log group.

    - The replication management function is executing a process for the target transaction volume.

    - The specified device is an LU (disk unit).

- This command can only be executed while none of the following commands are running:

    - swstsvrset (Managed Server configuration information setting command)

    - swstdevinfoset (Device information setting command)

    - swstsrsemtch (Resource match command)

    - swstresback (Resource backup command)

    - swstresrst (Resource restore command)

9.2.1.6 swstbkpoldisp (Backup policy display command)

NAME

swstbkpoldisp - displays a backup policy

SYNOPSIS

Displaying a backup policy by specifying the transaction volume

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstbkpoldisp [Device-Name]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstbkpoldisp [-h Server-Name] [Device-Name]
```

Displaying a backup policy by specifying a log group

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstbkpoldisp -n Log-Group-Name[/RDB-Name]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstbkpoldisp -n [-h Server-Name] Log-Group-Name[/RDB-Name]
```


DESCRIPTION

This command displays a backup policy that has been set for a transaction volume.
The following information will be displayed:

Title	Description
Server	Displays a Managed Server name.
Device	Displays the AdvancedCopy Manager device name.
Interval-Days	Displays a number of interval days.

Title	Description
Save-Number	Displays a number of preservation generations.
Mount-Point (Method)	Displays the mount point of a device. The file system type of the mount point is displayed inside the parentheses. If the device is a transaction volume used for Symfoware, the names of Symfoware database spaces and log groups as well as the RDB system names are displayed, followed by "SymfoWARE" in parentheses.

- For transaction volumes used for Symfoware, the display in the "Mount-Point (Method)" column has the following format:

```
DB-Name  DBS-Name  /  LogGroup  /  RDB-Name  (SymfoWARE)
(1)      (2)          (3)          (4)         (5)

(1) Displays the database name.
(2) Displays the database space name.
(3) Displays the log group name.
(4) Displays the RDB system name. If there is no RDB system name, the display shows "-".
(5) Displays that the volume is a Symfoware volume.
```

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on a Management Server.

-n

Specifies that the item being displayed is a log group.
In the operand, specify the log group name and the RDB system name.

OPERANDS

Device-Name

Specify an AdvancedCopy Manager device name.
Refer to "1.5 Managing a Device on AdvancedCopy Manager" for the devices to be specified.
If the device name is omitted, a backup policy for a transaction volume with a registered backup policy will be displayed.

Log-Group-Name

Specify the name of the log group.

RDB-Name

Specify the RDB system name after the log group name, delimited by a slash ("/").
If there is no RDB system name, you can omit it.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Display a backup policy of a transaction volume:

```
# /opt/FJSVswsts/bin/swstbkpoldisp /dev/sdg3
Server Device    Interval-Days Save-Number Mount-Point (Method)
Work2  /dev/sdg3 30            2           /mnt1       (ext2)
#
```

- Display the backup policy for log group (LOG01/RDB1):

```
# /opt/FJSVswsts/bin/swstbkpoldisp -n LOG01/RDB1
Server Device    Interval-Days  Save-Number  Mount-Point (Method)
job2   /dev/sdd1  30             2            DB1.DBS6/L0G01/RDB1 (SymfoWARE)
job2   /dev/sdd2  30             2            DB2.DBS2/L0G01/RDB1 (SymfoWARE)
  :
#
```

9.2.1.7 swstbkpoldel (Backup policy deletion command)

NAME

swstbkpoldel - deletes the backup policy

SYNOPSIS

Deleting a backup policy by specifying the transaction volume

To be executed on the Managed Server

/opt/FJSVswsts/bin/swstbkpoldel *Device-Name*

To be executed on the Management Server

/opt/FJSVswsts/bin/swstbkpoldel [-h *Server-Name*] *Device-Name*

Deleting a backup policy by specifying a log group

To be executed on the Managed Server

/opt/FJSVswsts/bin/swstbkpoldel -n *Log-Group-Name*[/*RDB-Name*]

To be executed on the Management Server

/opt/FJSVswsts/bin/swstbkpoldel -n [-h *Server-Name*] *Log-Group-Name*[/*RDB-Name*]

DESCRIPTION

An existing backup policy can be deleted using this command.
Delete the backup history information of the transaction volume prior to the execution of this command.

If you delete a backup policy with a log group specified, the backup policy will be deleted for all transaction volumes in that log group.

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on a Management Server.

-n

Specifies that deletion is going to be made for a log group.
In the operand, specify the log group name and the RDB system name. If there is no RDB system name, you can omit it.

OPERANDS

Device-Name

Specify an AdvancedCopy Manager device name.

Refer to "1.5 Managing a Device on AdvancedCopy Manager" for the devices to be specified.

Log-Group-Name

Specify the name of the log group.

RDB-Name

Specify the RDB system name after the log group name, delimited by a slash ("/").

If there is no RDB system name, you can omit it.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- To delete the backup policies of a transaction volume:

```
# /opt/FJSVswsts/bin/swstbkpoldel /dev/sdg3
/dev/sdg3 swstbkpoldel completed
#
```

- To delete the backup policy for log group (LOG01/RDB1):

```
# /opt/FJSVswsts/bin/swstbkpoldel -n LOG01/RDB1
/dev/sdd1 swstbkpoldel completed
/dev/sdd2 swstbkpoldel completed
    :
LOG01/RDB1 swstbkpoldel completed
#
```

NOTES

- A backup policy cannot be deleted if:

    - The transaction volume to be deleted has backup history information.

    - A transaction volume whose backup policies are to be deleted has already started backup synchronous processing.

    - The replication management function is executing a process for the target transaction volume.

- This command can only be executed while none of the following commands are running:

    - swstsvrset (Managed Server configuration information setting command)

    - swstdevinfoset (Device information setting command)

    - swstsrsemtch (Resource match command)

    - swstresback (Resource backup command)

    - swstresrst (Resource restore command)

9.2.2 Operation commands

This section describes operation commands for backup management.

9.2.2.1 swstbackup (Backup execution command)

NAME

swstbackup - backup

SYNOPSIS

Executing a backup by specifying the transaction volume

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstbackup [-suspend|-T] [-k] Device-Name [-Xdevmap Device-Map-File]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstbackup [-h Server-Name] [-suspend|-T] [-k] Device-Name [-Xdevmap Device-Map-File]
```

Executing a backup by specifying a log group

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstbackup -n [-suspend|-T] [-k] Log-Group-Name[/RDB-Name] [-Xdevmap Device-Map-File]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstbackup -n [-h Server-Name] [-k] [-suspend|-T] Log-Group-Name[/RDB-Name] [-Xdevmap Device-Map-File]
```

DESCRIPTION

When using this command, the operation changes according to different states of the Advanced Copy that is being executed.

- When synchronous processing (EC) is not performed
  Snapshot processing (OPC) is performed and the processing of information, such as the TimeStamp and backup volume of a copy target, is registered into backup history information.

- When synchronous processing (EC) is performed
  The state of synchronous processing is checked and, in the case of an equivalent maintenance state, backup synchronous processing is suspended. Information such as time and backup volume of a copy target is simultaneously registered into backup history information. A backup of this form is called synchronous high-speed backup. Note that it aborts when an equivalent maintenance state is not reached. To use the "Chapter 5 Backup Operation using Suspend/Resume Function" for synchronous high-speed backup, execute the backup command with -suspend option specified, and suspend backup synchronous processing.

In the above cases, a backup is established, and access to the backup volume is available immediately after the execution command.

AdvancedCopy Manager executes backup pre-processing and post-processing scripts before and after backup.

For information on customizing these scripts, refer to "Appendix A Pre-processing and Post-processing of Backup and Restoration".

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on a Management Server.

-k

Specifies implementation of backup with reference dump.
If you omit this option, backup will be implemented with normal dump.

When you back up with reference dump, you have to execute Symfoware's "rdbrtr" command in order to set inhibition of updates (disable writing of data) for all DSI (function that expresses the storage structures for tables) on the transaction volumes to be backed up. It cannot specify this option for the transaction volumes other than Symfoware.

-n

Specifies that the item being backed up is a log group.
In the operand, specify the log group name and the RDB system name. If there is no RDB system name, you can omit it.

-Xdevmap *Device-Map-File*

Specifies a target backup volume to perform the snapshot fast backup. In the operand, specify a "device map file" in which a combination of a transaction volume and a target backup volume are described. If this option is not specified, the target backup volume will be automatically selected by the backup function. To execute this command on a remote basis using the -h option, specify a device map file (of a Managed Server) using an absolute pathname. A file name specified in Device-Map-File must not include any national character.

-suspend

This results in synchronous high-speed backup using the "Chapter 5 Backup Operation using Suspend/Resume Function".

-T

Specifies that the differential snapshot high-speed backup be executed.
This option cannot be specified during synchronous processing.
This option is valid only for copying within a box when the ETERNUS Disk storage system supports the QuickOPC function.
If this option is not specified, ordinary snapshot processing (ie, OPC without using the QuickOPC function) is started. This option must always be specified to perform a backup operation using differential snapshots.

OPERANDS

Device-Name

Specify an AdvancedCopy Manager device name.
Refer to "1.5 Managing a Device on AdvancedCopy Manager" for the devices to be specified.

Log-Group-Name

Specify the name of the log group.

RDB-Name

Specify the RDB system name after the log group name, delimited by a slash ("/").
If there is no RDB system name, you can omit it.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Back up a transaction volume:

```
# /opt/FJSVswsts/bin/swstbackup /dev/sdg3
/dev/sdg3 swstbackup completed
#
```

- Back up log group (LOG01/RDB1):

```
# /opt/FJSVswsts/bin/swstbackup -n LOG01/RDB1
LOG01/RDB1 swstbackup completed
#
```

NOTES

- To preserve data integrity, this command performs the backup only after inhibiting access to both transaction and backup volumes from all other processes. Thus, drive letters, logical volume names or device labels of devices are unassigned. This command terminates with an error if a drive letter, logical volume name or device label of a transaction volume is assigned or if:

    - The mount point has directories to which other volumes are also mounted.

    - The volume has a file in use.

- When the -bind option in the mount command is specified to mount part of the file hierarchy of the transaction volume at another mount point, use backup manager's authority to unmount the transaction volume.

- Do not execute the backup execution command in a state where another volume is mounted at the mount point of the target volume. For such a state, unmount the other volume before executing the backup execution command.

- To mount a transaction volume from a server other than the server on which this command is entered, unmount the transaction volume, using the authority of the backup administrator. If the transaction volume must not be unmounted, create backup copies with the following procedure:

    1. Execute the sync command to synchronize the transaction volume and file system.

    2. Execute this command.

    3. Execute the fsck command for the backup volume, and check the file system.

- Backup cannot be executed in the following situations:

    - No backup policy has been specified for the specified transaction volume.

    - The backup that was executed one generation earlier was unsuccessful.

    - Restore is being performed to the transaction volume that is to be backed up.

    - Synchronous backup processing is performed in a transaction volume or a log group to be processed, but the state of equivalence is not reached;

    - Restoration is being executed using backup history information that will cause a generation overflow (this may occur when restoration is performed using the backup volume of the oldest history).

    - Any of the box information (box identifier, OLU number, EXTENT start position, and EXTENT size) of a transaction volume has been changed.

    - No unused backup volume was available for use as the copying destination.

    - The replication management function is executing a process for the target transaction volume.

    - The replication management function is executing a process for the target backup volume.

    - The replication management function is executing a process for backup records exceeding the number of preservation generations.

    - The mount point of a transaction volume set at registration has been changed.

    - A backup volume defined as the copy destination is mounted.

    - An error was detected during backup of Symfoware database spaces.

    - When during backup of transaction volumes with Symfoware's reference dump specified, multi-database spaces are being used.

    - When the -T option is specified during the operation in which SDX objects are backed up in units of logical volumes

    - When the -n option was specified for a log group in which GDS and non-GDS volumes are mixed.

- You cannot perform backup if:

    - The backup volume specified by the device map file cannot be paired with the transaction volume. This occurs if:

        - The specified backup volume is in use.

        - The transaction volume and the backup volume are not the same partition size.

- The device map file contains an error (refer to "3.4.9.1 Describing a device map file").

- The device map file name includes a national character.

- When an ordinary volume is specified as the transaction volume and a GDS logical volume as the backup volume.

- When a GDS logical volume is specified as the transaction volume and an ordinary volume as the backup volume.

- At this time, if existing backup history information satisfies the condition of the number of saved generations specified by swstbkpolset (Backup policy setting command), then information on the oldest generation is automatically deleted from the backup history information during the backup process, and the backup volume is cleared. However, if -suspend or -T option is specified for the backup history information, then the backup volume with the Suspend or tracking state is not cleared and the backup volume remains in use.

- The transaction volume on which the backup command is executed, one other transaction volume, and the suspended or tracking backup volume cannot be specified as backup destinations in the device map file.

- For details on how to use SDX objects of GDS, refer to "Configuration and conditions of unsupported SDX objects" and "10.1.5.1 Notes on SDX object operations (backup management)".

- The pre-processing and post-processing scripts are not executed when the mirror partition of the SDX object is TEMP.

- This command can be executed only while none of the following commands are running:

    - swstsvrset (Managed Server configuration information setting command)

    - swstdevinfoset (Device information setting command)

    - swstsrsemtch (Resource match command)

    - swstresback (Resource backup command)

    - swstresrst (Resource restore command)

9.2.2.2 swstbackstat (Backup execution status display command)

NAME

swstbackstat - displays the completion status of a backup

SYNOPSIS

Displaying the backup execution status by specifying the transaction volume

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstbackstat [Device-Name]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstbackstat [-h Server-Name] [Device-Name]
```

Displaying the backup execution status by specifying a log group

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstbackstat -n Log-Group-Name[/RDB-Name]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstbackstat -n [-h Server-Name] Log-Group-Name[/RDB-Name]
```

DESCRIPTION

This command displays the completion status of a backup that is being performed using swstbackup (Backup execution command).

The following information will be displayed:

Title	Description
Server	Displays a Managed Server name.
Transaction-Disk	Displays the AdvancedCopy Manager device name of a transaction volume.
Backup-Disk	Displays a backup volume name. If no backup has been collected, "----" will be displayed.
Status	Displays a backup completion status using one of the following character strings: "----" : No backup has been collected (ie, no history exists). "succeeded" : Backup has been completed. "executing" : Copying using OPC is in progress. "failed" : Copying using OPC has been interrupted due to an error. "gds-error" : Copying with GDS ended abnormally.
Execute	Displays how much has been copied in percentage terms if "executing" is displayed in the Status field. Displays "----" unless "executing" is displayed.

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on a Management Server.

-n

Specifies that the item being displayed is a log group.
In the operand, specify the log group name and the RDB system name. If there is no RDB system name, you can omit it.

OPERANDS

Device-Name

Specifies the "1.5 Managing a Device on AdvancedCopy Manager" corresponding to a transaction volume.
If this operand is omitted, the backup execution statuses of all the transaction volumes will be displayed.

Log-Group-Name

Specify the name of the log group.

RDB-Name

Specify the RDB system name after the log group name, delimited by a slash ("/").
If there is no RDB system name, you can omit it.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Display the actual copying execution status:

```
# /opt/FJSVswsts/bin/swstbackstat
Server Transaction-Disk Backup-Disk Status    Execute
Work1  /dev/sdb3         /dev/sdd3   succeeded ----
Work1  /dev/sdb4         /dev/sdd4   executing 70%
Work1  /dev/sdb5         /dev/sdd5   executing 30%
```

```
:
#
```

NOTES

- Do not use the OPC stop function of ETERNUS Web GUI while the snapshot fast backup is in progress. If you do use the OPC stop function of ETERNUS Web GUI, it will display "succeeded" whether or not the backup has actually been successful.

- If "failed" is displayed in the "Status" field, a hardware error may have occurred. Refer to "8.3 Troubleshooting: Hardware error during backup" for information about this.

- If "gds-error" is displayed in the Status field, a problem with GDS is assumed.
  Remedy the cause of the GDS problem, and take one of the following steps:

    - If backup synchronous processing is in progress, execute swstcancelsync (Backup synchronous processing cancel command) to cancel the backup synchronous processing.

    - If snapshot backup is in progress, execute swsthistdel (History information deletion command) to delete the backup history information.

    - If restoration is in progress, execute swstsrsemtch (Resource match command), and perform restoration again.

9.2.2.3 swstrestore (Restore execution command)

NAME

swstrestore - restore

SYNOPSIS

For restore of a regular transaction volume

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstrestore [-g Generation-No | -v Version-No] [-r Restore-Device-Name] [-
Xgds-softcopy] Device-Name
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstrestore [-h Server-Name] [-g Generation-No | -v Version-No] [-r Restore-
Device-Name] [-Xgds-softcopy] Device-Name
```

For recover of a Symfoware database space by specifying a transaction volume

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstrestore [-g Generation-No | -v Version-No] [-l ArchiveLogFileList-
FileName] [-w Work-Directory] [-nolog Work-RecCtrlFile-Directory] [-Xgds-softcopy] Device-Name
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstrestore [-h Server-Name] [-g Generation-No | -v Version-No] [-l
ArchiveLogFileList-FileName] [-w Work-Directory] [-nolog Work-RecCtrlFile-Directory] [-Xgds-
softcopy] Device-Name
```

For bundle recovery by specifying multiple transaction volumes

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstrestore -bundle [-g Generation-No | -v Version-No] [-a | -e EndPoint] [-
l ArchiveLogFileList-FileName] [-w Work-Directory] [-nolog Work-RecCtrlFile-Directory] [-Xgds-
softcopy] DeviceListFile
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstrestore -bundle [-h Server-Name] [-g Generation-No | -v Version-No] [-a
| -e EndPoint] [-l ArchiveLogFileList-FileName] [-w Work-Directory] [-nolog Work-RecCtrlFile-
Directory] [-Xgds-softcopy] DeviceListFile
```

For recovery of a Symfoware database space by specifying a log group

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstrestore -n [-g Generation-No | -v Version-No] [-a | -e EndPoint] [-l
ArchiveLogFileList-FileName] [-w Work-Directory] [-nolog Work-RecCtrlFile-Directory] [-Xgds-
softcopy] Log-Group-Name[/RDB-Name]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstrestore -n [-h Server-Name] [-g Generation-No | -v Version-No] [-a | -e
EndPoint] [-l ArchiveLogFileList-FileName] [-w Work-Directory] [-nolog Work-RecCtrlFile-
Directory] [-Xgds-softcopy] Log-Group-Name[/RDB-Name]
```

DESCRIPTION

The data that exists in backup history information is restored using OPC.
When this command is executed, AdvancedCopy Manager performs the restoration pre-processing and post-processing scripts before and after copying data using OPC. For more information on customizing these scripts, refer to "Appendix A Pre-processing and Post-processing of Backup and Restoration".

You can restore data to a device that is not a transaction volume. When you do so, the device must not be registered as a transaction volume and the device must have the same amount of space available as the specified transaction volume.

For transaction volumes of Symfoware and logical volumes of GDS, restoring to devices other than the original transaction volumes is not available.

During recovery of database spaces, access to the relevant database spaces must not be permitted. In order to prohibit access, execute the "rdbinh" or the "rdbexspc" command provided by Symfoware. For details on these commands, refer to the "Symfoware Server RDB Operations Guide".

If recovery of a database space fails due to insufficient space in the working directory, you can use the -w option to retry recovery temporarily specifying a different directory.

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on a Management Server.

-g *Generation-No*

Specifies the relative generation number of data to be restored.
Check the relative generation number using swsthistdisp (History information display command).

-v *Version-No*

Specifies the absolute generation number of data to be restored.
Check the absolute generation number using swsthistdisp (History information display command).

-r *Restore-Device-Name*

Specifies a restore destination device name to restore data to any device other than a transaction volume.

Cannot be specified for a Symfoware transaction volume.

-a

Specifies that the entire database resources managed by a recovery control file shall be recovered up to the point in time of backup collection.

Cannot be specified at the same time as the -e option.
Cannot be specified when recovering individual transaction volumes in a log group.
Cannot be specified for the transaction volumes other than Symfoware.

-e *EndPoint*

Specifies that the entire database resources managed by a recovery control file shall be recovered up to a specific recovery point. For details on recovery points, refer to the Symfoware Server manuals.
Cannot be specified at the same time as the -a option.
Cannot be specified at the same time as the -nolog option.
Cannot be specified when recovering individual transaction volumes in a log group.
Cannot be specified for the transaction volumes other than Symfoware.

-l *ArchiveLogFileList-FileName*

Specifies the absolute path of the file stating the names of the backup locations for the archive logs that are required for recovery of Symfoware.
The file stating the name of the backup locations for the archive logs must be stored on the Managed Server that contains the transaction volumes to be restored.
If you execute this command on a Management Server, the file must be stored on the Managed Server that contains the transaction volumes to be restored.
Cannot be specified at the same time as the -nolog option.
Cannot be specified for the transaction volumes other than Symfoware.

-w *Work-Directory*

Specifies the absolute path to the name of a temporary working directory.
The directory must be located on the Managed Server that contains the transaction volumes to be restored.
If you execute this command on a Management Server, the directory must be located on the Managed Server that contains the transaction volumes to be recovered.
If this option is omitted, the working directory specified in swstsvrset (Managed Server configuration information setting command) will be used.
Cannot be specified for the transaction volumes other than Symfoware.

-n

Specifies that the item being restored is a log group.
In the operand, specify the log group name and the RDB system name. If there is no RDB system name, you can omit it.

-bundle

Specifies bundle recovery of multiple database spaces.
In the operand, specify the file (DeviceListFile) stating the device names corresponding with the database spaces to be recovered in a batch.

-nolog *Work-RecCtrlFile-Directory*

During recovery of Symfoware, archive logs are not applied.
If you want to specify an output directory for the recovery control files that are required when applying the log, and then apply the log by Symfoware's "rdbmrrcv" command, you can specify this directory.
If you perform recovery with this option specified, the database will not be available for use until you apply the log with Symfoware's "rdbmrrcv" command.

-Xgds-softcopy

Performs restoration by software-copy for restoration in linkage with GDS Snapshot. If this option is not specified, restoration is performed with OPC.

Note: If neither the -g or -v option is specified, the latest information from the information history management will be restored.

OPERANDS

Device-Name

Specifies the AdvancedCopy Manager device name corresponding to a transaction volume.
Refer to "1.5 Managing a Device on AdvancedCopy Manager" for the devices to be specified.

Log-Group-Name

   Specify the name of the log group.

RDB-Name

   Specify the RDB system name after the log group name, delimited by a slash ("/").
   If there is no RDB system name, you can omit it.

DeviceListFile

   Specify the device list file that states the device names corresponding with the database spaces to be recovered in a batch.
   If you execute this command on a Management Server, this file must be stored on the Managed Server that contains the transaction volumes to be recovered.
   For information on how to describe device list files, refer to "4.1.5.3.1 Describing a device list file".

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Restore data with relative generation number 2. For information on the relative generation number, refer to data displayed using swsthistdisp (History information display command):

```
# /opt/FJSVswsts/bin/swstrestore -g 2 /dev/sdb6
/dev/sdb6 swstrestore completed
#
```

- Restore data with relative generation number 2 for log group (LOG01/RDB1):

```
# /opt/FJSVswsts/bin/swstrestore -n -g 2 LOG01/RDB1
LOG01/RDB1 swstrestore completed
#
```

- If the "Chapter 5 Backup Operation using Suspend/Resume Function" is used for the backup operation on the transaction volume, the suspended synchronous processing is deleted and the restore command is executed. For information on the suspended synchronous processing, refer to the data displayed by swstsyncstat (Backup synchronous processing progress display command). Execute swstcancelsync (Backup synchronous processing cancel command) to cancel the synchronous processing:

```
# /opt/FJSVswsts/bin/swstsyncstat /dev/sdb6
Server Transaction-Disk Backup-Disk Status    Execute
Work2  /dev/sdb6        /dev/sdd6    suspend   ----
Work2  /dev/sdb6        /dev/sde5    executing 75%
# /opt/FJSVswsts/bin/swstcancelsync -all /dev/sdb6
/dev/sdb6 swstcancelsync completed
# /opt/FJSVswsts/bin/swstsyncstat /dev/sdb6
Server Transaction-Disk Backup-Disk Status Execute
Work2  /dev/sdb6        ----         ----   ----
# /opt/FJSVswsts/bin/swstrestore /dev/sdb6
/dev/dsk/sdb6 swstrestore completed
#
```

NOTES

- On a transaction volume on which a file system has been constructed, restore a specific file as follows:

   1. Mount a backup volume. Check the backup volume by executing the Backup History List view/swsthistdisp (History information display command).

   2. Copy the file to be restored by executing, for example, the cp command.

3. Unmount the backup volume.

- If a transaction volume is mounted from a server other than the server on which this command is entered, execute this command after unmounting the transaction volume, using the authority of the backup administrator.

- When the -bind option in the mount command is specified to mount part of the file hierarchy of the transaction volume at another mount point, use the backup administrator's authority to unmount the transaction volume.

- Do not execute the restore execution command in a state where another volume is mounted at the mount point of the target volume. In this case, unmount the other volume before executing the restore execution command.

- Restoration or recovery cannot be executed in the following situations:

  - When both relative and absolute generation numbers have been specified.

  - When a device registered as a transaction volume has been specified in a parameter with the -r option.

  - When a device registered as a backup volume has been specified in a parameter with the -r option, the device has been registered in the backup registration information of a certain transaction volume. If you restore data to a backup volume without registered backup data, and the backup volume is used to back up a certain transaction volume, then data will be overwritten. It is recommended that a backup volume should NEVER be specified as a restore destination device.

  - When you perform restoration immediately after starting backup where the actual copying using OPC is still in progress, (except when the combination of transaction volume and backup volume is the same for backup and restoration). A restoration should always be performed after the actual copying using OPC is completed. Alternatively, on a transaction volume on which a file system has been constructed, perform restoration according to the method described in the Point "In a transaction volume on which a file system has been constructed, restore a specific file" above.

  - A transaction volume currently in use with synchronous backup processing is specified.

  - When any of the box information (box identifier, OLU number, EXTENT start position, and EXTENT size) of a transaction volume to be restored has been changed.

  - The replication management function is executing a process for the target transaction volume.

  - The replication management function is executing a process for the device specified in the -r option.

  - The mount point set at registration of a transaction volume has been changed.

  - The -Xgds-softcopy option is specified for a device that is not a GDS logical volume.

  - A GDS logical volume is specified as the restoration destination volume of an ordinary transaction volume.

  - The -n or -bundle option was specified when GDS volumes and non-GDS volumes are mixed in a log group.

  - Different backup methods (normal dump or reference dump) to recover data exist in the same log group, when you specify a log group or you perform recovery of multiple volumes in a batch,

  - The -r option was specified when you perform recovery of the transaction volumes of Symfoware.

  - When an SDX object has been specified in a parameter with the -r option.

  - When the -a, -e, -l, -w, -bundle, or -nolog option was specified for setting transaction volumes other than Symfoware.

  - When a directory name specified in the -w, -l, -bundle, or -nolog option contains double-byte characters (as used in Asian languages).

  - An error was detected when during recovery of Symfoware.

  - When the file stating the names of the archive log backup files contains an archive log backup file that is backed up on an MT (tape device).

- The suspended or tracking backup volume cannot be used as the restore destination volume.

- The restore command cannot be executed on a transaction volume during synchronous processing (eg, during copying by EC, under the equivalency maintained status, in the Suspend state).
  If the following errors occur during restore processing, then execute swstcancelsync (Backup synchronous processing cancel

command) (with the -all option specified) to cancel all of the synchronous processing of the transaction volume, and execute the restore command:

  - "swst0639 This Command cannot be run. This is because a copy is currently running"

  - "swst0634 OPC cannot be run. Error code=26 (SCSI command busy)"

- When the archive log files are stored on a different device at the time of recovering database spaces, the file which states the names of the backup locations for the archive logs and is required for recovery, must be stored beforehand on the Managed Server that contains the transaction volumes to be recovered. For information on how to describe this file, refer to the "Symfoware Server RDB Operations Guide".

- Recovery cannot be performed if the specified "file that includes the name of the archive log's backup location" contains an archive log backup file that is backed up on an MT (tape device). For recovery using an archive log that is backed up on an MT, use one of the following methods:

  - Deploy all backed up archive logs from the MT to the working directory.

  - Perform recovery without applying the log by specifying the -nolog option, and apply the archive log afterwards.

    In order to apply the archive log after recovery without the log, execute the Symfoware's "rdbmrrcv" command. For details on applying archive logs, refer to the "Symfoware Server RDB Operations Guide".

- If you perform "recovery of a specific point in time by specifying the end point of recovery" or "recovery of the point in time of backup" during recovery of database spaces, processing to restore Symfoware's management information will also be implemented. Since this processing is implemented as a portion of recovery processing, it takes time for commands to be processed.

- For details on how to use SDX objects of GDS, refer to "Configuration and conditions of unsupported SDX objects" and "10.1.5.1 Notes on SDX object operations (backup management)".

- This command can only be executed while none of the following commands are running:

  - swstsvrset (Managed Server configuration information setting command)

  - swstdevinfoset (Device information setting command)

  - swstsrsemtch (Resource match command)

  - swstresback (Resource backup command)

  - swstresrst (Resource restore command)

9.2.2.4 swstcancelrest (Restore cancellation command)

NAME

swstcancelrest - cancels OPC-based restoration

SYNOPSIS

Canceling a restore by specifying transaction volume

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstcancelrest [-g Generation-No | -v Version-No] [-r Restore-Device-Name] [-
emergency] Device-Name
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstcancelrest [-h Server-Name][-g Generation-No | -v Version-No][-r Restore-
Device-Name] [-emergency] Device-Name
```

Canceling a restore by specifying a log group (Symfoware)

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstcancelrest -n [-g Generation-No | -v Version-No ] [-emergency] Log-Group-
Name[/RDB-Name]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstcancelrest -n [-h Server-Name] [-g Generation-No | -v Version-No ] [-
emergency] Log-Group-Name[/RDB-Name]
```

DESCRIPTION

This command cancels OPC-based restoration.

A hardware error will be displayed when swstrestore (Restore execution command) is executed while data is being copied using OPC.
Fix the cause of the hardware error and repeat restoration using swstrestore (Restore execution command).

Check whether a fault occurred using swstreststat (Restore execution status display command).

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on a Management Server.

-g *Generation-No*

Specifies the relative generation number of data being restored.
This option cannot be specified at the same time as the -v option.
Check the relative generation number using swstreststat (Restore execution status display command).

-v *Version-No*

Specifies the absolute generation number of data being restored.
This option cannot be specified at the same time as the -g option.
Check the absolute generation number using swstreststat (Restore execution status display command).

-r *Restore-Device-Name*

Specifies the device name if data is restored to any other device than a transaction volume.
Check the device name using swstreststat (Restore execution status display command).

It cannot specify this option for the transaction volumes of Symfoware.

-emergency

Emergency operation mode is available for operations.
In such cases, OPC sessions are neither confirmed nor cancelled.

-n

Specifies that the restore is being cancelled for a log group.
In the operand, specify the log group name and the RDB system name. If there is no RDB system name, you can omit it.

Note: The restoration of the latest history information will be cancelled if neither the -g or -v option is specified.

OPERANDS

Device-Name

Specify an AdvancedCopy Manager device name.
Refer to "1.5 Managing a Device on AdvancedCopy Manager" for the devices to be specified.

Log-Group-Name

Specify the name of the log group.

RDB-Name

Specify the RDB system name after the log group name, delimited by a slash ("/").
If there is no RDB system name, you can omit it.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Cancel restoration of the transaction volume:

```
# /opt/FJSVswsts/bin/swstcancelrest -g 1 /dev/sdb6
/dev/sdb6 swstcancelrest completed.
#
```

- Restoration of the latest history information on the transaction volume is cancelled in emergency operation mode:

```
# /opt/FJSVswsts/bin/swstcancelrest -emergency /dev/sdb6
/dev/sdb6 swstcancelrest completed.
#
```

- Cancel restoration of log group (LOG01/RDB1):

```
# /opt/FJSVswsts/bin/swstcancelrest -n  -g  1 LOG01/RDB1
LOG01/RDB1 swstcencelrest completed.
#
```


NOTES

- If you cancel OPC-based restoration using this command, the transaction volume will be in an incomplete state and can no longer be used because part of the data has already been overwritten by the copy process. Perform restoration again using swstrestore (Restore execution command).

  If the transaction volume is a GDS logical volume and restoration is cancelled, the state of the transaction volume becomes INVALID. In such cases, refer to the "PRIMECLUSTER(TM) Global Disk Services User's Guide (Linux)," and restore the logical volume.

- You cannot cancel restoration if:

  - The history specified in the -g or -v option does not exist.

  - Restoration is not performed on the device specified in the -r option.

  - Restoration is not performed from the specified history. In this case, processing is terminated with an information message "swst0303 Restoration for specified restoration device is not performed". The end status after this message is a normal end.

  - The -n option was specified when GDS volumes and non-GDS volumes are mixed in a log group.

  - The replication management function is executing a process for the target transaction volume

- However, OPC sessions are neither confirmed nor cancelled by execution in emergency operation mode. If an OPC session is in the error suspended state, then unmanaged sessions remain. In this case, ETERNUS Web GUI must be used to cancel these sessions.

- This command can only be executed while none of the following commands are running:

  - swstsvrset (Managed Server configuration information setting command)

  - swstdevinfoset (Device information setting command)

9.2.2.5 swstreststat (Restore execution status display command)

NAME

swstreststat - displays the execution status of restoration

SYNOPSIS

Displaying the restore execution status by specifying the transaction volume

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstreststat [-g Generation-No | -v Version-No] [-Xgds-softcopy] [Device-Name]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstreststat [-h Server-Name] [-g Generation-No | -v Version-No] [-Xgds-softcopy] [Device-Name]
```

Displaying the restore execution status by specifying a log group

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstreststat -n [-g Generation-No | -v Version-No]  [-Xgds-softcopy] Log-Group-Name[/RDB-Name]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstreststat -n [-h Server-Name] [-g Generation-No | -v Version-No]  [-Xgds-softcopy] Log-Group-Name[/RDB-Name]
```

DESCRIPTION

This command displays the execution status of restoration.
The following information will be displayed:

Title	Description
Server	Displays a Managed Server name.
Transaction-Disk	Displays the AdvancedCopy Manager device name of a transaction volume.
Generation	Displays the relative generation number of restored backup data. "----" will be displayed if no generation number is specified in the -g or -v option and no backup history exists.
Version	Displays the absolute generation number of restored backup data. "----" will be displayed if no generation number is specified in the -g or -v option and no backup history exists.
Backup-Disk	Displays the name of a backup volume from which data should be copied and restored. "----" will be displayed if restoration is not in progress.
Restore-Device	Displays a restoration target device name unless data is copied and restored to a transaction volume. "----" will be displayed if restoration is not in progress, or if data is restored to a transaction volume.

Title	Description
Status	Displays the execution status: "----" : No copying using OPC or software-copy by GDS is in progress. "executing" : Copying using OPC or software-copy by GDS is in progress. "succeeded" : Copying using OPC or software-copy by GDS has been completed. "failed" : Copying using OPC has been interrupted due to an error. "gds-error" : Copying with GDS ended abnormally.
Execute	Displays how much has been copied in percentage terms if "executing" is displayed in the Status field. Displays "----" unless "executing" is displayed.

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on a Management Server.

-g *Generation-No*

Specifies the relative generation number of data whose restore execution status is to be displayed.

-v *Version-No*

Specifies the absolute generation number of data whose restore execution status is to be displayed.

-n

Specifies that the execution status of restoration being displayed is for a log group.
In the operand, specify the log group name and the RDB system name. If there is no RDB system name, you can omit it.

-Xgds-softcopy

Displays the restore execution status of a GDS logical volume.
Use this option to check the copy status when restoration using the software-copy function of GDS is in progress.
When the option is specified, the Status field display when copying is completed is "----", which is the same as before copying started.

Note: The restore execution statuses of all the histories will be displayed if both the -g and -v options are omitted.

OPERANDS

Device-Name

Specify an AdvancedCopy Manager device name.
Refer to "1.5 Managing a Device on AdvancedCopy Manager" for the devices to be specified.
If this operand is omitted, the restore execution statuses of all the transaction volumes will be displayed.

Log-Group-Name

Specify the name of the log group.

RDB-Name

Specify the RDB system name after the log group name, delimited by a slash ("/").
If there is no RDB system name, you can omit it.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Display the restore execution statuses of transaction volumes:

```
# /opt/FJSVswsts/bin/swstreststat
Server Transaction-Disk Generation Version Backup-Disk   Restore-Device Status    Execute
Work2  /dev/sdb3        1          10      /dev/sdd3     /dev/sdc3      executing 75%
Work2  /dev/sdb4        ----       ----    ----          ----          ----      ----
Work2  /dev/sdb5        2          12      /dev/sdd5     ----          succeeded ----
Work2  /dev/sdb6        3          13      /dev/sdd6     /dev/sdc6      executing 75%
:
#
```

NOTES

- Do not use the OPC stop function of ETERNUS Web GUI while the snapshot fast backup is in progress. If you use the OPC stop function of ETERNUS Web GUI, this command displays "succeeded" whether or not the restoration has actually been successful.

- If "failed" is displayed in the "Status" field, a hardware error may have occurred. Refer to "8.3 Troubleshooting: Hardware error during backup" for information about this.

- If no generation number has been specified in the -v or -g option, the restore execution statuses of all the histories will be displayed. If, at this time, restoration is not in progress from any of the histories, "----" will be displayed in all the information fields other than "Server" and "Transaction-Disk".

- If "gds-error" is the status displayed in the Status field, a problem with GDS is assumed. Correct the cause of the GDS problem, and take action as follows:

  - If backup synchronous processing is in progress, execute swstcancelsync (Backup synchronous processing cancel command) to cancel the backup synchronous processing.

  - If snapshot backup is in progress, execute swsthistdel (History information deletion command).

  - If restoration is in progress, execute swstsrsemtch (Resource match command), and perform restoration again.

9.2.2.6 swsthistdisp (History information display command)

NAME

swsthistdisp - displays the backup history information

SYNOPSIS

Displaying the backup history information by specifying the transaction volume

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swsthistdisp [Device-Name]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swsthistdisp [-h Server-Name] [Device-Name]
```

Displaying the backup history information by specifying a log group

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swsthistdisp -n Log-Group-Name[/RDB-Name]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swsthistdisp -n [-h Server-Name] Log-Group-Name[/RDB-Name]
```

DESCRIPTION

This command displays history information in relation to data that has been backed up.

The following information will be displayed:

Title	Description
Server	Displays a Managed Server name.
Device	Displays the AdvancedCopy Manager device name.
Mount-Point	Displays the mount point of a device. The file system type of the mount point is displayed inside the parentheses. If the "Device" is a transaction volume used for Symfoware, the names of Symfoware database spaces and log groups as well as the RDB system names are displayed, followed by "SymfoWARE" in parentheses.
Generation	Displays a relative generation number. This item will not be displayed if no backup history information exists.
Version	Displays an absolute generation number. This item will not be displayed if no backup history information exists.
Backup-Date	Displays a TimeStamp reflecting when backup has been completed. This item will not be displayed if no backup history information exists.
Backup-Disk	Displays the name of a backup volume to which data has been backed up. This item will not be displayed if no backup history information exists.
Status	Displays an actual copying execution status using one of the following character strings: "succeeded" : Copying using OPC has been completed. "executing" : Copying using OPC is in progress. "failed" : Copying using OPC has been interrupted due to an error. "gds-error" : Copying with GDS ended abnormally. This item will not be displayed if no backup history information exists.
Execute	Displays how much has been copied in percentage terms if "executing" is displayed in the Status field. Displays "----" unless "executing" is displayed. This item will not be displayed if no backup history information exists.
ArcSerial	Displays the serial number of the archive log. The history information for backups performed by reference dump displays "reference". Information other than for Symfoware transaction volumes is not displayed. This item will not be displayed if no backup history information exists.

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on a Management Server.

-n

Specifies that the backup history information being displayed is for a log group.
In the operand, specify the log group name and the RDB system name. If there is no RDB system name, you can omit it.

OPERANDS

Device-Name

Specify an AdvancedCopy Manager device name.
Refer to "1.5 Managing a Device on AdvancedCopy Manager" for the devices to be specified.
If this operand is omitted, the backup history information of all the transaction volumes will be displayed.

Log-Group-Name

Specify the name of the log group.

RDB-Name

Specify the RDB system name after the log group name, delimited by a slash ("/").
If there is no RDB system name, you can omit it.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Display the backup history information of all the transaction volumes:

```
# /opt/FJSVswsts/bin/swsthistdisp
Server=StrgSV01 Device=/dev/sdb6 Mount-Point=/mnt/Tran1 (ext2)
Generation Version Backup-Date      Backup-Disk   Status     Execute
1          10      2000/11/12 22:00 /dev/sdd6     succeeded ----
2          9       2000/11/11 22:00 /dev/sdf6     succeeded ----
:
#
```

NOTES

- If "failed" is displayed in the "Status" field, a hardware error may have occurred. Refer to "8.3 Troubleshooting: Hardware error during backup" for information about this.

- If "gds-error" is the status displayed in the Status field, a problem with GDS is assumed. Correct the cause of the GDS problem, and take action as follows:

  - If backup synchronous processing is in progress, execute swstcancelsync (Backup synchronous processing cancel command) to cancel the backup synchronous processing.

  - If snapshot backup is in progress, execute swsthistdel (History information deletion command).

  - If restoration is in progress, execute swstsrsemtch (Resource match command), and perform restoration again.

9.2.2.7 swsthistdel (History information deletion command)

NAME

swsthistdel - deletes backup history information

SYNOPSIS

Deleting the backup history information by specifying the transaction volume

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swsthistdel -g Generation-No | -v Version-No | -z [-emergency] Device-Name
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swsthistdel [-h Server-Name] -g Generation-No | -v Version-No | -z [-emergency] Device-Name
```

Deleting the backup history information by specifying a log group

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swsthistdel -n -g Generation-No | -v Version-No | -z [-emergency] Log-Group-Name[/RDB-NAME]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swsthistdel -n [-h Server-Name] -g Generation-No | -v Version-No | -z [-emergency] Log-Group-Name[/RDB-NAME]
```

DESCRIPTION

This command deletes backup history information.
The backup volume of the deleted history is released and becomes an empty backup volume.

When you delete the history information for Symfoware transaction volumes, the recovery control files will also be deleted.

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can only be specified on a Management Server.

-g *Generation-No*

Specifies the relative generation number of backup history information to be deleted.
This option cannot be specified at the same time as the -v or -z option.
Check the relative generation number using swsthistdisp (History information display command).

-v *Version-No*

Specifies the absolute generation number of backup history information to be deleted.
This option cannot be specified at the same time as the -g or -z option.
Check the absolute generation number using swsthistdisp (History information display command).

-z

Specifies that all the backup history information is to be deleted.
This option cannot be specified at the same time as the -g or -v option.

-emergency

Emergency operation mode is available for operations.
In such cases, OPC sessions are neither confirmed nor cancelled.

-n

Specifies that the backup history being deleted is for a log group.
In the operand, specify the log group name and the RDB system name. If there is no RDB system name, you can omit it.

OPERANDS

Device-Name

    Specify an AdvancedCopy Manager device name.
    Refer to "1.5 Managing a Device on AdvancedCopy Manager" for the devices to be specified.

Log-Group-Name

    Specify the name of the log group.

RDB-Name

    Specify the RDB system name after the log group name, delimited by a slash ("/").
    If there is no RDB system name, you can omit it.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Delete a history with absolute generation number 10 of a transaction volume:

```
# /opt/FJSVswsts/bin/swsthistdel -v 10 /dev/sdb6
/dev/sdb6 swsthistdel completed
#
```

- The history information of relative generation number 1 of the transaction volume is deleted in the emergency operation mode:

```
# /opt/FJSVswsts/bin/swsthistdel -g 1 -emergency /dev/sdb6
/dev/sdb6 swsthistdel completed
#
```

- Delete the backup history of absolute generation number 10 for log group (LOG01/RDB1):

```
# /opt/FJSVswsts/bin/swsthistdel -n -v 10 LOG01/RDB1
/dev/sdd1 swsthistdel completed
/dev/sdd2 swsthistdel completed
    :
LOG01/RDB1 swsthistdel completed
#
```

NOTES

- You cannot delete backup history information if:

    - The backup history information specified in the -g or -v option does not exist.

    - The specified backup history information is currently being used in a restoration.

    - The replication management function is executing a process for the target transaction volume.

    - The replication management function is executing a process for the target backup volume.

- However, OPC sessions are neither confirmed nor cancelled by execution in emergency operation mode. If an OPC session is in the error suspended state, then unmanaged sessions remain. In this case, ETERNUS Web GUI must be used to cancel these sessions.

- Deleting the suspended history information does not cancel the synchronous processing (in an EC session).
  Deleting the tracking history information does not cancel the tracking processing (in an OPC session).

- In synchronous high-speed backup with -suspend option specified, the backup volume remains in use even after the history information is deleted.

In differential snapshot high-speed backup with -T option specified, the backup volume remains in use even after the history information is deleted.

- If the history information is deleted while the OPC physical copy has not been completed in a backup operation mode other than the differential snapshot high-speed backup, physical copying of the deleted history is interrupted. In this case, the data in the backup volume is incomplete (ie, copy has stopped part way through). The backup volume is put in the unused state upon deletion of the history information and becomes ready for use in the next backup.

  In the differential snapshot high-speed backup mode, the physical copying of the deleted history is continued even after the history information on the incomplete OPC physical copy is deleted. Even after deletion of the history information, the backup volume is in the busy state and becomes ready for use in the next backup of the relevant transaction volume.

- For details on how to use SDX objects of GDS, refer to "Configuration and conditions of unsupported SDX objects" and "10.1.5.1 Notes on SDX object operations (backup management)".

- This command can be executed only while none of the following commands are running:

  - swstsvrset (Managed Server configuration information setting command)

  - swstdevinfoset (Device information setting command)

  - swstsrsemtch (Resource match command)

  - swstresback (Resource backup command)

  - swstresrst (Resource restore command)

9.2.2.8 swststartsync (Backup synchronous processing start command)

NAME

swststartsync - starts backup synchronous processing

SYNOPSIS

Starting backup synchronous processing by specifying the transaction volume

  To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swststartsync [-Xgds-softcopy | -Xgds-selectcopy] Device-Name [-Xdevmap
Device-Map-File]
```

  To be executed on the Management Server

```
/opt/FJSVswsts/bin/swststartsync [-h Server-Name] [-Xgds-softcopy | -Xgds-selectcopy] Device-
Name [-Xdevmap Device-Map-File]
```

Starting backup synchronous processing by specifying a log group

  To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swststartsync [-Xgds-softcopy | -Xgds-selectcopy] -n Log-Group-Name[/RDB-
NAME] [-Xdevmap Device-Map-File]
```

  To be executed on the Management Server

```
/opt/FJSVswsts/bin/swststartsync -n [-h Server-Name] [-Xgds-softcopy | -Xgds-selectcopy] Log-
Group-Name[/RDB-NAME] [-Xdevmap Device-Map-File]
```

DESCRIPTION

This command starts backup synchronous processing (ie, it starts disk-to-disk copy using EC).

Even if it is already in use, the transaction volume where the synchronous processing started s well as the suspended backup volume can be used as backup destinations.

Except for the transaction volume where the synchronous processing started and the suspended backup volume, transaction volumes cannot be specified as backup destinations in the device map file.

If you execute this command with specifying a log group name, backup synchronous processing will be executed for all transaction volumes that are included in the specified log group. If execution of this command terminates with an error, backup synchronous processing that was performed so far will be cancelled. However, if the status was "suspend" before command execution, even after terminating with an error, synchronous processing will not be cancelled.
Also, you can run this command specifying individual transaction volumes that are included in a log group.

For a GDS transaction volume, synchronous processing by software-copy is also possible.

OPTIONS

-h *Server-Name*

Specify the name of a Managed Server.
This option can be specified only on the Management Server.

-Xdevmap *Device-Map-File*

Specifies a target backup volume to perform the backup synchronous processing. In the operand, specify a device map file in which a combination of a transaction volume and a target backup volume is described. If this option is not specified, a target backup volume will be automatically selected. To execute this command on a remote basis using the -h option, specify a device map file of a Managed Server using an absolute pathname. A file name specified in *Device-Map-File* must not include any national characters.

-n

Specifies that backup synchronous processing is being started for a log group.
In the operand, specify the log group name and the RDB system name. If there is no RDB system name, you can omit it.

-Xgds-softcopy

Specifies the start of backup synchronous processing of an SDX object and use of the software-copy function instead of the Advanced Copy function.
If you specify this option to start differential copying when the Advanced Copy function is being used, Advanced Copy processing stops and software-copy processing starts.
You can specify the option when the function in linkage with GDS Snapshot is being used.

-Xgds-selectcopy

Specifies that GDS select the copy processing method of an SDX object.
You can specify the option when the function in linkage with GDS Snapshot is being used.

OPERANDS

Device-Name

Specify an AdvancedCopy Manager device name.
Refer to "1.5 Managing a Device on AdvancedCopy Manager" for the devices to be specified.

Log-Group-Name

Specify the name of the log group.

RDB-Name

Specify the RDB system name after the log group name, delimited by a slash ("/").
If there is no RDB system name, you can omit it.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- To start the backup synchronous processing for a transaction volume:

```
# /opt/FJSVswsts/bin/swststartsync /dev/sdb6
/dev/sdb6 swststartsync completed.
#
```

- To start backup synchronous processing for log group (LOG01/RDB1):

```
# /opt/FJSVswsts/bin/swststartsync -n LOG01/RDB1
LOG01/RDB1 swststartsync completed.
#
```

- Backup synchronous processing of a logical volume (/dev/sfdsk/CLS01/dsk/VOL01) of an SDX object is started. The copy processing method is software-copy:

```
# /opt/FJSVswsts/bin/swststartsync -Xgds-softcopy /dev/sfdsk/CLS01/dsk/VOL01
/dev/sfdsk/CLS01/dsk/VOL01 swststartsync completed.
#
```

NOTES

- When you execute the backup synchronous processing start command for a transaction volume or log group to which the backup synchronous processing is already being executed, the message "swst0301 Backup synchronous processing is already in progress." is output. (When you specify a log group, the message id is "swst0302".)
  And then the command is terminated normally.

- You cannot start backup synchronous processing in the following cases. In any of the cases below, perform the action according to the "System administrator response" message.

  - When specifying a device that does not support the EC function.

  - When the backup volumes required for starting the backup synchronous processing cannot be obtained.

  - When specifying a device that is being restored by OPC.

  - When mounting the backup volume to which data is copied.

  - The -Xgds-softcopy option is specified for a device that is not a GDS logical volume. Alternatively, the -Xgds-selectcopy option is specified for a device that is not a GDS logical volume.

  - The -n option was specified when GDS volumes and non-GDS volumes are mixed in a log group.

  - The replication management function is executing a process for the target transaction volume.

- If a suspended backup volume is the transaction volume at the start of the synchronous processing, then the suspended backup volume is selected as the backup destination with the highest priority, and the synchronous processing restarts (ie, differential copy) starts. Otherwise, the synchronous processing (ie, full copy) starts.

- Backup synchronous processing with the output destination backup volume specified cannot be started in the following circumstances:

  - The backup volume specified in the device map file cannot be used as the backup volume to be paired with the transaction volume. This means that either of the following conditions is not satisfied:

    - The specified backup volume is in the unused state.

    - The transaction volume and backup volume have the same partition size.

  - The target transaction volume is incorrectly described in the device map file (Refer to "3.4.9.1 Describing a device map file").

  - Any special or foreign characters, operators or reserved words are used in the in the device map file names.

  - An ordinary volume is specified as the transaction volume and a GDS logical volume as the backup volume.

  - A GDS logical volume is specified as the transaction volume and an ordinary volume as the backup volume.

- For details on how to use SDX objects of GDS, refer to "Configuration and conditions of unsupported SDX objects" and "10.1.5.1 Notes on SDX object operations (backup management)".

- This command cannot be executed when any of the following commands are running:

  - swstsvrset (Managed Server configuration information setting command)

  - swstdevinfoset (Device information setting command)

  - swstsrsemtch (Resource match command)

  - swstresback (Resource backup command)

  - swstresrst (Resource restore command)

9.2.2.9 swstcancelsync (Backup synchronous processing cancel command)

NAME

swstcancelsync - cancels the backup synchronous processing

SYNOPSIS

Canceling backup synchronous processing by specifying the transaction volume

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstcancelsync [-emergency] [-bd Backup-Device-Name | -all] Device-Name
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstcancelsync [-h Server-Name] [-emergency] [-bd Backup-Device-Name | -all]
Device-Name
```

Canceling backup synchronous processing by specifying a log group

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstcancelsync [-emergency] -n Log-Group-Name[/RDB-NAME] [-all]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstcancelsync [-h Server-Name] [-emergency] -n Log-Group-Name[/RDB-NAME] [-
all]
```

DESCRIPTION

The backup synchronous processing can be cancelled during copying by EC or under the equivalency maintained status.

The synchronous processing can be cancelled with the backup volume during copying by EC, under the equivalency maintained status, or while the Suspend state is set. Specifying cancellation of all of the synchronous processing cancels all of the backup synchronous processing including those in the Suspended state.

If the system detects an error while backup synchronous processing is being executed, copying in progress at that time terminates abnormally. In this case, use this command to cancel the backup synchronous processing. Use swstsyncstat (Backup synchronous processing progress display command) to check for the error.

You can execute this command specifying individual transaction volumes that are included in a log group for which to cancel backup synchronous processing.

OPTIONS

-h *Server-Name*

Specify the name of a Managed Server.
This option can be specified only with a Management Server.

-emergency

Emergency operation mode is available for operations.
In such cases, EC sessions are neither confirmed nor cancelled.

-bd *Backup-Device-Name*

The synchronous processing on a specific backup volume can be cancelled.
Specify the copy destination backup volume name. In the operand, specify the AdvancedCopy Manager device name for the specified device.
Note that the backup volume name and the -all option cannot be specified together.

-all

All of the synchronous processing of the transaction volume can be specified to be cancelled. Note that this specification and the -bd option cannot be used together.

-n

Specifies that backup synchronous processing is being cancelled for a log group.
In the operand, specify the log group name and the RDB system name. If there is no RDB system name, you can omit it.

OPERANDS

Device-Name

Specify an AdvancedCopy Manager device name.
Refer to "1.5 Managing a Device on AdvancedCopy Manager" for the devices to be specified.

Log-Group-Name

Specify the name of the log group.

RDB-Name

Specify the RDB system name after the log group name, delimited by a slash ("/").
If there is no RDB system name, you can omit it.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Cancelling backup synchronous processing to a transaction volume:

```
# /opt/FJSVswsts/bin/swstcancelsync /dev/sdb6
/dev/sdb6 swstcancelsync completed.
#
```

- Cancelling backup synchronous processing of transaction volume in the emergency operation mode:

```
# /opt/FJSVswsts/bin/swstcancelsync -emergency /dev/sdb6
/dev/sdb6 swstcancelsync completed.
#
```

- To cancel backup synchronous processing for log group (LOG01/RDB1):

```
# /opt/FJSVswsts/bin/swstcancelsync -n LOG01/RDB1
LOG01/RDB1 swstcancelsync completed.
#
```

- Cancelling backup synchronous processing of the backup volume from the transaction volume:

```
# /opt/FJSVswsts/bin/swstsyncstat /dev/sdb6
Server Transaction-Disk Backup-Disk Status    Execute
Work2  /dev/sdb6        /dev/sdc6   suspend   ----
Work2  /dev/sdb6        /dev/sdd6   executing 75%
# /opt/FJSVswsts/bin/swstcancelsync /dev/sdb6 -bd /dev/sdd6
/dev/sdb6 swstcancelsync completed.
# /opt/FJSVswsts/bin/swstsyncstat /dev/sdb6
Server Transaction-Disk Backup-Disk Status  Execute
Work2  /dev/sdb6        /dev/sdc6   suspend ----
#
```

- Cancelling all of the backup synchronous processing of the transaction volume:

```
# /opt/FJSVswsts/bin/swstsyncstat /dev/sdb6
Server Transaction-Disk Backup-Disk Status    Execute
Work2  /dev/sdb6        /dev/sdc6   suspend   ----
Work2  /dev/sdb6        /dev/sdd6   executing 75%
# /opt/FJSVswsts/bin/swstcancelsync /dev/sdb6 -all
/dev/sdb6 swstcancelsync completed.
# /opt/FJSVswsts/bin/swstsyncstat /dev/sdb6
Server Transaction-Disk Backup-Disk Status Execute
Work2  /dev/sdb6        ----        ----   ----
#
```

NOTES

- However, EC sessions are neither confirmed nor cancelled by execution in the emergency operation mode. If an EC session is in the error suspended state or the halt state, then unmanaged sessions remain. In this event, ETERNUS Web GUI must be used to cancel these sessions.

- You cannot cancel synchronous backup processing in the following case:

    - The -n option was specified when GDS volumes and non-GDS volumes are mixed in a log group.

    - The replication management function is executing a process for the target transaction volume.

- For details on how to use SDX objects of GDS, refer to "Configuration and conditions of unsupported SDX objects" and "10.1.5.1 Notes on SDX object operations (backup management)".

- This command cannot be executed when any of the following commands are running:

    - swstsvrset (Managed Server configuration information setting command)

    - swstdevinfoset (Device information setting command)

    - swstsrsemtch (Resource match command)

    - swstresback (Resource backup command)

    - swstresrst (Resource restore command)

9.2.2.10 swstsyncstat (Backup synchronous processing progress display command)

NAME

swstsyncstat - displays the execution status of the backup synchronous processing

SYNOPSIS

Displaying the execution status of backup synchronous processing by specifying the transaction volume

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstsyncstat [Device-Name]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstsyncstat [-h Server-Name] [Device-Name]
```

Displaying the execution status of backup synchronous processing by specifying a log group

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstsyncstat -n Log-Group-Name[/RDB-Name]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstsyncstat -n [-h Server-Name] Log-Group-Name[/RDB-Name]
```

DESCRIPTION

The execution status of the backup synchronous processing is displayed (during copying, under the equivalency maintained status, or in the suspended status).
If one transaction volume has multiple EC sessions, then the statuses of all of these EC sessions are displayed.
The table below lists the meaning of each title.

Title	Description
Server	Displays the name of the Managed Server.
Transaction-Disk	Displays the AdvancedCopy Manager device name of a transaction volume.
Backup-Disk	Displays the name of the backup volume. When backup synchronous processing is not being executed, the system displays "----".
Status	Displays the status of backup synchronous processing: "----" : Backup synchronous processing is not being executed. "executing" : Backup synchronous processing is being executed but equivalency maintained status has not been established. "equivalent" : Equivalency maintained status is established. "suspend" : Suspend status is established. "failed" : Backup synchronous processing was abnormally terminated. "nosession" : No session exists (eg, due to inconsistent resource information). "gds-error" : Copying with GDS ended abnormally.
Execute	Displays the progress of the backup synchronous processing as a percentage (0 to 100%). "----" is displayed when the status is Suspend and at times when the backup synchronous processing is not running.

OPTIONS

-h *Server-Name*

Specify the name of a Managed Server.
This option can be specified only with the Management Server.

-n

Specifies that the execution status of backup synchronous processing being displayed is for a log group.

In the operand, specify the log group name and the RDB system name. If there is no RDB system name, you can omit it.

OPERANDS

Device-Name

Specify an AdvancedCopy Manager device name.

Refer to "1.5 Managing a Device on AdvancedCopy Manager" for the devices to be specified.

When you omit the device name, the system displays the status of backup synchronous processing for all of the transaction volumes.

Log-Group-Name

Specify the name of the log group.

RDB-Name

Specify the RDB system name after the log group name, delimited by a slash ("/").

If there is no RDB system name, you can omit it.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- To display the status of backup synchronous processing to a transaction volume:

```
# /opt/FJSVswsts/bin/swstsyncstat /dev/sdb6
Server Transaction-Disk Backup-Disk Status    Execute
Work1  /dev/sdb6        /dev/sdd5   suspend   ----
Work1  /dev/sdb6        /dev/sdd6   executing 75%
#
```

- To display the status of backup synchronous processing for all of the transaction volumes:

```
# /opt/FJSVswsts/bin/swstsyncstat
Server Transaction-Disk Backup-Disk Status     Execute
Work1  /dev/sdb3        /dev/sdd1   suspend    ----
Work1  /dev/sdb3        /dev/sdd3   executing  75%
Work1  /dev/sdb3        /dev/sdd4   suspend    ----
Work1  /dev/sdb5        /dev/sdd5   equivalent 100%
Work1  /dev/sdb6        /dev/sdd6   failed     ----
Work1  /dev/sdb4        ----        ----       ----
#
```

- To display information on the execution status of backup synchronous processing for log group (LOG01/RDB1):

```
# /opt/FJSVswsts/bin/swstsyncstat  -n LOG01/RDB1
Server Transaction-Disk  Backup-Disk Status    Execute
job2   /dev/sdd1         /dev/sdz1   executing 75%
job2   /dev/sdd2         /dev/sdz2   executing 10%
   :
#
```

NOTES

- If "failed" is indicated as the status of the backup synchronous processing (in the Status field), a hardware failure has possibly occurred.
  Refer to "8.3 Troubleshooting: Hardware error during backup" for more information about this.

- If "nosession" is indicated as the status of the backup synchronous processing (in the Status field), resource information having inconsistency is a possible cause. Use swstsrsemtch (Resource match command) to remove the inconsistency from the resource information.

- If "gds-error" is the status displayed in the Status field, a problem with GDS is assumed. Correct the cause of the GDS problem, and take action as follows:

  - If backup synchronous processing is in progress, execute swstcancelsync (Backup synchronous processing cancel command) to cancel the backup synchronous processing.

  - If snapshot backup is in progress, execute swsthistdel (History information deletion command).

  - If restoration is in progress, execute swstsrsemtch (Resource match command), and perform restoration again.

9.2.2.11 swstexecstat (Execution status display command)

NAME

swstexecstat - displays the execution status of other commands executed

SYNOPSIS

Displaying an execution status by specifying the transaction volume

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstexecstat [-j] [Device-Name]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstexecstat [-j] [-h Server-Name] [Device-Name]
```

Displaying an execution status by specifying a log group

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstexecstat -n [-j] Log-Group-Name[/RDB-NAME]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstexecstat -n [-h Server-Name] [-j] Log-Group-Name[/RDB-NAME]
```

DESCRIPTION

This command displays the execution status of other commands executed on a transaction volume. The execution status of OPC/EC will be displayed if OPC/EC is being executed.
The following information will be displayed:

Title	Description
Server	Displays a Managed Server name.
Device	Displays the AdvancedCopy Manager device name.
Last-Backup-Date	Displays the time when the last backup execution command was accepted.
Interval	Displays how many days have passed after the specified number of interval days since the last backup date has been exceeded. This item will be displayed as "DELAY=number-of-days-passing". "OK" will be displayed if the number of interval days has not been exceeded.
Status	Displays whether any of the following commands is being executed on a transaction volume. If so, "EXEC" will be displayed. Otherwise, "IDLE" will be displayed. - swstbkpolset (Backup policy setting command)

Title	Description
	- swstbkpoldel (Backup policy deletion command)
	- swstbackup (Backup execution command)
	- swstrestore (Restore execution command)
	- swsthistdel (History information deletion command)
	- swststartsync (Backup synchronous processing start command)
	- swstcancelsync (Backup synchronous processing cancel command)
	- swstcancelrest (Restore cancellation command)
	- swstcanceltrk (Tracking cancel command)
Mount-Point (Method)	Displays the mount point of a device. The file system type of the mount point is displayed inside the parentheses.
	If the "Device" is a transaction volume used for Symfoware, the names of Symfoware database spaces and log groups as well as the RDB system names are displayed, followed by "SymfoWARE" in parentheses.
Execute	Displays the name of the command being executed or the progress of copying.
	"----" : Neither a command nor OPC/EC is being executed. "swststartsync" : The backup synchronous processing start command is being processed. "swstcancelsync" : The backup synchronous processing cancel command is being processed. "swstbackup" : The backup execution command is being processed. "swstrestore" : The restore execution command is being processed. "swstcancelrest" : The restore cancellation command is being processed. "swstbkpolset" : The backup policy setting command is being processed. "swstbkpoldel" : The backup policy deletion command is being processed. "swsthistdel" : The history information deletion command is being processed. "swstcanceltrk" : The tracking cancel command is being processed. "sync(xxx%)" : The progress of backup synchronous processing by EC or the software-copy function of GDS. "sync(failed)" : EC ended abnormally. "snapshot(xxx%)" : The progress of OPC backup. "snapshot(failed)" : OPC backup ended abnormally. "restore(xxx%)" : The progress of restoration. "restore(failed)" : OPC restoration ended abnormally. "gds-error": GDS copying ended abnormally.
	During synchronous high-speed backup with -suspend specified, no operation ("----") is indicated in the Execute field.

If multiple processes are performed on transaction volumes (for example, when the backup synchronous processing start command is issued while a real copy of the snapshot backup is being handled), the display in the Execute field conforms to the following rules:

- The display of a command name takes priority if both copying and command executions are in progress.

- When a disk copy of the snapshot backup and a disk copy of the backup synchronous processing are being handled simultaneously, the progress of the disk copy of the backup synchronous processing is displayed. A disk copy for restoration is not performed during disk copy of the snapshot backup or backup synchronous processing due to a restriction of ETERNUS Disk storage system.

The following information will be displayed:

Title	Description
Server	Displays a Managed Server name.
Device	Displays the AdvancedCopy Manager device name.

Title	Description
Execute-Date	Displays the command execution date and time (TimeStamp).
Status	Displays whether any of the following commands are being executed on a transaction volume. If so, "EXEC" will be displayed. Otherwise, "IDLE" will be displayed. - swstbkpolset (Backup policy setting command) - swstbkpoldel (Backup policy deletion command) - swstbackup (Backup execution command) - swstrestore (Restore execution command) - swsthistdel (History information deletion command) - swststartsync (Backup synchronous processing start command) - swstcancelsync (Backup synchronous processing cancel command) - swstcancelrest (Restore cancellation command) - swstcanceltrk (Tracking cancel command)
Mount-Point (Method)	Displays the mount point of a device. The file system type of the mount point is displayed inside the parentheses. If the "Device" is a transaction volume used for Symfoware, the names of Symfoware database spaces and log groups as well as the RDB system names are displayed, followed by "SymfoWARE" in parentheses.
Execute	Displays the name of the command being executed or the progress of copying. "----" : Neither a command nor OPC/EC is being executed. "swststartsync" : The backup synchronous processing start command is being processed. "swstcancelsync" : The backup synchronous processing cancel command is being processed. "swstbackup" : The backup execution command is being processed. "swstrestore" : The restore execution command is being processed. "swstcancelrest" : The restore cancellation command is being processed. "swstbkpolset" : The backup policy setting command is being processed. "swstbkpoldel" : The backup policy deletion command is being processed. "swsthistdel" : The history information deletion command is being processed. "swstcanceltrk" : The tracking cancel command is being processed. "sync(xxx%)" : The progress of backup synchronous processing by EC or the software-copy function of GDS. "sync(failed)" : EC ended abnormally. "snapshot(xxx%)" : The progress of OPC backup. "snapshot(failed)" : OPC backup ended abnormally. "restore(xxx%)" : The progress of restoration. "restore(failed)" : OPC restoration ended abnormally. "gds-error": GDS copying ended abnormally. During synchronous high-speed backup with -suspend specified, no operation ("----") is indicated in the Execute field.

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on a Management Server.

-j

Specifies to display information in another format.

-n

> Specifies that the command execution status being displayed is for a log group.
> In the operand, specify the log group name and the RDB system name. If there is no RDB system name, you can omit it.

OPERANDS

Device-Name

> Specifies the AdvancedCopy Manager device name corresponding to a transaction volume.
> If this operand is omitted, the execution statuses of all the transaction volumes will be displayed.

Log-Group-Name

> Specify the name of the log group.

RDB-Name

> Specify the RDB system name after the log group name, delimited by a slash ("/").
> If there is no RDB system name, you can omit it.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Display the execution statuses of transaction volumes:

```
# /opt/FJSVswsts/bin/swstexecstat
Server Device      Last-Backup-Date Interval Status Mount-Point (Method) Execute
Work1  /dev/sdb3 2000/12/11 12:20 OK        IDLE   /usr1 (ext2)          ----
Work1  /dev/sdb4 2000/12/10 12:20 DELAY=1   IDLE   /usr2 (ext2)          sync(22%)
Work1  /dev/sdb5 2000/12/09 12:20 DELAY=2   IDLE   /usr3 (ext2)          snapshot(45%)
:
#
```

- An example of using the -j option is as follows:

```
# /opt/FJSVswsts/bin/swstexecstat -j
Server Device      Execute-Date     Status Mount-Point (Method) Execute
Work1  /dev/sdb3 2000/12/11 12:20 IDLE   /usr1 (ext2)         ----
Work1  /dev/sdb4 2000/12/10 12:20 IDLE   /usr2 (ext2)         sync(22%)
Work1  /dev/sdb5 2000/12/09 12:20 IDLE   /usr3 (ext2)         snapshot(45%)
:
#
```

NOTES

- If "sync (failed)", "snapshot(failed)", or "restore(failed)" is displayed in the **Execute** field, a hardware error may have occurred. After fixing the cause of the hardware error, perform the following:

    - For "sync (failed)":
      Execute swstcancelsync (Backup synchronous processing cancel command) to cancel the backup synchronous processing.

    - If "snapshot(failed)" is displayed:
      Delete backup history information using swsthistdel (History information deletion command).

    - If "restore(failed)" is displayed:
      Cancel restoration using swstcancelrest (Restore cancellation command).

- If "gds-error" is the status displayed in the Status field, a problem with GDS is assumed. Correct the cause of the GDS problem, and take action as follows:

    - If backup synchronous processing is in progress, execute swstcancelsync (Backup synchronous processing cancel command) to cancel the backup synchronous processing.

    - If snapshot backup is in progress, execute swsthistdel (History information deletion command).

    - If restoration is in progress, execute swstsrsemtch (Resource match command), and perform restoration again.

9.2.2.12 swstcanceltrk (Tracking cancel command)

NAME

swstcanceltrk - stops tracking processing

SYNOPSIS

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstcanceltrk [-copy | -emergency] Device-Name (-bd Backup-Device-Name | -all)
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstcanceltrk -h Server-Name [-copy | -emergency] Device-Name (-bd Backup-
Device-Name | -all)
```

DESCRIPTION

This command stops tracking processing.

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on the Management Server.

-copy

Specifies that physical copying, if active, should also be stopped (the OPC session should be cancelled).
This option cannot be specified for a backup volume that has already contained a backup history. For a volume whose copy destination volume has contained a backup history, delete the history information in advance using swsthistdel (History information deletion command).

-emergency

Specifies that the command should operate in the emergency operation mode.
The command does not check tracking processing and does not stop tracking processing or the OPC session.

-bd *Backup-Device-Name*

Stops tracking processing for the specified backup volume. Specify the target backup volume name as an operand. This option cannot be specified at the same time as the -all option.

-all

Stops the entire tracking processing specified for the transaction volumes. This option cannot be specified at the same time as the -bd option.

OPERANDS

Device-Name

Specify an AdvancedCopy Manager device name.
Refer to "1.5 Managing a Device on AdvancedCopy Manager" for the devices to be specified.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- The following example stops the tracking processing that was specified for the backup volume from the transaction volume:

```
# /opt/FJSVswsts/bin/swstcanceltrk /dev/sdb6 -bd /dev/sdd6
/dev/sdb6 swstcanceltrk completed.
#
```

- The following example stops tracking processing and physical copying that were specified for the transaction volume:

```
# /opt/FJSVswsts/bin/swsttrkstat /dev/sdb6
Server Transaction-Disk Backup-Disk Status    Execute Update
job2   /dev/sdb6        /dev/sdd6   executing 33%     -----
# /opt/FJSVswsts/bin/swstcanceltrk -copy /dev/sdb6 -bd /dev/sdd6
/dev/sdb6 swstcanceltrk completed.
#
```

NOTES

- If the command is executed in emergency operation mode, no OPC session is cancelled. A session that is not managed remains in operation if the command is entered during tracking. The session must be cancelled using ETERNUS Web GUI or GDS.

- Tracking processing cannot be stopped under the following condition. In this case, take action according to "System administrator response" provided for the message output.

  - When tracking processing is not performed for the specified copy source and destination volumes.

- This command cannot be executed while one of the following commands is being executed:

  - swstsvrset (Managed Server configuration information setting command)

  - swstdevinfoset (Device information setting command)

  - swstsrsemtch (Resource match command)

  - swstresback (Resource backup command)

  - swstresrst (Resource restore command)

9.2.2.13 swsttrkstat (Tracking status display command)

NAME

swsttrkstat - displays the status of tracking processing

SYNOPSIS

Displaying the execution status of tracking by specifying the transaction volume

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swsttrkstat [Device-Name]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swsttrkstat [-h Server-Name] [Device-Name]
```

Displaying the execution status of tracking by specifying a log group

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swsttrkstat -n Log-Group-Name[/RDB-Name]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swsttrkstat -n [-h Server-Name] Log-Group-Name[/RDB-Name]
```

DESCRIPTION

This command displays the status of tracking processing.
When multiple types of synchronous processing are running on one transaction volume, the command displays the statuses of all types of tracking processing.
The displayed data is explained below.

Title	Description
Server	Displays a Managed Server name.
Transaction-Disk	Displays the AdvancedCopy Manager device name of a transaction volume.
Backup-Disk	Displays a backup volume name. Displays "----" when the volume is not in the tracking status.
Status	Displays the execution status. "----" : Not in the process of tracking "executing" : In the process of physical copying and tracking "tracking" : In the process of tracking "failed" : Physical copying or tracking terminated abnormally "nosession" : No session exists. (eg, resource information is inconsistent.)
Execute	Displays the progress of copying as a percentage when "executing" is displayed in the Status column. "----" is displayed when "tracking" is displayed in the Status column.
Update	Displays the amount of updated data after logical copying as a percentage when "tracking" is displayed in the Status column. "----" is displayed when "executing" is displayed in the Status column.

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on the Management Server.

-n

Specifies that the execution status of tracking being displayed is for a log group.
In the operand, specify the log group name and the RDB system name. If there is no RDB system name, you can omit it.

OPERANDS

Device-Name

> Specify an AdvancedCopy Manager device name.
> Refer to "1.5 Managing a Device on AdvancedCopy Manager" for the devices to be specified.
> If this operand is omitted, the command displays the statuses of tracking processing for all transaction volumes.

Log-Group-Name

> Specify the name of the log group.

RDB-Name

> Specify the RDB system name after the log group name, delimited by a slash ("/").
> If there is no RDB system name, you can omit it.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- The following example displays the status of tracking processing of a transaction volume:

```
# /opt/FJSVswsts/bin/swsttrkstat /dev/sdb6
Server Transaction-Disk Backup-Disk Status    Execute Update
job1   /dev/sdb6        /dev/sdd6   tracking ----    12%
#
```

- The following example displays the status of tracking processing of all transaction volumes:

```
# /opt/FJSVswsts/bin/swsttrkstat
Server Transaction-Disk Backup-Disk Status    Execute Update
job1   /dev/sdb3        /dev/sdd3   nosession ----   ----
job1   /dev/sdb4        /dev/sdd4   executing 75%    ----
job1   /dev/sdb5        /dev/sdd5   failed    ----   ----
job1   /dev/sdb6        /dev/sdd6   tracking  ----   12%
job1   /dev/sdc1        ----        ----      ----   ----
   :
#
```

- To display the execution status of tracking for log group (LOG01/RDB1):

```
# /opt/FJSVswsts/bin/swsttrkstat  -n LOG01/RDB1
Server Transaction-Disk Backup-Disk Status    Execute Update
job2   /dev/sdc4        /dev/sde4   tracking ----    12%
job2   /dev/sdc5        /dev/sde5   tracking ----    15%
   :
#
```

NOTES

- When the execution status (ie, the value displayed in the "Status" column) is "failed," a hardware problem has occurred Fix cause of the problem and then re-execute the differential snapshot backup.

- When the execution status (ie, the value displayed in the "Status" column) is "nosession," the resource information is probably inconsistent. Use swstsrsemtch (Resource match command) to remove any inconsistencies from the resource information.

9.2.3 Maintenance commands

This section describes the backup management maintenance commands.

9.2.3.1 swstsrsemtch (Resource match command)

NAME

swstsrsemtch - recovers resource information

SYNOPSIS

To be executed on the Managed Server

```
/opt/FJSVswsts/bin/swstsrsemtch [-x]
```

To be executed on the Management Server

```
/opt/FJSVswsts/bin/swstsrsemtch [-x] [-h Server-Name]
```

DESCRIPTION

This command recovers resource information if a system failure, etc. has made the resource information inconsistent.

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on a Management Server.

-x

Achieves consistency with a repository. Specify this option if a backup management files and a repository become inconsistent.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Recover resource information:

```
# /opt/FJSVswsts/bin/swstsrsemtch
swstsrsemtch completed
#
```

NOTES

- This command can only be executed if none of the following commands are running:

  - swstsvrset (Managed Server configuration information setting command)

  - swstbkpolset (Backup policy setting command)

  - swstbkpoldel (Backup policy deletion command)

9.2.3.2 swstresback (Resource backup command)

NAME

swstresback - backup of backup management file and recovery control file

SYNOPSIS

```
/opt/FJSVswsts/bin/swstresback Backup-Directory
```

DESCRIPTION

A backup management file and a recovery control file are backed up to the specified directory.

- Backup management file
  Management information, such as Managed Server configuration information

- Recovery control file
  File that was created when backing up Symfoware database spaces

The information backed up by this command can be restored by swstresrst (Resource restore command).

OPERANDS

Backup-Directory

Specifies the name of a directory to which the backup management file and the recovery control file will be saved.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Backup of the backup management file and the recovery control file:

```
# /opt/FJSVswsts/bin/swstresback /home/backup
swstresback completed
#
```

NOTES

- You cannot perform backup if:

    - The specified directory does not exist.

    - Copying of data to the specified directory failed because of insufficient free space, etc.

- This command can be executed only while none of the following commands are running:

    - swstsvrset (Managed Server configuration information setting command)

    - swstbkpolset (Backup policy setting command)

    - swstbkpoldel (Backup policy deletion command)

    - swstbackup (Backup execution command)

    - swstrestore (Restore execution command)

    - swstcancelrest (Restore cancellation command)

    - swststartsync (Backup synchronous processing start command)

    - swstcancelsync (Backup synchronous processing cancel command)

    - swstdevinfoset (Device information setting command)

    - swsthistdel (History information deletion command)

    - swstcanceltrk (Tracking cancel command)

    - swstresback (Resource backup command)

    - swstresrst (Resource restore command)

    - swstsrsemtch (Resource match command)

9.2.3.3 swstresrst (Resource restore command)

NAME

swstresrst - restore of backup management file and recovery control file

SYNOPSIS

```
/opt/FJSVswsts/bin/swstresrst Backup-Directory
```

DESCRIPTION

A backup management file and a recovery control file are restored from the specified directory.

- Managed resources information
  Management information, such as Managed Server configuration information

- Recovery control file
  File that was created when backing up Symfoware database spaces

OPERANDS

Backup-Directory

The backup location directory name directed by swstresback (Resource backup command) is specified.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Restore the backup management file and the recovery control file:

```
# /opt/FJSVswsts/bin/swstresrst /home/backup
swstresrst completed
#
```

NOTES

- Before the backup management file and the recovery control file is restored by this command, swstsvrset (Managed Server configuration information setting command) must be executed. For the name of the target directory for output of the recovery control file that is specified by the -c option of swstsvrset (Managed Server configuration information setting command), specify the directory that was set when you performed the resource backup.

- You cannot perform restoration if:

    - The specified directory does not exist.

    - No backup management files exist in the specified directory.

    - Copying of data from the specified directory for restoration failed because of insufficient free space, or for some other reason.

- This command can be executed only while none of the following commands are running:

    - swstsvrset (Managed Server configuration information setting command)

    - swstbkpolset (Backup policy setting command)

    - swstbkpoldel (Backup policy deletion command)

    - swstbackup (Backup execution command)

    - swstrestore (Restore execution command)

    - swstcancelrest (Restore cancellation command)

    - swststartsync (Backup synchronous processing start command)

    - swstcancelsync (Backup synchronous processing cancel command)

    - swstdevinfoset (Device information setting command)

    - swsthistdel (History information deletion command)

    - swstcanceltrk (Tracking cancel command)

    - swstresback (Resource backup command)

    - swstresrst (Resource restore command)

    - swstsrsemtch (Resource match command)

9.3 Configuration Management Commands

This section explains the operation of the configuration management command. This manages the information on a server device which performs backup and replication, or is a candidate for being used for backup or replication.

 Note

....

Operating environment for command execution

The table below lists the requirements for the operating environment in which to execute the different commands.

Table 9.14 Operating environment for executing configuration management commands

Command	Operating environment of AdvancedCopy Manager
stgcmmodnode	To execute this command on a Management Server (Management Server transaction), the Management Server (Management Server transaction) must be running.
stguserset	In cluster operation, the Management Server transaction must be running.
stgxfwcmaddsrv	The Management Server (Management Server transaction) and target Managed Server (Managed Server transaction) must be running.
stgxfwcmmodsrv	The Management Server (Management Server transaction) and target Managed Server (Managed Server transaction) must be running.
stgxfwcmdelsrv	The Management Server (Management Server transaction) must be running.
stgxfwcmsetdev	The Management Server (Management Server transaction) and target Managed Server (Managed Server transaction) must be running.
stgxfwcmdeldev	The Management Server (Management Server transaction) and target Managed Server (Managed Server transaction) must be running.
stgxfwcmdispsrv	The Management Server (Management Server transaction) must be running.
stgxfwcmdispdev	The Management Server (Management Server transaction) must be running.
stgxfwcmdisppat	The Management Server (Management Server transaction) must be running.
stgenvdisp	This command is executable when both Management Server (Management Server transaction) and Managed Server (Managed Server transaction) are not running. Note that when the cluster service (named differently in respective clustered systems) is not running in cluster operation, this command cannot display certain information.
stgmgrinfoset	The package-backup-of-the-management-information execution server must be running.
stgmgrinfobkup	The Management Server (Management Server transaction), package-backup-of-the-management-information execution server, and package-backup-of-the-management-information target servers must be running.
stgmgrinfodisp	The package-backup-of-the-management-information execution server must be running.
stgxfwcmsetsdx	Managed Server transactions must be running for cluster operations.
stgxfwcmsetmode	Managed Server transactions must be running for cluster operations.
stgdbloginf	This command is executable when both Management Server (Management Server transaction) and Managed Server (Managed Server transaction) are not running.
stgdbdmp	This command is executable when both Management Server (Management Server transaction) and Managed Server (Managed Server transaction) are not running.
stgdbrcv	This command is executable when both Management Server (Management Server transaction) and Managed Server (Managed Server transaction) are not running.
stgdbbackup	The Management Server (Management Server transaction) must be running.
stgdbrestore	The Management Server (Management Server transaction) must be running.
stgdbset	This command is executable when both Management Server (Management Server transaction) and Managed Server (Managed Server transaction) are not running.
stgdbunset	This command is executable when both Management Server (Management Server transaction) and Managed Server (Managed Server transaction) are not running.

Command	Operating environment of AdvancedCopy Manager
stgrepocnv	This command is executable when both Management Server (Management Server transaction) and Managed Server (Managed Server transaction) are not running.
stgforcedelsrv	The Management Server (Management Server transaction) must be running. This command is executable when target Managed Server is not running.

Note

In cluster operation

In cluster operation, execute a command by taking the following steps:

1. If necessary, specify the logical node name to the transaction in the environment variable SWSTGNODE. Specify the environment variable as follows:

```
# SWSTGNODE=logical-node-name
# export SWSTGNODE
```

2. Execute the command.

The table below lists the commands available in cluster operation, indicates whether environment variables need to be configured, and shows the nodes on which commands are executable. For notes and cautions on command execution, refer to the Remarks column.

Table 9.15 Commands when cluster operation is used

Command	Environment variable	Command execution node	Remarks
stgcmmodnode	Required	Transaction operating node	-
stguserset	Not required	Transaction operating node	Executable only on Management Server transaction
stgxfwcmaddsrv	Not required	Transaction operating node	Executable only on Management Server transaction
stgxfwcmmodsrv	Not required	Transaction operating node	Executable only on Management Server transaction
stgxfwcmdelsrv	Not required	Transaction operating node	Executable only on Management Server transaction
stgxfwcmsetdev	Not required	Transaction operating node	Executable only on Management Server transaction
stgxfwcmdeldev	Not required	Transaction operating node	Executable only on Management Server transaction
stgxfwcmdispsrv	Not required	Transaction operating node	Executable only on Management Server transaction
stgxfwcmdispdev	Not required	Transaction operating node	Executable only on Management Server transaction
stgxfwcmdisppat	Not required	Transaction operating node	Executable only on Management Server transaction
stgenvdisp	Not required	Transaction operating node	-
stgmgrinfobkup	Required	Transaction operating node	-
stgmgrinfodisp	Required	Transaction operating node	-
stgmgrinfoset	Required	Transaction operating node	-
stgxfwcmsetsdx	Required	Transaction operating node	-

Command	Environment variable	Command execution node	Remarks
stgxfwcmsetmode	Required	Transaction operating node	-
stgdbloginf	Not required	Transaction operating node	Executable only on Management Server transaction
stgdbdmp	Not required	Transaction operating node	Executable only on Management Server transaction
stgdbrcv	Not required	Transaction operating node	Executable only on Management Server transaction
stgdbbackup	Not required	Transaction operating node	Executable only on Management Server transaction
stgdbrestore	Not required	Transaction operating node	Executable only on Management Server transaction
stgdbset	Not required	Transaction operating node	Executable only on Management Server transaction
stgdbunset	Not required	Transaction operating node	Executable only on Management Server transaction
stgrepocnv	Not required	Transaction operating node	Executable only on Management Server transaction
stgforcedelsrv	Not required	Transaction operating node	Executable only on Management Server transaction

9.3.1 stgcmmodnode (Management Server information change command)

NAME

stgcmmodnode - changes or displays Management Server information

SYNOPSIS

```
/opt/FJSVswstf/bin/stgcmmodnode [-i IP-Address] [-p PORT-Number] [-n Server-Name]
```

DESCRIPTION

This command has the following two functions:

- Displays Management Server information specified on the Managed Server on which this command is currently executed

- Changes Management Server information.

This command changes the following information on a Management Server:

- IP address of a Management Server

- PORT number used by a communication daemon that runs on a Management Server

- Server name of a Management Server

Refer to the following for the change procedure which uses this command:

- "8.5.1 Changing the IP address of Management Server"

- "8.5.3 Changing the port number of Management Server or Managed Server"

OPTIONS

None

Displays the server information currently being managed.

-i *IP-Address*

Specifies the IP address of a Management Server, to change the IP address of the Management Server.

-p *PORT-Number*

Specifies the port number of a Management Server, to change the port number of the Management Server.

-n *Server-Name*

Specifies the server name of a Management Server to change the server name of the Management Server. This option cannot be specified if this command is executed on a Management Server.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Display server information:

```
# /opt/FJSVswstf/bin/stgcmmodnode
swstf2413 IP Address = 10.10.10.10
swstf2413 PORT Number = 1226
swstf2413 Storage Management Server's Name = fujitsu-storage-management-server
#
```

- Change the IP address of a Management Server to 10.12.13.14:

```
# /opt/FJSVswstf/bin/stgcmmodnode -i 10.12.13.14
swstf2401 Storage Management Server registry updated.
#
```

NOTES

- This command can be executed only by root users.

- For the port number of a Management Server, you can specify a value between 1024 and 65535.

- With this command, you must specify the same value as the one specified on a Management Server. An error will occur in the operation if an invalid value is specified.

- To execute this command, the AdvancedCopy Manager daemons must be stopped.
  After successful execution, you must restart the AdvancedCopy Manager daemons.

9.3.2 stguserset (Repository access user change command)

NAME

stguserset - changes the user name and the password

SYNOPSIS

```
/opt/swstorage/bin/stguserset User-Name Password
```

DESCRIPTION

This command changes the user name and the password of a "repository access user" specified when the ETERNUS SF Manager was installed. This command must be executed on a Management Server.

The change in the username and the password is executed according to the following procedures:

1. Close the Web Console.

2. Execute Repository access user change command (stguserset).

3. Stop RMI daemons.

4. Start RMI daemons.

OPERANDS

User-Name

Specifies a user name to be used for repository access.
Specify a username, of up to 18 bytes, that is specifiable under the operating system.

Password

Specifies a user password to be used for repository access.
Specify a username, of up to 18 bytes, that is specifiable under the operating system.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Change the user name and the password to be used for repository access to a1234 and Yfd1fz3, respectively:

```
# /opt/swstorage/bin/stguserset a1234 Yfd1fz3
Password for accessing the repository has been registered.
#
```

NOTES

- This command can be executed only by root users.

- We recommend that you use a high root or superuser instead of the general user frequently used for password changes for security reasons.

9.3.3 stgxfwcmaddsrv (Server information addition command)

NAME

stgxfwcmaddsrv - adds Managed Server information

SYNOPSIS

```
/opt/FJSVswstf/bin/stgxfwcmaddsrv [-n Server-Name] -i IP-Address -p Port-Number
```

DESCRIPTION

This command adds Managed Server information to be managed by AdvancedCopy Manager. This command can only be executed from the Management Server.

OPTIONS

-n *Server-Name*

Specifies the name of the Managed Server to be added. The name must be less than 255 bytes and must be a unique server name so that AdvancedCopy Manager can manage the server. The usable characters are 1-byte alphanumeric characters, minus sign (-), underscore (_), number sign (#), and period (.), with the first character being a 1-byte alphanumeric character. A 2-byte character causes an error.

Note that the server name is case-sensitive. When specifying the server name in other commands or Web Console, enter the server name exactly as entered in this command, including the case.

Non-cluster operation: Host name must be recognized by the operating system
Cluster operation: AdvancedCopy Manager logical node name must be specified during cluster setup

-i *IP-Address*

Specify the IP address of the Managed Server to be added in IPv4 format.
Specify a physical IP address if the Managed Server is managed as a physical server. Specify a logical IP address if the IP address is made redundant and managed (as a logical server), the same as under the cluster environment.

-p *Port-Number*

Specify the port number that is set on the Managed Server side to communicate with the Managed Server to be added. (Value from 1,024 to 65,535)
Under the cluster environment, specify the port number that is allocated to the service name stgxfws_? (Logical node name of AdvancedCopy Manager specified during cluster setup).

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Add a server by using the host name:

```
# /opt/FJSVswstf/bin/stgxfwcmaddsrv -i 10.10.10.10 -p 1226
stgxfwcmaddsrv completed
#
```

- Add a logical server so that the cluster setup is performed by specifying a server name:
  (Logical IP address: 10.10.10.20; port number allocated to stgxfws_logical-node-name: 1227)

```
# /opt/FJSVswstf/bin/ stgxfwcmaddsrv -n Server1 -i 10.10.10.20 -p 1227
stgxfwcmaddsrv completed
#
```

NOTES

- This command can be executed only by root users.

- Please perform the command in the location specified at the time of product installation.

- Do not specify broadcast IP address to the IP address of the Managed Server to be added.

- Do not specify a server which another Management Server is managing as a Managed Server.

- Because a Management Server and a logical server for a Management Server transaction are added during database setup and cluster setup, they are not added by this command.
  When processing takes a long time, a period (.) is displayed until it is complete.

- This command cannot be executed while the following commands are being executed:

    - stgxfwcmaddsrv (Server information addition command)

    - stgxfwcmmodsrv (Server information change command)

    - stgxfwcmdelsrv (Server information deletion command)

    - stgxfwcmsetdev (Device information collection/reflection command)

    - stgxfwcmdeldev (Device information deletion command)

- This command cannot be executed while the following operations are being performed in the initial window by using the Web Console:

    - Addition of a new server

    - Change of a server

    - Deletion of a server

    - Acquisition and Reflection of Information on All Devices

    - Acquisition and Reflection of Information on a Specific Device

    - Deletion of device information

9.3.4 stgxfwcmmodsrv (Server information change command)

NAME

stgxfwcmmodsrv - changes the following server information

SYNOPSIS

```
/opt/FJSVswstf/bin/stgxfwcmmodsrv -n Server-Name [-s New-Server-Name] [-i New-IP-Address] [-p New-
port-Number] | [-f server-information-change-instruction-filename [-y]]
```

DESCRIPTION

This command changes the following server information managed by AdvancedCopy Manager. It is necessary to execute this command on the Management Server.

- A Managed Server IP address

- A Managed Server port number

- The server name of a Management Server

- The server name of a Managed Server

Server information can also be changed as a batch while maintaining the consistency of multiple items of information between multiple servers in the managed system by specifying the server information change instruction file.

The following server information can be changed:

- Management Server IP address

- Managed Server IP address

- Management Server port number

- Managed Server port number

- Server name of the Management Server

- Server names of Managed Servers

OPTIONS

-n *Server-Name*

Specify a server name that is currently being managed.
Specify a server name displayed by stgxfwcmdispsrv (Server information display command) exactly as displayed, including the case.

-s *New-Server-Name*

Specify a new server name in less than 255 bytes. The usable characters are: 1-byte alphanumeric characters, minus sign (-), underscore (_), number sign (#), and period (.), with the first character being a 1-byte alphanumeric character. (A 2-byte character causes an error.)
Note that the new server name is case-sensitive. When specifying the server name in other commands or Web Console, enter the server name exactly as entered in this command, including the case.
Specify at the same time as the -n option.

-i *New-IP-Address*

Specify the IP address of the Managed Server to be added by IPv4 format.
Specify a physical IP address if the Managed Server is managed as a physical server. Specify a logical IP address if the IP address is made redundant and managed (as a logical server), as in a cluster environment.
Specify at the same time as the -n option.

-p *New-port-Number*

Specify a port number required for post-change communication. The values are from 1024 to 65535)
Under the cluster environment, specify the port number that is allocated to service name stgxfws_? (Logical node name of AdvancedCopy Manager specified during cluster setup)
For non-cluster operations, specify the port number assigned to service name stgxfws.
Specify at the same time as the -n option.

-f *server-information-change-instruction-filename*

Specify the absolute path of the server information change instruction filename.
This option can only be specified at the same time as than the -y option.
Format for the server information change instruction file

- In the first line, enter the server name of a currently managed server, the information class to be changed, and the new value, delimited by spaces or tabs.

- The information classes to be changed are as follows:
  S: Server name
  I: IP address
  P: Port number

- If multiple information classes, such as the IP address and the port number, need to be changed at the same server, enter the IP address change instruction, then enter the port number change instruction on the second line.

```
Manager    I  10.10.20.10
Server1    I  10.10.20.20
Server1    P  9888
Server2    S  DstServer
```

-y

Specify this option to continue processing if the management system contains multiple servers and an error occurs at a specific server. Remove the cause of the error at the server where the error occurred, and perform the required change procedure separately for that

server.
Specify at the same time as the -f option.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- The IP address of a Managed Server (with server name: Server1) is changed to 10.20.30.40:

```
# /opt/FJSVswstf/bin/stgxfwcmmodsrv -n Server1 -i 10.20.30.40
stgxfwcmmodsrv completed
#
```

- The server name of a Managed Server (with server name: Server1) is changed to Agent5:

```
# /opt/FJSVswstf/bin/stgxfwcmmodsrv -n Server1 -s Agent5
stgxfwcmmodsrv completed
#
```

- The server information change instruction file is specified to change server information as a batch:

```
# /opt/FJSVswstf/bin/stgxfwcmmodsrv -f /acm/modfile
stgxfwcmmodsrv completed
#
```

- The server information change instruction file is specified to change server information as a batch, even if an error occurs at a specific server:

```
# /opt/FJSVswstf/bin/stgxfwcmmodsrv -f /acm/modfile -y

swstf2459 Resource adjustment command terminated abnormally. Server name=[Server1]
swsrp0625 This command cannot be executed because execution of other AdvancedCopy Manager commands
is in progress.

stgxfwcmmodsrv completed, excluding the server where the error occurs
#
```

NOTES

- Since stgcmmodnode (Management Server information change command) is used to change the following server information, this information can only be changed if the server information change instruction file is specified:

    - The IP address of a Management Server

    - The port number of a Management Server

- If the server information is going to be changed by using this command, the machine definition information must be changed in advance. Refer to the procedure given in "8.5 Changing the Operating Environment", for information on executing the processing.

- This command can be executed only by root users.

- Please perform the command in the location specified at the time of product installation.

- When processing takes a long time, a period (.) is displayed until it is complete.

- The following notes apply to changing the Management Server information (except when the -f option is specified):

    - Use stgcmmodnode (Management Server information change command) to change the IP address and port number, then restart the various AdvancedCopy Manager services.

- After Management Server information is changed, execute stgcmmodnode (Management Server information change command) at the Managed Servers to set the changed Management Server information.

- The following notes apply to changing the Managed Server information:

  - Stop the various AdvancedCopy Manager services before changing the IP address. Then, change the IP address in the software that provides the system or logical IP address and restart the AdvancedCopy Manager daemon.

  - Stop the various AdvancedCopy Manager services before changing the port number. Then, change the port number in the system and restart the AdvancedCopy Manager daemon.

- The following notes apply when the -f option is specified:

  - The server name or IP address that is to be changed cannot be in use by another server that is currently managed.

  - The server name to be changed is case-sensitive. Specify a new server name displayed by stgxfwcmdispsrv (Server information display command) exactly as displayed, including the case.

  - The new server name is case-sensitive. When specifying the server name in other commands or Web Console, enter the server name exactly as entered in this command, including the case.

  - The Management Server that executes this command and the AdvancedCopy Manager services and Managed Server transactions at all Managed Servers managed by Management Server transactions must be running.

  - Backup management commands and replication management commands must not be executed at the Management Server that executes this command nor in the AdvancedCopy Manager services and Managed Server transactions at all Managed Servers managed by Management Server transactions.
  If commands are executed, they may terminate abnormally. If the server information change command terminates abnormally, perform recovery as indicated in the error message. If backup management or replication management commands are accidentally executed during execution of the server information change command, and if the backup management and replication management commands terminate abnormally, re-execute the backup management or replication management commands after normal termination of the server information change command.

- The following notes apply when the -y option is specified:

  - If this option is specified, processing will be forcibly continued if a communication or other error occurs for some servers. Therefore, consistency of the overall system cannot be maintained. Usually, do not specify this option.
  Specify this option only when absolutely necessary due to problems at some servers.

  - Take the action indicated in the error message at the server which output the message. Normal operation is not possible until this action is taken, so take action promptly.

- This command cannot be executed while the following commands are being executed:

  - stgxfwcmaddsrv (Server information addition command)

  - stgxfwcmmodsrv (Server information change command)

  - stgxfwcmdelsrv (Server information deletion command)

  - stgxfwcmsetdev (Device information collection/reflection command)

  - stgxfwcmdeldev (Device information deletion command)

- This command cannot be executed while the following operations are being performed in the initial window by using the Web Console:

  - Addition of a new server

  - Change of a server

  - Deletion of a server

  - Acquisition and Reflection of Information on All Devices

  - Acquisition and Reflection of Information on a Specific Device

  - Deletion of device information

9.3.5 stgxfwcmdelsrv (Server information deletion command)

NAME

stgxfwcmdelsrv - deletes the Managed Server information

SYNOPSIS

```
/opt/FJSVswstf/bin/stgxfwcmdelsrv -n Server-Name
```

DESCRIPTION

This command deletes the Managed Server information that AdvancedCopy Manager is currently managing. It is necessary to execute this command on the Management Server.

OPTIONS

-n *Server-Name*

Specify the Managed Server name to be deleted.

Specify a server name displayed by stgxfwcmdispsrv (Server information display command) exactly as displayed, including the case.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- A Managed Server (server name: Server1) information is deleted:

```
# /opt/FJSVswstf/bin/stgxfwcmdelsrv -n Server1
stgxfwcmdelsrv completed
#
```

NOTES

- This command can be executed only by root users.

- Perform this command in the location specified at the time of product installation.

- The Management Server information cannot be deleted.

- When processing takes a long time, a period (.) is displayed until it is complete.

- This command cannot be executed while the following commands are being executed:

    - stgxfwcmaddsrv (Server information addition command)

    - stgxfwcmmodsrv (Server information change command)

    - stgxfwcmdelsrv (Server information deletion command)

    - stgxfwcmsetdev (Device information collection/reflection command)

    - stgxfwcmdeldev (Device information deletion command)

- This command cannot be executed when the following operations are being performed in the initial window by using the Web Console:

    - Addition of a new server

    - Change of a server

    - Deletion of a server

    - Acquisition and Reflection of Information on All Devices

    - Acquisition and Reflection of Information on a Specific Device

    - Deletion of device information

9.3.6 stgxfwcmsetdev (Device information collection/reflection command)

NAME

stgxfwcmsetdev - collects and reflects the device information

SYNOPSIS

```
/opt/FJSVswstf/bin/stgxfwcmsetdev -n Server-Name [-d Device-Name] [-f File-Name | -y]
```

DESCRIPTION

This command collects the latest device information detected by the server and displays the difference between this and the information managed by AdvancedCopy Manager. After this, it displays the specified device difference information in the management information for AdvancedCopy Manager. It is necessary to execute this command on the Management Server.

The device information currently being managed and the device information detected are compared, stored, deleted, updated, and used in the following cases:

  - When operation is started using AdvancedCopy Manager

  - When device information being managed is changed

OPTIONS

-n *Server-Name*

Specify the server name whose device information is to be collected or reflected.
Specify a server name displayed by stgxfwcmdispsrv (Server information display command) exactly as displayed, including the case.

-d *Device-Name*

Specify the device name whose information is to be collected or reflected.

If this specification is omitted, all devices on the specified server are assumed to be specified.

-f *File-Name*

Use the absolute path to specify the file name to which the difference between the device information managed currently and the device information collected is output. Only the device difference information is output to the file, and the reflection processing for the subsequent management information is not executed.
This option cannot be specified with the -y option.

-y

The difference between the device information managed by AdvancedCopy Manager and the latest device information detected is reflected unconditionally in the management information, without an inquiry being made about that difference.
This option cannot be specified with the -f option.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- The information for each device of a server is collected or reflected:

```
# /opt/FJSVswstf/bin/stgxfwcmsetdev -n Server1
>Device detected
  1:/dev/sdf0
  2:/dev/sdf1
  3:/dev/sdf2
  4:/dev/sdf6 (/storage-2)
>Specify device to be added. [* or ?-? or ?,? or q]
2,4

>Device not detected
  1:/dev/sdm0 (/aaa)
  2:/dev/sdm1

>Specify a device requested to be deleted. [* or ?-? or ?,? or q]
*
>Information of the following device has changed. The information will be changed.
  1:/dev/sdf4
>Do you want to continue the processing? [y or n] y
stgxfwcmsetdev completed
#
```

- The information for a specific device managed by the server is collected or reflected:

```
# /opt/FJSVswstf/bin/stgxfwcmsetdev -n Server1 -d /dev/sdl0
> The following device information has been changed. This information is updated.
1:/dev/sdl0 (/zzz)
> Do you want to continue the processing? [y or n] y
stgxfwcmsetdev completed
#
```

NOTES

- The following message is displayed if the number of devices is high:

```
 - hit Enter for more or s to select  -
```

Press the 'Enter' key to continue to display more devices. To interrupt the display and specify either an addition or a deletion, press 's' key.

- Use the following methods to request addition or deletion:

  - Enter "*" to specify all.

  - Enter numbers separated by a hyphen to specify a range. (Example: 1-3)

  - Enter numbers separated by a comma to specify separate numbers. (Example: 2, 4)

  - A combination of range and separate number specifications can be entered. (Example: 1-3, 5)

  - Enter no data to specify no devices to add or delete (only press Enter in this case).

  - Enter "q" to stop the processing.

- This command can be executed only by root users.

- Please perform the command in the location specified at the time of product installation.

- Obtaining disk information for HP-UX servers with GPT disks using AdvancedCopy Manager is not possible.

- When processing takes a long time, a period (.) is displayed until it is complete.

- The following message may be displayed when a device requested to be added is specified. It is an indication of an attempt to register a device that is in the box and may have invalid settings. Make sure that this problem is fixed before continuing the processing (This message is also displayed when the -y option is specified.)

```
>The setting of license or the table size is not done, or the table size is
insufficient, to the cabinet where the partition that makes up the following
device exists.
>Continue? [y or n]
```

- The following message may be displayed when a device has been requested to be added.

```
> The partitions that comprise the following device are already managed.
> Do you want to continue the processing? [y or n]
```

This message may be displayed in the following case.
Check the following before continuing the processing: whether the device will be subject to duplicated registration. This message is also displayed when the -y option is specified.

  - In the case two or more logical servers running on the same physical server are managed on the clustered system operating environment.

  - In the case a device with the other name that indicates a physical disk in ETERNUS Disk storage system is reflected into the management information of AdvancedCopy Manager.

- The amount of time required to process this command is proportional to the total number of devices defined on the specified server when the device information is collected. Therefore, when the number of devices is large, execute the command at a time of low CPU and I/O loads.
Note that it takes about 0.5 second for each device (partition) when there is no load on the system.

- This command cannot be executed while the following commands are being executed:

  - stgxfwcmaddsrv (Server information addition command)

  - stgxfwcmmodsrv (Server information change command)

  - stgxfwcmdelsrv (Server information deletion command)

  - stgxfwcmsetdev (Device information collection/reflection command)

  - stgxfwcmdeldev (Device information deletion command)

- This command cannot be executed by using the Web Console when the following operations are being performed:

  - Addition of a new server

  - Change of a server

  - Deletion of a server

  - Acquisition and Reflection of Information on All Devices

  - Acquisition and Reflection of Information on a Specific Device

  - Deletion of device information

9.3.7 stgxfwcmdeldev (Device information deletion command)

NAME

stgxfwcmdeldev - deletes the device information

SYNOPSIS

```
/opt/FJSVswstf/bin/stgxfwcmdeldev -n Server-Name -d Device-Name
```

DESCRIPTION

This command deletes the device information currently being managed by AdvancedCopy Manager. It is necessary to execute this command in the Management Server.

stgxfwcmsetdev (Device information collection/reflection command) compares the device information currently being managed and the device information detected before it executes the processing. However, this command deletes the device information irrespective of whether the device information currently being managed has been detected.

OPTIONS

-n *Server-Name*

Specify the name of the server that is managing the device to be deleted.
Specify a server name displayed by stgxfwcmdispsrv (Server information display command) exactly as displayed, including the case.

-d *Device-Name*

Specify the name of the device to be deleted.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- The device information for a Managed Server is deleted:

```
# /opt/FJSVswstf/bin/stgxfwcmdeldev -n Server1 -d /dev/sdn7
stgxfwcmdeldev completed
#
```

NOTES

- This command can be executed only by root users.

- Please perform this command in the location specified at the time of product installation.

- When processing takes a long time, a period (.) is displayed until it is complete.

- This command cannot be executed while the following commands are being executed:

  - stgxfwcmaddsrv (Server information addition command)

  - stgxfwcmmodsrv (Server information change command)

  - stgxfwcmdelsrv (Server information deletion command)

  - stgxfwcmsetdev (Device information collection/reflection command)

  - stgxfwcmdeldev (Device information deletion command)

- This command cannot be executed by using the Web Console if the following operations are being performed:

    - Addition of a new server

    - Change of a server

    - Deletion of a server

    - Acquisition and Reflection of Information on All Devices

    - Acquisition and Reflection of Information on a Specific Device

    - Deletion of device information

9.3.8 stgxfwcmdispsrv (Server information display command)

NAME

stgxfwcmdispsrv - displays the server information

SYNOPSIS

```
/opt/FJSVswstf/bin/stgxfwcmdispsrv [-n Server-Name]
```

DESCRIPTION

This command displays the server information managed by AdvancedCopy Manager. It is necessary to execute this command on the Management Server.

The contents to be displayed are as follows:

Title	Description
Server	A server name is displayed.
Server-Type	A server type is displayed. "Management-server" : Management Server "Server" : Managed Server
Tape	"---" is always displayed.
OS-Type	An operating system type is displayed.
OS-Version	An operating system version is displayed.
Version	The version of this product is displayed.
IP-Address	The IP address for communication with the server is displayed.
Port-Number	The port number for communication with the server is displayed.

OPTIONS

-n *Server-Name*

Specify the server name whose information is to be displayed.
Specify a server the same name entered at stgxfwcmaddsrv (Server information addition command) or stgxfwcmmodsrv (Server information change command) exactly, including the case.

If this option is omitted, the information for all servers managed by AdvancedCopy Manager is displayed.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- The information for all servers is displayed:

```
# /opt/FJSVswstf/bin/stgxfwcmdispsrv
Server  Server-Type       Tape OS-Type OS-Version    Version IP-Address  Port-Number
Manager Management-Server --   Linux   2.6.18-194.el5 V15.1   10.10.10.10 1226
Server1 Server            --   Linux   2.6.18-194.el5 V15.1   10.10.10.20 1226
#
```

- The information for a Managed Server(server name: Server1) is displayed:

```
# /opt/FJSVswstf/bin/stgxfwcmdispsrv -n Server1
Server  Server-Type Tape OS-Type OS-Version    Version IP-Address  Port-Number
Server1 Server      --   Linux   2.6.18-194.el5 V15.1   10.10.10.20 1226
#
```

NOTES

- In "OS-Version", the OS version that is displayed refers to the following time:
  In the case of the Management Server: the time when the database was created.
  In the case of the Managed Server: the time when the Managed Server was added.
  Even if the OS is upgraded after this, the OS version displayed will retain the version at the database creation time and the time the Managed Server was added.

- This command can be executed only by root users.

- Please perform this command in the location specified at the time of product installation.

- When processing takes a long time, a period (.) is displayed until it is complete.

9.3.9 stgxfwcmdispdev (Device information display command)

NAME

stgxfwcmdispdev - displays the Device information

SYNOPSIS

```
/opt/FJSVswstf/bin/stgxfwcmdispdev -n Server-Name [-d Device-Name [-sm|-grp]]
```

DESCRIPTION

This command displays information on a device being managed by AdvancedCopy Manager. The list of devices which use the same copy area as the specified device, the specified device itself, and the devices which exist in the same logical group can be displayed.
It is necessary to execute this command on the Management Server.

The following contents are displayed. Note that "SDX" or "SDX object" in the table are SynfinityDISK or PRIMECLUSTER Global Disk Services (GDS) terms.

Title	Description
S	The status of each device is indicated. The following is displayed according to the device status: * : Device not set up (The AdvancedCopy license setting may be undefined, cache table size may be unspecified, or the cache table may be too small in the box containing this device.) D : Device deletion error U : Device update error
Device	**Solaris server:** A RAW device name is displayed. **Windows server:** A device name managed internally by AdvancedCopy Manager is displayed. The name is displayed in the g?d? and g?d?p? format. g? : Disk array number (0 if the device is not installed in the ETERNUS Disk storage system) d? : Disk number(*1) p? : Logical disk number(*2) (*1)This is a number given uniquely within the ETERNUS Disk storage system. This number differs from the disk number in Disk Management. (*2)This is the sequential number of a logical disk on the physical disk. The logical disk indicates the following contents in Disk Management. The number is given by counting up from the left to the right in Disk Management. - In the case of a basic partition, the logical disk is that entire partition. - In the case of an extended partition, the logical disk is a logical drive in that partition. In the case of an SDX object, the name is displayed in the disk class name or mirror volume name g?d?p? format. Disk class name : Disk class name of SDX object Mirror volume name : Mirror volume name of SDX object **HP-UX server:** A RAW device name or a volume group name under LVM is displayed. **Linux server:** When a RAW device has been created, its name is displayed; otherwise, a block device name is displayed. **AIX server:** A RAW device name or a volume group name under LVM is displayed.
Block-Device	**Solaris server:** A block device name is displayed. **Windows server:** The name is displayed in the disk?p? format. disk? : Disk number displayed by disk administrator p? : Logical disk number (*2) **HP-UX server:** A block device name or a volume group name under LVM is displayed. **Linux server:** A block device name is displayed. **AIX server:** A block device name or a volume group name under LVM is displayed.

Title	Description

Mount-Point	Solaris, HP-UX, AIX and Linux server:
	
	    A mount point name is displayed.
	
	Windows server:
	
	    A drive character is displayed.

File-System	A file system name is displayed.
Device-Type	A device type is displayed. Normal/SDX(slice)/LVM/VxVM/SDX(volume)
Size(MB)	The capacity (in megabytes) allocated to a device is displayed.

RDB-Name	Windows and Linux server:
	
	    The Symfoware RDB system name is displayed.
	
	Solaris, HP-UX and AIX server:
	
	    "---" is always displayed.

Log-Group-Name	Windows and Linux server:
	
	    The Symfoware log group name is displayed.
	
	Solaris, HP-UX and AIX server:
	
	    "---" is always displayed.

DB-Space-Name	Windows and Linux server:
	
	    The Symfoware database space name is displayed.
	
	Solaris, HP-UX and AIX server:
	
	    "---" is always displayed.

OPTIONS

-n *Server-Name*

Specify the server name that manages the device whose information is to be displayed.

-d *Device-Name*

Specify the device name whose information is to be displayed.
If this option is omitted, the information for all devices managed by the specified server is displayed.

-sm

With the -d option also specified, this option specifies that a list of the devices using the same copy area as the specified device be displayed.

-grp

The list of the devices which exist in the same logical group as a specified device is displayed by specifying it simultaneously with the -d option. The device type of the specified device can be specified with the d option only for LVM/VxVM.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- The information for a device of a server is displayed:

```
# /opt/FJSVswstf/bin/stgxfwcmdispdev -n Server1 -d /dev/sdl0
S Device     Block-Device Mount-Point File-System Device-Type Size(MB) RDB-Name Log-Group-Name DB-
Space-Name
  /dev/sdl0 /dev/sdl0     /zzz        ext2        Normal       407      ---      ---            ---
#
```

- The information for each device of a server is displayed:

```
# /opt/FJSVswstf/bin/stgxfwcmdispdev -n Server1
S Device                        Block-Device                   Mount-Point     File-System Device-Type
Size(MB) RDB-Name Log-Group-Name DB-Space-Name
  /dev/sdl0                     /dev/sdl0                      /boot           ext2        Normal
141     ---      ---            ---
  /dev/sdu0                     /dev/sdu0                      ---             ---         Normal
407     ---      ---            ---
* /dev/sdu1                     /dev/sdu1                      ---             ---         Normal
407     ---      ---            ---
  /dev/sdu2                     /dev/sdu2                      ---             ---         Normal
28      ---      ---            ---
D /dev/sfdsk/Class2/rdsk/Vol201 /dev/sfdsk/Class2/dsk/Vol201 /ACMMnt/Vol201 ext3        SDX
304     ---      ---            ---
  /dev/sfdsk/Class2/rdsk/Vol202 /dev/sfdsk/Class2/dsk/Vol202 /ACMMnt/Vol202 ext3        SDX
400     ---      ---            ---
U /dev/sfdsk/Class2/rdsk/Vol203 /dev/sfdsk/Class2/dsk/Vol203 /ACMMnt/Vol203 ext3        SDX
512     ---      ---            ---
  /dev/sfdsk/Class2/rdsk/Vol204 /dev/sfdsk/Class2/dsk/Vol204 /ACMMnt/Vol204 ext3        SDX
608     ---      ---            ---
#
```

- Device of the server is displayed and device information which uses the same copy area is displayed:

```
# /opt/FJSVswstf/bin/stgxfwcmdispdev -n Server1 -d /dev/vx/rdsk/dg1/vol1 -sm
Device             Mount-Point
/dev/vx/rdsk/dg1/vol1 /storagedata
/dev/vx/rdsk/dg1/vol2 ---
/dev/vx/rdsk/dg1/vol3 /zzz
#
```

- Device information which exists in the same logical group as device of the server is displayed:

```
# /opt/FJSVswstf/bin/stgxfwcmdispdev -n Server1 -d /dev/vx/rdsk/dg1/vol1 -grp
Device             Mount-Point
/dev/vx/rdmp/c1t0d0s2 ---
/dev/vx/rdmp/c1t0d1s2 ---
/dev/vx/rdsk/dg1/vol1 /storagedata
/dev/vx/rdsk/dg1/vol2 ---
/dev/vx/rdsk/dg1/vol3 /zzz
/dev/vx/rdsk/dg1/vol4 ---
#
```

NOTES

- The command displays the above details based on information registered in AdvancedCopy Manager. Therefore, when the configuration has been changed and there is a device not registered in a logical group, and it is registered afterwards, the correct information is not displayed. In this case, please re-attempt the acquisition of device information about the registration of the device that was not previously registered, and the device which does the "Acquisition/reflection of information on all devices" processing.

- "RDB-Name", "Log-Group-Name" and "DB-Space-Name" are displayed under the following conditions:

    - The Managed Server's OS is Linux or Windows.

    - The Symfoware Server Advanced Backup Controller V10.0 or later is installed in the Managed Server.

    - The target device to be displayed exists in the ETERNUS Disk storage system.

    - When collecting the device information, the target RDB operating environment is running.

- This command can be executed only by root users.

- Please perform the command in the location specified at the time of product installation.

- When processing takes a long time, a period (.) is displayed until it is complete.

9.3.10 stgxfwcmdisppat (Partition information display command)

NAME

stgxfwcmdisppat - displays the information for the partitions

SYNOPSIS

```
/opt/FJSVswstf/bin/stgxfwcmdisppat -n Server-Name -d Device-Name
```

DESCRIPTION

This command displays the information for the partitions that comprise a device managed by AdvancedCopy Manager. It is necessary to execute this command in the Management Server.
The following contents are displayed. Note that "SDX" or "SDX object" in the table are SynfinityDISK or PRIMECLUSTER Global Disk Services (GDS) terms.

Title	Description
Partition	Solaris server: A RAW device name is displayed. In the case of an SDX object, the name is displayed in the RAW device name c?t?d? format. Windows server: When the device is installed in the ETERNUS Disk storage system, the same contents as stgxfwcmdispdev (Device information display command) are displayed. When the device is not installed in the ETERNUS Disk storage system, the name is displayed in the following formats. MBR disks: S?p? S? : Signature (8-digit hexadecimal number) p? : Logical disk number (*1) GPT disks: G? : DiskId of disk (32-bit hexadecimal character string) p? : Logical disk number (*1) (*1) This is the sequence number of a logical disk on the physical disk. The logical disk indicates the following contents in Disk Management. The number is given by counting up from the left to the right in Disk Management. - In the case of a basic partition, the logical disk is the entire partition. - In the case of an extended partition, the logical disk is a logical drive in that partition

Title	Description
	HP-UX server:
	A RAW device name or a volume group name under LVM is displayed.
	Linux server:
	When a RAW device has been created, its name is displayed. Otherwise, a block device name is displayed.
	In the case of an SDX object (i.e., a slice), it displays in the form of RAW device name:sd?.
	Refer to "1.5.2 SDX objects of PRIMECLUSTER GDS" for details.
	- RAW device name: The udev device name indicated by the sdxinfo DEVNAME value
	- RAW device name: The sdxinfo DEVNAME value
	AIX server:
	A RAW device name or a volume group name under LVM is displayed.
Identifier	**Solaris server:**
	A block device name is displayed.
	In the case of an SDX object, the name is displayed in the block device name c?t?d? format.
	Windows server:
	A device name managed internally by AdvancedCopy Manager is displayed.
	The same contents as Device of stgxfwcmdispdev (Device information display command) are displayed.
	HP-UX server:
	A block device name or a volume group name under LVM is displayed.
	Linux server
	A block device name is displayed.
	In the case of an SDX object (i.e., a slice), it displays in the form of a block device name:sd?.
	Refer to "1.5.2 SDX objects of PRIMECLUSTER GDS" for details.
	- RAW device name: The udev device name indicated by the sdxinfo DEVNAME value
	- RAW device name: The sdxinfo DEVNAME value
	AIX server:
	A block device name or a volume group name under LVM is displayed.
Physical-Disk	**Solaris, HP-UX, Linux, and AIX server:**
	"---" is always displayed.
	Windows server:
	A physical device name is displayed.
	The same contents as 'Block-Device' of stgxfwcmdispdev (Device information display command) are displayed.
Disk-Type	When the device is installed in the ETERNUS Disk storage system, "GR" is displayed.
	When the device is not installed in the ETERNUS Disk storage system, "---" is displayed.
LU-Number	When the device is installed in the ETERNUS Disk storage system, a logical unit number is displayed.
	When the device is not installed in the ETERNUS Disk storage system, "---" is displayed.

OPTIONS

-n *Server-Name*

Specify the server name that manages a device whose information is to be displayed.

Specify a server name displayed by stgxfwcmdispsrv (Server information display command) exactly as displayed, including the case.

-d *Device-Name*

    Specify a device name whose information is to be displayed.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- The partition information for a device of a server is displayed:

```
# /opt/FJSVswstf/bin/stgxfwcmdisppat -n Server1 -d /dev/sdl0
Partition Identifier Physical-Disk Disk-Type LU-Number
/dev/sdl0 /dev/sdl0  ---           GR        10
#
```

NOTES

- This command can be executed only by root users.

- Please perform the command in the location specified at the time of product installation.

- When processing takes a long time, a period (.) is displayed until it is complete.

9.3.11 stgenvdisp (Environment information display command)

NAME

stgenvdisp - displays the values of the environment settings

SYNOPSIS

```
/opt/FJSVswstf/bin/stgenvdisp [-n Server-Name]
```

DESCRIPTION

This command displays the values of the environment settings in AdvancedCopy Manager.
The contents of the displayed information are described below.

Title	Description
Server type	The server type is displayed as follows: Management server : Management Server or Management Server transaction Server : Managed Server or Managed Server transaction
Cluster system	The type of clustered system is displayed as follows: VERITAS Cluster Server : VERITAS Cluster Server MSCS : MSCS or WSFC HACMP : High Availability Cluster Multi-Processing GENERAL : Clustered system on Solaris, Linux, or HP-UX environment
Node type	The node type is displayed as follows:

Title	Description
	Primary : Primary node Secondary : Secondary node
Node name	The logical node name is displayed. Local-Service indicates a local transaction.
Service name	The transaction name is displayed.
Cluster name	The cluster name is displayed.
Mount point	The mount point for AdvancedCopy Manager shared data is displayed.
Disk resource	The name of the shared disk resource used for AdvancedCopy Manager shared data is displayed.
Drive	The drive letter used for AdvancedCopy Manager shared data is displayed.
Network name	The network name used for AdvancedCopy Manager is displayed.
External connection	The external connection network name is displayed.
Subnet mask	The subnet mask is displayed.
IP address resource	IP address resource information is displayed.
IP by cluster setup	The logical IP address specified in the cluster setting command is displayed. However, note that the displayed IP address is not the active logical IP address if the IP address of either the Management Server transaction or Managed Server transaction is changed during operation.
Manager IP address	The IP address of the Management Server is displayed. However, if no server is added, then "????" is displayed.
Repository user	The name of the user accessing the repository is displayed.
DB log file	Windows Version only: The name of the DB log file storage destination folder is displayed.
DB dictionary file	Windows Version only: The name of the DB dictionary file storage destination folder is displayed.
Repository database	Windows Version only: The name of the repository DB storage destination folder is displayed.
DB file	Solaris Version: The name of the DB file system device or directory is displayed. Linux Version: The DB file directory name is displayed.
RDB dictionary	Solaris Version: The RDB dictionary raw device name is displayed. Linux Version: The RDB dictionary directory name is displayed.
RDB log file	Solaris Version: The name of the RDB log file raw device or directory is displayed. Linux Version: The RDB log file directory name is displayed.
Database space	Solaris Version: The name of the raw device containing the repository data storage DB is displayed.

Title	Description
	Linux Version: The name of the directory containing the repository data storage DB is displayed.
Database size	Windows and Linux Version only: The database size is displayed.
Communication	The communication service port number (the service name in the services file: stgxfws) is displayed. In a cluster environment, the transaction communication service port number (the service name in the services file: stgxfws_logical-node-name) is displayed.
Remote access	The remote access service port number (the service name in the services file: rmiswstf) is displayed.
Exclusive control	The exclusive control service port number (the service name in the services file: semswstf) is displayed.
Repository	Windows and Solaris Version only: The repository service port number (the service name in the services file: rdaswstf) is displayed.
Program install directory	Displays the "Program directory" specified during installation.
Fixed configuration directory	Displays the "Environment directory" specified during installation.
Modifying configuration directory	Displays the "Work directory" specified during installation.
Version	Displays the version and level of AdvancedCopy Manager.
Language	Displays the AdvancedCopy Manager language type. JAPANESE : Japanese version ENGLISH : English version
Code	Solaris and Linux Version only: Displays the encoding specified during installation.

OPTIONS

-n *Server-Name*

Specifies the logical node name (logical node name specified in the cluster setup command) of the transaction.

In a cluster environment

If this option is specified, information on only the specified transactions is output. Otherwise, information on all transactions is output.

In a non-cluster environment

This option cannot be specified. If the option is specified, then message swstf6002 is output and processing ends. Otherwise, information is output.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Information in a non-clustered system is displayed:

[For the Management Server]

```
# /opt/FJSVswstf/bin/stgenvdisp
  Server type           : Management server
  Manager IP address    : 10.10.10.10
  Repository user       : acm
  DB file               : /acmdb/dbfile
  RDB dictionary        : /acmdb/rdbdic
  RDB log file          : /acmdb/rdblog
  Database space        : /acmdb/dbsp
  Database size         : 65MB
  Communication         : 1226
  Remote access         : 9851
  Exclusive control     : 9852
-------------------------------------------------------------------------
  Version               :  V15.1
  Language              : ENGLISH
  Program install directory        : /opt
  Fixed configuration directory    : /etc/opt
  Modifying configuration directory  : /var/opt
#
```

[For the Managed Server]

```
# /opt/FJSVswstf/bin/stgenvdisp
  Server type           : Server
  Manager IP address    : 10.10.10.10
  Communication         : 1226
-------------------------------------------------------------------------
  Version               : V15.1
  Language              : ENGLISH
  Program install directory        : /opt
  Fixed configuration directory    : /etc/opt
  Modifying configuration directory  : /var/opt
#
```

- Information on all transactions in a cluster environment is displayed.

[For the Management Server]

```
# /opt/FJSVswstf/bin/stgenvdisp
  Server type           : Management server
  Cluster system        : GENERAL
  Node type             : Primary
  Node name             : nodemgr
  Service name          : manager_service
  Mount point           : /acmmgrdisk
  IP by cluster setup   : 10.10.10.10
  Manager IP address    : 10.10.10.10
  Repository user       : acm
  DB file               : /acmmgrdisk/acmdb
  RDB dictionary        : /acmmgrdisk/acmdb
  RDB log file          : /acmmgrdisk/acmdb
  Database space        : /acmmgrdisk/acmdb
  Database size         : 65MB
  Communication         : 11000
  Remote access         : 9851
  Exclusive control     : 9852
-------------------------------------------------------------------------
  Server type           : Server
  Node name             : Local-Service
  Manager IP address    : 10.10.10.10
  Communication         : 1226
-------------------------------------------------------------------------
  Version               : V15.1
  Language              : ENGLISH
```

```
   Program install directory        : /opt
   Fixed configuration directory     : /etc/opt
   Modifying configuration directory : /var/opt
#
```

[For the Managed Server]

```
# /opt/FJSVswstf/bin/stgenvdisp
   Server type               : Server
   Cluster system            : GENERAL
   Node type                 : Primary
   Node name                 : nodeagt1
   Service name              : agent_service1
   Mount point               : /acmagt1disk
   Manager IP address        : 10.10.10.10
   Communication             : 11000
   -------------------------------------------------------------------------
   Server type               : Server
   Cluster system            : GENERAL
   Node type                 : Secondary
   Node name                 : nodeagt2
   Service name              : agent_service2
   Mount point               : /acmagt2disk
   Manager IP address        : 10.10.10.10
   Communication             : 12000
   -------------------------------------------------------------------------
   Server type               : Server
   Node name                 : Local-Service
   Manager IP address        : 10.10.10.10
   Communication             : 1226
   -------------------------------------------------------------------------
   Version                   : V15.1
   Language                  : JAPANESE
   Program install directory        : /opt
   Fixed configuration directory     : /etc/opt
   Modifying configuration directory : /var/opt
#
```

NOTES

- To output cluster operation information, execute the command from an active node. Executing the command from a standby node displays only partial information. To output cluster operation information while a transaction is stopped, either set the shared disk containing the AdvancedCopy Manager shared data of the transaction to the online state, or mount the shared disk and then execute the command.

- When this command is executed in a clustered system, the environment variable SWSTGNODE need not be set.

- This command can be executed only by root users.

- "????" is displayed for any information that cannot be displayed because of an error.

- Some items may not be output, depending on the server type (whether Management Server or Managed Server) or cluster type.

9.3.12 stgmgrinfoset (Management information package backup server registration/deletion command)

NAME

stgmgrinfoset - registers. deletes or displays the management information

SYNOPSIS

```
/opt/FJSVswstf/bin/stgmgrinfoset [-del | -d Directory-Name -c management-server-encoding]
```

DESCRIPTION

This command can be used to register, delete or display information about the management information package backup execution server(the server on which stgmgrinfobkup (Management information package backup command) is executed).

- This command registers the management information package backup execution server.

- The command deletes the management information package backup execution server.

- The command displays information about the management information package backup execution server.

OPTIONS

None

Displays information about the management information package backup execution server.

-del

Deletes the management information package backup execution server.

-d *Directory-Name*

Specify the absolute path to the directory where the backup data is to be stored.
Only alphanumeric characters and underscores can be used in the directory name.

-c *management-server-encoding*

Specify the encoding designated when the Management Server was installed. Specify ASCII

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Register the management information package backup execution server:

  (assume that the path to the directory where backup data is to be stored is /stgbkupdata, and the encoding used on the Management Server is ASCII):

  ```
  # /opt/FJSVswstf/bin/stgmgrinfoset -d /stgbkupdata -c ASCII
  swstf3000 The server which performs package backup of the management information on AdvancedCopy
  Manager was registered.
  #
  ```

- Delete the management information package backup execution server:

  ```
  # /opt/FJSVswstf/bin/stgmgrinfoset -del
  swstf3001 The server which performs package backup of the management information on AdvancedCopy
  Manager was deleted. Delete a directory. Directory=/stgbkupdata
  #
  ```

- Display information about the management information package backup execution server:

  ```
  # /opt/FJSVswstf/bin/stgmgrinfoset
  Directory    : /stgbkupdata
  ```

```
Manager code : ASCII
#
```

NOTES

- This command can be executed only by root users.

- To change the backup data storage directory, first use the -del option to delete the current directory, then register a new backup data storage directory.

- Register only one Managed Server as the management information package backup execution server.

- Set the name of the logical node in the environment variable SWSTGNODE when the Managed Server registered as the management information package backup execution server is in a clustered system.

- Create the directory specified in the -d option in advance.

9.3.13 stgmgrinfobkup (Management information package backup command)

NAME

stgmgrinfobkup - backup the management information

SYNOPSIS

```
/opt/FJSVswstf/bin/stgmgrinfobkup [-n Server-Name] [-t Backup-Type]
```

DESCRIPTION

This command backup the management information stored on the Managed Server and the Management Server.

OPTIONS

-n *Server-Name*

Creates backups of the management information stored on the specified server.
If this option is not specified, the management information on all Managed Server (including the Management Server) is backed up.

-t *Backup-Type*

Specifies the types of data that are to be backed up. Specify one of the values listed below. Do not specify more than one value in the -t option:

DBData : Database
RegiData : Registry
MgrInfoData : Policy file for management information package backup
BkupData : Backup management
ReplData : Replication management

If this option is not specified, backups are created as defined in the policy file. For details on the policy file, refer to the section "8.1.2.1 Preliminaries".

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Back up the management information on all the Managed Server as defined in the policy file:

```
# /opt/FJSVswstf/bin/stgmgrinfobkup
swstf3014 Backup of management information is started. Start time=04/07/02 11:00:00, Server
name=acm01
swstf3020 Backup of management information was ended normally. End time=04/07/02 11:02:00, Server
name=acm01
swstf3014 Backup of management information is started. Start time=04/07/02 11:02:00, Server
name=acm02
swstf3020 Backup of management information was ended normally. End time=04/07/02 11:05:00, Server
name=acm02
swstf3021 The normal end of the backup of management information was carried out by all servers.
#
```

- Back up the management information on Managed Server as defined in the policy file:

```
# /opt/FJSVswstf/bin/stgmgrinfobkup -n acm01
swstf3014 Backup of management information is started. Start time=04/07/02 11:00:00, Server
name=acm01
swstf3020 Backup of management information was ended normally. End time=04/07/02 11:02:00, Server
name=acm01
swstf3021 The normal end of the backup of management information was carried out by all servers.
#
```

- Back up the registries on all of the Managed Server:

```
# /opt/FJSVswstf/bin/stgmgrinfobkup -t RegiData
swstf3014 Backup of management information is started. Start time=04/07/02 11:00:00, Server
name=acm01
swstf3020 Backup of management information was ended normally. End time=04/07/02 11:02:00, Server
name=acm01
swstf3014 Backup of management information is started. Start time=04/07/02 11:02:00, Server
name=acm02
swstf3020 Backup of management information was ended normally. End time=04/07/02 11:05:00, Server
name=acm02
swstf3021 The normal end of the backup of management information was carried out by all servers.
#
```

- Back up the backup management located on Managed Server:

```
# /opt/FJSVswstf/bin/stgmgrinfobkup -n acm01 -t BkupData
swstf3014 Backup of management information is started. Start time=04/07/02 11:00:00, Server
name=acm01
swstf3020 Backup of management information was ended normally. End time=04/07/02 11:02:00, Server
name=acm01
swstf3021 The normal end of the backup of management information was carried out by all servers.
#
```

NOTES

- This command can be executed only by root users.

- Be sure to check the information (ie, a message indicating completion or an error) that is output by this command.

- This command can be executed only on the management information package backup execution server.

- Set the name of the logical node in the environment variable SWSTGNODE when the management information package backup execution server is in a clustered system.

- Do not execute other AdvancedCopy Manager commands or Web Console operations while this command is running. Conversely, do not execute this command while another AdvancedCopy Manager command or a Web Console operation is running.

- Do not execute this command again if it is already running.

- If this command was executed under one of the following circumstances from a Managed Server on which AdvancedCopy Manager Version 14.2 or earlier is installed, backup of authentication data will not be implemented.

    - When AuthData was specified in the -t option

    - When the value "YES" was specified for the AuthData key in the Management Server section of the policy file for batch backup of management data

- To save backup data to a tape or similar external medium, save all files and directories under the backup data storage directory.

- Registry backups are made only when RegiData is specified in the -t option. A registry backup will not be made if the -t option is not specified.

- A backup of the policy file for management information package backup is made only when MgrInfoData is specified in the -t option. No backup is made when the -t option is not specified. Additionally, specify the name of the management information package backup execution server in the -n option.

- Backup is not possible for any management information on Managed Server with Softek AdvancedCopy Manager Version 10.4 (V10.0L50) or earlier.

9.3.14 stgmgrinfodisp (Management information package backup status display command)

NAME

stgmgrinfodisp - displays the status of execution of the stgmgrinfobkup

SYNOPSIS

```
/opt/FJSVswstf/bin/stgmgrinfodisp [-n server-name]
```

DESCRIPTION

This command displays the status of execution of stgmgrinfobkup (Management information package backup command).
The types of information listed below are output.
If no backups are created when the management information package backup command is executed or if the management information package backup command ends in error, then no data is output.

Item	Description
Server	Server name
RegiData	Date and time at which the management information package backup command created a backup of a registry (ie, the date and time that the management information package backup command was executed) In yy/mm/dd hh:mm:ss format
DBData	Date and time at which the management information package backup command created a backup of the database (ie, the date and time that the management information package backup command was executed) yy/mm/dd hh:mm:ss
BkupData	Date and time at which the management information package backup command created a backup of the backup management (ie, the date and time that the management information package backup command was executed) yy/mm/dd hh:mm:ss

Item	Description
ReplData	Date and time at which the management information package backup command created a backup of replication management (ie, the date and time that the management information package backup command was executed) yy/mm/dd hh:mm:ss
MgrInfoData	Date and time at which the management information package backup command created a backup of the policy file for management information package backup (ie, the date and time that the management information package backup command was executed) yy/mm/dd hh:mm:ss

OPTIONS

-n *Server-Name*

Specifies the name of the server where the backup status is to be output.
If this option is not specified, the status of backups on the Managed Server for which the management information package backup command was executed is displayed.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Output the information about all the Managed Server:

```
# /opt/FJSVswstf/bin/stgmgrinfodisp
Server:        acm01
   DBData:      04/08/02 16:11:26
   RegiData:    04/08/01 09:00:00
   ReplData:    04/08/02 16:32:22
Server:        acm02
   RegiData:    04/08/01 09:00:00
   BkupData:    04/08/03 16:00:00
#
```

- Output the information about Managed Server named acm01:

```
# /opt/FJSVswstf/bin/stgmgrinfodisp -n acm01
Server:        acm01
   DBData:      04/08/02 16:11:26
   RegiData:    04/08/01 09:00:00
   ReplData:    04/08/02 16:32:22
#
```

NOTES

- This command can be executed only by root users.

- This command can be executed only on the management information package backup execution server.

- When the management information package backup execution server is in a clustered system, set the name of the logical node in the environment variable SWSTGNODE.

9.3.15 stgxfwcmsetsdx (SDX object management unit configuration command)

NAME

stgxfwcmsetsdx - sets SDX object management unit

SYNOPSIS

```
/opt/FJSVswstf/bin/stgxfwcmsetsdx [-slice | -volume]
```

DESCRIPTION

The unit of operation of the SDX object treated with AdvancedCopy Manager (Refer to the "3.8.1 Units of SDX object operations") is set. Moreover, the unit of operation of a present SDX object is displayed.

The display form is 'Unit of management (SWSTGNODE = <configuration value>).'

- Unit of management

- Slice in case of each slice

- Volume in case of each logical volume

One of these values are used value-environment variable SWSTGNODE is set. The value is Null when environment variable SWSTGNODE is not set.

OPTIONS

The unit of present of the SDX object management is displayed.

-slice

The unit of management of the SDX object treated with AdvancedCopy Manager is set to the unit of the slice.

-volume

The unit of management of each logical volume of the SDX object treated with AdvancedCopy Manager is set.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- The unit of present of the SDX object management in non-cluster environment is displayed:

```
# /opt/FJSVswstf/bin/stgxfwcmsetsdx
Volume (SWSTGNODE=null)
#
```

- The unit of the SDX object of the Management Server service (The logical node name: aserv) is set to the unit of the slice:

```
# SWSTGNODE=aserv; export SWSTGNODE
# /opt/FJSVswstf/bin/stgxfwcmsetsdx -slice
stgxfwcmsetsdx completed (SWSTGNODE=aserv)
#
```

NOTES

- This command can be executed only by root users.

- When the cluster is operated, each Managed Server service is set.

9.3.16 stgxfwcmsetmode (Information collection mode configuration command)

NAME

stgxfwcmsetmode - sets the collection mode for information

SYNOPSIS

```
/opt/FJSVswstf/bin/stgxfwcmsetmode -udev {on | off | disp}
/opt/FJSVswstf/bin/stgxfwcmsetmode -udev_gds {on | off | disp}
/opt/FJSVswstf/bin/stgxfwcmsetmode -virt {kvm | none | disp}
```

DESCRIPTION

This command sets the collection mode for information collected by the server. It also displays the collection mode that is currently set.

For operation in a virtual environment, this command sets the Virtual Environment Support Function.
It also displays the status of Virtual Environment Support Function that is currently set.

OPTIONS

-udev {on | off | disp}

For ordinary devices, specifies whether to use the device generated by udev. Refer to "1.5.1 Normal devices" for details. The following values can be specified:

on : The device generated by udev is used.
off : The device generated by udev is not used.
disp : Checks whether the device generated by udev is to be used or not.

Information output at time of disp specification

udev : on The device generated by udev is used.
udev : off The device generated by udev is not used.

-udev_gds {on | off | disp}

For PRIMECLUSTER GDS SDX objects (only when operating in slice units), specifies whether to use the device generated by udev. Refer to "1.5.2 SDX objects of PRIMECLUSTER GDS" for details. The following values can be specified:

on : The device generated by udev is used.
off : The device generated by udev is not used.
disp : Checks whether the device generated by udev is to be used or not.

Information output at time of disp specification

udev : on The device is used.
udev : off The device is not used.

-virt {kvm | none | disp}

Specifies whether to use the Virtual Environment Support Function. The following values can be specified:

kvm : The KVM Support Function is used.

none : The Virtual Environment Support Function is not used.

disp : Check the current setup status of Virtual Environment Support Function.

Information output when disp is specified

virt : kvm The KVM Support Function is used.

virt : none The Virtual Environment Support Function is not used.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- To specify that the device generated by udev is not used:

```
# /opt/FJSVswstf/bin/stgxfwcmsetmode -udev off
stgxfwcmsetmode completed (SWSTGNODE=nodemgr)
#
```

- To check whether the device generated by udev is to be used or not:

```
# /opt/FJSVswstf/bin/ stgxfwcmsetmode -udev disp
udev : on (SWSTGNODE=null)
#
```

- To specify that the KVM Support Function is used:

```
# /opt/FJSVswstf/bin/stgxfwcmsetmode -virt kvm
stgxfwcmsetmode completed (SWSTGNODE=nodeagt)
#
```

- To check the current setup status of Virtual Environment Support Function:

```
# /opt/FJSVswstf/bin/stgxfwcmsetmode -virt disp
virt:kvm (SWSTGNODE=nodeagt)
#
```

NOTES

- If this command is used to change the set value at a time when information concerning managed devices exists, a message indicating that device information has been changed may be posted the next time device information is collected or displayed. Therefore, the backup management and replication management definition information must be deleted before the setting is changed.

- For backup operations in link with Symfoware, if you use the udev names in mind at the time of definition of the Symfoware database spaces, do not make the setting "The device generated by udev is not used". If the settings contain any mistakes, information related to Symfoware will possibly not be retrieved correctly.

- If this command is executed in an environment where udev is not installed, the command returns an error.

9.3.17 stgdbloginf (Repository update area busy rate check command)

NAME

stgdbloginf - check repository update area busy rate

SYNOPSIS

```
/opt/swstorage/bin/stgdbloginf
```

DESCRIPTION

This command checks the busy rate of the area required when a repository is updated.
It is necessary to execute this command in the Management Server.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

```
# /opt/swstorage/bin/stgdbloginf
80%
```

NOTES

- This command can be executed only by root users.

9.3.18 stgdbdmp (Database save command)

NAME

stgdbdmp - saves the database

SYNOPSIS

```
/opt/swstorage/bin/stgdbdmp -d Directory-Name [-y]
```

DESCRIPTION

This command saves the database to the specified directory.
It is necessary to execute this command in the Management Server.

OPTION

-d *Directory-Name*

Specifies the name of the directory where the database area is saved.
This directory must already have been created.

-y

If the specified directory contains contents that have already been saved, this option overwrites that content when the database area is saved.
If this option is omitted and if the specified directory contains contents that have already been saved, the save processing is not performed and processing ends.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

Save the database in a directory (/storage/dmpdir):

```
# /opt/swstorage/bin/stgdbdmp -d /storage/dmpdir -y
////// AdvancedCopy Manager database save //////
====================================================================
DB data will be saved.
====================================================================
Editing of the work file used for save processing is complete.
.
Update of the DB space is suppressed.
Forced switching of archive log files is complete.
.........................
Save processing of DSI data is complete.
Archive log files have been completely discarded.
Save processing of RDB dictionary data is complete.
Suppression of DB space updates has been completely canceled.
====================================================================
Save processing of DB data is complete.
====================================================================
Database save processing ended normally.
#
```

NOTES

- This command can be executed only by root users.

- Execute this command under the locale specified at the time of introducing the Management Server. If Symfoware is introduced in advance, execute the command under the locale specified at the time of introducing Symfoware.

- If no data is saved by this command, the respective database cannot be recovered with the "8.1.1.3.3 Recovering a database" if an I/O error occurs in the database space.

- If the Management Server operates in a clustered system, it is recommended that the directory used to save the database be prepared on a shared disk.

- Processing may end abnormally if another application or command is executing while this command is being processed. In this event, retry processing after the other operation has finished.

9.3.19 stgdbrcv (Database recovery command)

NAME

stgdbrcv - recovers the database

SYNOPSIS

```
/opt/swstorage/bin/stgdbrcv -m mode [-d Backup-Directory]
```

DESCRIPTION

This command recovers a saved database.

It is necessary to execute this command in the Management Server.

OPTION

-m *Mode*

> This option specifies the recovery mode. Specify the mode that suits the fault that occurred in the database area.
>
> Refer to "Table 8.3 Relationship between qdg messages and recovery modes" for details on faults and the corresponding recovery modes.

-d *Backup-Directory*

> Specify the name of the directory storing the saved data.
>
> This option must be specified if 0, 1, 2, or 4 is specified in the -m option.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

```
# /opt/swstorage/bin/stgdbrcv -m 1 -d /storage/dmpdir
////// AdvancedCopy Manager database recovery //////
==================================================================
DSI recovery (FW mode) will start.
==================================================================
Editing of the work file used for recovery is complete.
DSI access inhibition is complete.
........................
DSI data recovery is complete.
==================================================================
DSI recovery (FW mode) is complete.
==================================================================
Database recovery ended normally.
#
```

NOTES

- This command can be executed only by root users.

- Execute this command under the locale specified at the time of introducing the Management Server. If Symfoware is introduced in advance, execute the command under the locale specified at the time of introducing Symfoware.

- After the database recovery command terminates normally, obtain the most recently saved data as described in "8.1.1.3.1 Saving a database". At this time it is recommended to specify another directory rather than the saved data storage directory used in the recovery command.

- After the database recovery command terminates normally, refer to "Chapter 2 Starting and Stopping Daemons"and restart (stop and then start) the AdvancedCopy Manager daemon.

9.3.20 stgdbbackup (Database backup command)

NAME

stgdbbackup - backups the database

SYNOPSIS

```
/opt/swstorage/bin/stgdbbackup
```

DESCRIPTION

This command backs up the data stored in the database to the specified directory as a batch.
It is necessary to execute this command in the Management Server.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

Refer to "8.2.1.1 Database backup".

NOTES

- This command can be executed only by root users.

- The directory where the data is to be saved must already be created.

- Execute this command under the locale specified at the time of introducing the Management Server. If Symfoware is introduced in advance, execute the command under the locale specified at the time of introducing Symfoware.

9.3.21 stgdbrestore (Database restoration command)

NAME

stgdbrestore - restores the database

SYNOPSIS

```
/opt/swstorage/bin/stgdbrestore
```

DESCRIPTION

This command restores to the database as a batch the data that was saved by database backup.
It is necessary to execute this command in the Management Server.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

Refer to "8.2.1.2 Database restoration".

NOTES

- This command can be executed only by root users.

- During this process, the data already stored in the database is overwritten by the saved data. Therefore, the data currently stored in the database is lost.

- Execute this command under the locale specified at the time of introducing the Management Server. If Symfoware is introduced in advance, execute the command under the locale specified at the time of introducing Symfoware.

- After database restoration ends normally, execute the processing described in "8.1.1.3.1 Saving a database", and then save the post-restoration database area.

9.3.22 stgdbset (Database integrated setup command)

NAME

stgdbset - creats a database

SYNOPSIS

```
/opt/swstorage/bin/stgdbset
```

DESCRIPTION

This command starts database integrated setup and creates database environments and databases.
It is necessary to execute this command in the Management Server.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

Refer to "Database creation" of the "ETERNUS SF Installation and Setup Guide".

NOTES

- This command can be executed only by root users.

- In advance, the user must perform the tasks described in "Estimation of database area", "Directory creation" and "Kernel parameter tuning" in the "ETERNUS SF Installation and Setup Guide" and must check the IP address of the Management Server.

- Four directories can be created in the same partition, but preparation of the directory in a separate partition is recomended for reliability and update performance.

- Execute this command under the locale specified at the time of introducing the Management Server. If Symfoware is introduced in advance, execute the command under the locale specified at the time of introducing Symfoware.

- Before starting database integrated setup, check whether the setting for using a shadow password has been set for the system. If use of a shadow password is not set, change the setting so that a shadow password is used, and then start database integrated setup. In addition, after database creation, always operate the database with the setting for using a shadow password set.

9.3.23 stgdbunset (Database integrated unsetup command)

NAME

stgdbunset - deletes the database

SYNOPSIS

```
/opt/swstorage/bin/stgdbunset
```

DESCRIPTION

This command starts database integrated unsetup, deletes databases and database environments, and cancels repository settings.
It is necessary to execute this command in the Management Server.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

Refer to "Cancel of repository configurations" of the "ETERNUS SF Installation and Setup Guide".

NOTES

- This command can be executed only by root users.

- If the Management Server operates in a cluster configuration, deleting a Management Server task also cancels the repository setting. Therefore, this operation is not required.

- Execute this command under the locale specified at the time of introducing the Management Server. If Symfoware is introduced in advance, execute the command under the locale specified at the time of introducing Symfoware.

- Database integrated unsetup uses Symfoware functions. If other software performs Symfoware's monitoring, database integrated unsetup might fail. Therefore, stop Symfoware monitoring before starting database integrated unsetup.

9.3.24 stgrepocnv (Repository data migration command)

NAME

stgrepocnv - migrates the repository data

SYNOPSIS

```
/opt/FJSVswstf/bin/stgrepocnv -d Directory-Name
```

DESCRIPTION

This command recovers the migration target repository data to the current version repository.
It is necessary to execute this command in the Management Server.

OPTION

-d *Directory-Name*

Specify the absolute path of the directory where the migration target repository data is stored.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

```
/opt/FJSVswstf/bin/stgrepocnv -d /tmp/backdir
.................
stgrepocnv completed
```

NOTES

- This command can be executed only by root users.

- For Management Server transactions, execute this command at a currently operating server. The environment variable SWSTGNODE need not be set.

9.3.25 stgforcedelsrv (Managed Server forcibly deletion command)

NAME

stgforcedelsrv - deletes the information of Managed Server forcibly

SYNOPSIS

```
/opt/FJSVswstf/bin/stgforcedelsrv -n Managed-Server -b before-dir -a after-dir
```

DESCRIPTION

This command deletes the Managed Server information forcibly. This command must be executed on a Management Server.

As this command deletes the backup/replication settings of the Managed Server, it must be performed with caution. Before performing the command always refer to "8.6 Action when Managed Server information cannot be deleted".

OPTIONS

-n *Managed-Server*

Specify the server name of Managed Server to be deleted forcibly.

-b *before-dir*

Specify the full path name of the directory containing the database backup.
The command deletes the Managed Server information forcibly from the database in the directory specified with the -b option and stores the database after deletion in the directory specified with -a option.

-a *after-dir*

Specify the full path name of the directory that will store the database after the Managed Server information has been deleted. Please note the following important points:

- If the specified directory does not exist, it must be created before executing the command.

- Specify an empty directory, containing no subdirectories or files.

- Specify a directory that will not disappear after the server has rebooted.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

Delete the information of Managed Server (Server1). After deleting its information, the database is created under the "/var/tmp/afterDB" directory.

```
# /opt/FJSVswstf/bin/stgforcedelsrv -n Server1 -b /var/tmp/beforeDB -a /var/tmp/afterDB
Begin the forcing deletion? (y/n) : y
stgforcedelsrv completed
#
```

NOTES

- This command can be executed only by root users.

- Please perform the command in the location specified at the time of product installation.

- The Management Server information cannot be deleted.

- Do not perform AdvancedCopy Manager operations (backup, replication, modifications to the operating environment, etc.) while executing this command.

- If the command terminates abnormally, temporary files may remain under the directory specified with -a option. When the command terminates abnormally, delete all the files under the directory specified with the -a option before executing the command again.

- Several instances of the command cannot be executed simultaneously.

9.4 Replication Management Commands

This section describes operations using Replication Management Commands.

📖 Note

Only a root user can execute all the commands of replication management. When command execution is carried out by other users, a message "swsrp0601 This Command cannot be run. This is because the user is not a root user." is displayed, and the command aborts.

📖 Note

Operating environment for command execution

The table below lists the requirements for the operating environment in which to execute the different commands.

Table 9.16 Operating environment when executing replication management commands

Command	Operating environment of AdvancedCopy Manager
swsrpcancel	To execute this command on a Management Server with the -h option specified, Windows services required by the Management Server and target Managed Server must be operating.

Command	Operating environment of AdvancedCopy Manager
	When the Managed Server runs on Windows and is in EC operation, Windows services required by the copy destination Managed Server must be running (except when the -m or -t option is specified). For inter-server replication in EC equivalency status, the associated server must be operating (except when the -m option is specified).
swsrpchsync	To execute this command on a Management Server with the -h option specified, Windows services required by the Management Server and target Managed Server must be running.
swsrpdelvol	To execute this command on a Managed Server, Windows services required by the Management Server must be running. To execute this command on a Management Server with the -h option specified, Windows services required by the Management Server and target Managed Server must be running. For inter-server replication, Windows services required by the Managed Server in communication must be running (except when the -d option is specified).
swsrprecbuffstat	To execute this command on a Management Server with the -h option specified, Windows services required by the Management Server and target Managed Server must be running.
swsrprecbuffset	To execute this command on a Managed Server, Windows services required by the Management Server must be running. To execute this command on a Management Server with the -h option specified, Windows services required by the Management Server and target Managed Server must be running.
swsrpmake	To execute this command on a Management Server with the -h option specified, Windows services required by the Management Server and target Managed Server must be running. When the Managed Server runs on Windows and is engaged in intra-server replication and EC operation, Windows services required by the Managed Server must be running. For inter-server replication, Windows services required by the Managed Server in communication must be running (except when the -m or -t option is specified).
swsrprecoverres	To execute this command on a Management Server with the -h option specified, Windows services required by the Management Server and target Managed Server must be running. When the Managed Server runs on Windows, Windows services required by the Managed Server must be running. To specify the -r option, Windows services required by the Management Server must be running.
swsrprevsync	To execute this command on a Management Server with the -h option specified, Windows services required by the Management Server and target Managed Server must be running.
swsrpsetvol	To execute this command on a Managed Server, Windows services required by the Management Server must be running. To execute this command on a Management Server with the -h option specified, Windows services required by the Management Server and target Managed Server must be running. For inter-server replication, Windows services required by the Managed Server in communication must be running.
swsrpstartsync	To execute this command on a Management Server with the -h option specified, Windows services required by the Management Server and target Managed Server must be running. When the Managed Server runs on Windows, Windows services required by the Managed Server must be running. For inter-server replication, the associated server must be running (except a case where the -m option is specified).
swsrpstat	To execute this command on a Management Server with the -h option specified, Windows services required by the Management Server and target Managed Server must be running.
swsrpvolinfo	To execute this command on a Management Server with the -h option specified, Windows services required by the Management Server and target Managed Server must be running.

In cluster operation

In cluster operation, execute a command by taking the following steps:

1. If necessary, specify the logical node name to the transaction in environment variable SWSTGNODE. Specify the environment variable as follows:

```
# SWSTGNODE=logical-node-name
# export SWSTGNODE
```

2. Execute the command.

The table below lists the commands available in cluster operation, indicates whether environment variables must be configured and shows the nodes on which commands are executable. For notes and cautions on command execution, refer to the Remarks column.

Table 9.17 Commands when cluster operation is used

Command	Environment variable	Command execution node	Remarks
swsrpsetvol	Required	Transaction operation node	-
swsrpvolinfo	Required	Transaction operation node	-
swsrpdelvol	Required	Transaction operation node	-
swsrprecbuffstat	Required	Transaction operation node	-
swsrprecbuffset	Required	Transaction operation node	-
swsrpstartsync	Required	Transaction operation node	-
swsrpmake	Required	Transaction operation node	-
swsrpstat	Required	Transaction operation node	-
swsrpcancel	Required	Transaction operation node	-
swsrpchsync	Required	Transaction operation node	-
swsrprevsync	Required	Transaction operation node	-
swsrprecoverres	Required	Transaction operation node	-

9.4.1 Environment definition commands

This section describes environment definition commands for replication management.

9.4.1.1 swsrpsetvol (Replication volume information setting command)

NAME

swsrpsetvol - sets the replication source volumes and replication destination volumes

SYNOPSIS

To be executed on the Managed Server

For replication on a server

```
/opt/FJSVswsrp/bin/swsrpsetvol [-n] [-u] [-Xvtoc] [-Xgroup Group-Name]
Original-Volume-Name Replica-Volume-Name
```

For replication between servers

```
/opt/FJSVswsrp/bin/swsrpsetvol [-n] [-u] [-Xvtoc] -o (ORG|REP|BOTH) [-Xgroup Group-Name]
Original-Volume-Name Replica-Volume-Name
```

To be executed on the Management Server

For replication on a server

```
/opt/FJSVswsrp/bin/swsrpsetvol -h Server-Name [-n] [-u] [-Xvtoc] [-Xgroup Group-Name] Original-
Volume-Name Replica-Volume-Name
```

For replication between servers

```
/opt/FJSVswsrp/bin/swsrpsetvol -h Server-Name [-n] [-u] -o (ORG|REP|BOTH) [-Xvtoc] [-Xgroup
Group-Name] Original-Volume-Name Replica-Volume-Name
```

DESCRIPTION

This command sets the replication source volumes and replication destination volumes that will be used in a replication operation. Execute this command to declare that the replication operation will performed using the specified replication volumes.

OPTIONS

-h *Server-Name*

Specifies the name of a Managed Server that is subject to the replication operation.
This option can be specified only on a Management Server.
If this option is omitted, the Managed Server on which the command is entered will be the target of the processing.

-u

Specifies the direction of copying.
If this option is specified, copying is performed only in the direction from replication source volume to a replication destination volume.
Copying from a replication destination volume to a replication source volume is prohibited.
If this option is omitted, copying between the replication source and replication destination volumes is bidirectional.

-o {ORG|REP|BOTH}

Specifies the operation server.
Specify one of the following for this operand:

ORG : Sets the replication source server as the operation server.
REP : Sets the replication destination server as the operation server.
BOTH : Sets both the copy-source server and copy-destination server as operation servers.

On the server that is not specified as the operation server, only the following commands can be executed for a replication volume:

- swsrpvolinfo (Replication volume information display command)

- swsrpstat (Operation status display command)

- swsrpcancel (Replication cancellation command) (The forcible stop option is required.)

Specifying BOTH in the operand enables the execution of all commands on both the copy-source server and copy-destination server.
Note that the BOTH option is usable when the versions of both servers are Version 11.0 (V11.0L10) or later.
This option is valid only for replication between servers.

-n

If this option is omitted, replication volume information cannot be set when the sizes of the replication source volume and replication destination volume are different.
When this option is specified, the sizes of the replication source volume and replication destination volume are not checked. Use this option when volume sizes are different (for example, for replication between different OSs).
When the volume sizes of the replication source volume and replication destination volume are different, copying is performed as follows:

The address at which copying starts is the first sector of the replication destination volume.

The copy data size is the size of the smaller volume.

Therefore, ensure that the size of the copy destination volume is the same size or larger than the size of the copy source volume.

This option cannot be used for SDX object logical volumes (GDS Snapshot function linkage function).

-Xgroup *Group-Name*

Specifies that the copy source and copy destination volumes will be registered in a group.

If the specified group does not exist, a new group is created. Otherwise, the specified copy source and copy destination volumes are added to an existing group.

The characters that can be used for the group name are single-byte alphanumeric characters, the minus sign (-), underscore (_), number sign (#) and period (.). The first character must be a single-byte alphanumeric character. Up to 64 characters can be used.

-Xvtoc

If this option is not specified, Solaris slices that include VTOC cannot be set as replication source volumes or replication destination volumes.

Specification of this option avoids checking to see whether or not replication source and replication destination volumes include VTOC. This option is specified for the following replication source and replication destination volumes:

- Under Solaris, VTOC-labeled disk slices that contain VTOC

- Under Solaris, a VTOC-labeled disk second slice that indicates the entire disk

This option is ignored if it is set for volumes other than the above.

OPERANDS

Original-Volume-Name

Specifies a replication source volume name.

For the volume names of other servers, use the format "*Volume-Name@Managed-Server-Name*".

Specifies the AdvancedCopy Manager device name for the *Volume-Name*.

Replica-Volume-Name

Specifies a replication destination volume name.

For the volume names of other servers, use the format "*Volume-Name@Managed-Server-Name*".

Specifies the AdvancedCopy Manager device name for the *Volume-Name*.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Set replication volume information by specifying volume/dev/sdb3 of Managed Server SRC-SV as the replication source volume and volume /dev/sdb3 of Managed Server TARG-SV as the replication destination volume. In addition, specify one-way copying for the copy direction and specify the replication source server (SRC-SV) as the operation server. The example below shows the command executed on SRC-SV:

```
# /opt/FJSVswsrp/bin/swsrpsetvol -u -o ORG /dev/sdb3 /dev/sdb3@TARG-SV
swsrpsetvol completed
#
```

- When the same settings are made from TARG-SV, the example is as follows (be careful with the volume name):

```
# /opt/FJSVswsrp/bin/swsrpsetvol -u -o ORG /dev/sdb3@SRC-SV /dev/sdb3
swsrpsetvol completed
#
```

NOTES

- Replication volume information cannot be set if:

    - Replication volume information has already been set.

    - The specified replication source volume and replication destination volume are already set as the replication destination volume and the replication source volume respectively.

    - The -n option is not specified when the sizes of the replication source volume and the replication destination volume are different.

    - Two volumes that have the same box information (ie, box identifier, OLU number, EXTENT start position, and EXTENT size) as the box information of the specified two volumes have already been set as replication volumes.

    - The replication destination volume has been registered as a backup volume by the backup management function.

    - The replication source volume has been registered as a backup volume by the backup management function and bi-directional copying is specified as a replication volume attribute.

    - The specified replication destination volume and replication source volume are volumes belonging to other servers (at least one volume must be a volume on the local server).

    - If the replication source volume and the replication destination volume are located on another ETERNUS Disk storage system, the remote copy function of the ETERNUS Disk storage system is not correctly set.

    - Communication with the Management Server fails.

    - Communication with a destination server fails for replication between servers.

    - The license for the ETERNUS Disk storage system that contains the specified replication source/target volume is one of the following:

        - The license is not registered in the license management function.

        - The license is expired.

- When the physical information on volumes (case information on partition size, a drive character, and ETERNUS Disk storage system) is changed using functions such as ETERNUS Web GUI, GDS which OS management tools (fdisk command etc.) and ETERNUS Disk storage system provides, it is necessary to re-configure the duplicate volume information associated with the changed volume, according to the following procedures:

    1. Delete the replication volume information by swsrpdelvol (Replication volume information deletion command).

    2. In the Web Console, "information acquisition / reflection" of all devices are performed. Or stgxfwcmsetdev (Device information collection / reflection command) is executed in a Management Server.

    3. Reset the replication volume information by swsrpsetvol (Replication volume information setting command).

- To copy the GPT disks by each Logical Unit (disk), the replication source and destination volumes must be the same size in the unit of bytes whether or not the replication destination volume is mounted.

- To copy the MBR disks by each Logical Unit (disk) when the replication destination volume is mounted (used), the replication source and destination volumes must be the same size in the unit of bytes.
When the replication destination volume is not mounted (not used), the size of the replication destination volume can be larger than the size of the replication source volume.

- When GDS SDX objects are used: Refer to "Configuration and conditions of unsupported SDX objects", and "10.1.5.2 Notes on SDX object operations (replication management)".

- This command cannot be executed while any of the following commands are running:

    - swsrpsetvol (Replication volume information setting command)

    - swsrpdelvol (Replication volume information deletion command)

    - swsrpstartsync (Synchronous processing start command)

    - swsrpmake (Replication creation command)

    - swsrpcancel (Replication cancellation command)

9.4.1.2 swsrpvolinfo (Replication volume information display command)

NAME

swsrpvolinfo - displays the replication volume information

SYNOPSIS

To be executed on the Managed Server

```
/opt/FJSVswsrp/bin/swsrpvolinfo [-L] [Original-Volume-Name]
```

To be executed on the Management Server

```
/opt/FJSVswsrp/bin/swsrpvolinfo -h Server-Name [-L] [Original-Volume-Name]
```

DESCRIPTION

This command displays the replication volume information registered by swsrpsetvol (Replication volume information setting command). The following information will be displayed:

Title	Description
Server	Displays a Managed Server name.
Original-Volume	Displays a replication source volume name.
Size	Displays the volume size.
Replica-Volume	Displays a replication destination volume name.
Copy	Displays the copying direction. uni-direction : One-way copying from a replication source volume to a replication destination volume is allowed. bi-direction : Bidirectional copying is allowed.
Op-Server	Displays the operation server of a pair of replication volumes. original : The replication source server is the operation server. replica : The replication destination server is the operation server. both : Sets both the copy-source server and copy-destination server as operation servers.
Group (Only when the -L option is specified)	Displays the group name. "----"is displayed if the copy source and destination volumes are not registered in a group.

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on a Management Server.

-L

Specifies that information will be displayed using the extended format.

For the items displayed by this option, refer to the table in the description above.

The items that are displayed may be expanded to support new functions provided by future versions.

OPERANDS

Original-Volume-Name

Specifies a replication source volume name.

For the volume names of other servers, use the format "Volume-Name@Managed-Server-Name."

Specifies the AdvancedCopy Manager device name for the volume name.

Displays all the replication volume information for the specified replication source volume. If this operand is omitted, all the replication volume information existing on the Managed Server on which this command is executed is displayed.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Direct that all the replication volume information on a Management Server is to be displayed on Managed Server SRC-SV:

```
# /opt/FJSVswsrp/bin/swsrpvolinfo -h SRC-SV
Server Original-Volume    Size     Replica-Volume     Size     Copy          Op-Server
SRC-SV /dev/sdb3@SRC-SV   4.0 Gbyte /dev/sdb3@TARG-SV  4.0 Gbyte uni-direction original
SRC-SV /dev/sdc3@SRC-SV   4.0 Gbyte /dev/sdc3@TARG-SV  5.0 Gbyte bi-direction  original
SRC-SV /dev/sdd3@SRC-SV   9.1 Gbyte /dev/sdd3@TARG-SV  9.1 Gbyte bi-direction  replica
SRC-SV /dev/sde3@SRC-SV   7.5 Gbyte /dev/sde3@TARG-SV  7.5 Gbyte bi-direction  both
#
```

NOTES

- Replication volume information cannot be displayed if:

    - The specified volume has not been set as a replication source volume

- This command cannot be executed while the following command is running:

    - swsrprecoverres (Resource adjustment command)

9.4.1.3 swsrpdelvol (Replication volume information deletion command)

NAME

swsrpdelvol - deletes the replication volume

SYNOPSIS

To be executed on the Managed Server

```
/opt/FJSVswsrp/bin/swsrpdelvol [-e] [-d] Original-Volume-Name Replica-Volume-Name
```

To be executed on the Management Server

```
/opt/FJSVswsrp/bin/swsrpdelvol -h Server-Name [-e] [-d] Original-Volume-Name Replica-Volume-Name
```

DESCRIPTION

This command deletes the replication volume information that has been set.
If copy processing is being performed on a replication volume that is to be deleted, execute the command after the copy processing has stopped.

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on a Management Server.

-e

Emergency mode is available for operations.
Specifies deleting replica volume information without checking the execution state of an Advanced Copy. Use this option only when it is clear that the Advanced Copy has not executed.

Deletes replication volume information without confirming the release of the master-proxy relationship of an SDX object when the function in linkage with GDS Snapshot is used.

-d

Specifies that no communication should take place between the linked servers in replication between servers. When this option is specified, replication volume information is deleted without communication taking place between the linked servers.
Use the option only when recovery of one of the linked servers cannot be expected.
The option is valid only for replication between servers.
The option can be used on the server not specified as the operation server.

OPERANDS

Original-Volume-Name

Specifies a replication source volume name.
For the volume names of other servers, use the format "*Volume-Name@Managed-Server-Name*".
Specifies the AdvancedCopy Manager device name for the *Volume-Name*.

Replica-Volume-Name

Specifies a replication destination volume name.
For the volume names of other servers, use the format "*Volume-Name@Managed-Server-Name*".
Specifies the AdvancedCopy Manager device name for the *Volume-Name*.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Delete replication volume information on Managed Server SRC-SV:

```
# /opt/FJSVswsrp/bin/swsrpdelvol /dev/sdb3 /dev/sdb3@TARG-SV
swsrpdelvol completed
#
```

NOTES

- Replication volume information cannot be deleted if:

    - The specified replication source volume and the replication destination volume are not set as replication volumes.

    - The copy processing (synchronization processing and snapshot processing) is performed between replication volumes.

    - For server-to-server replication, the Managed Server that executed the command is not the operation server of the replication volume.

    - Communication with a Management Server fails.

    - For server-to-server replication, communication with a destination server fails.

    - However, OPC sessions are neither confirmed nor cancelled by execution in emergency operation mode. If an OPC session is in the error suspended state, then unmanaged sessions remain. In this event, ETERNUS Web GUI must be used to cancel these sessions.

- When GDS SDX objects are used: Refer to "Configuration and conditions of unsupported SDX objects", and "10.1.5.2 Notes on SDX object operations (replication management)".

- This command cannot be executed while any of the following commands are running:

    - swsrpsetvol (Replication volume information setting command)

    - swsrpdelvol (Replication volume information deletion command)

    - swsrpstartsync (Synchronous processing start command)

    - swsrpmake (Replication creation command)

    - swsrpcancel (Replication cancellation command)

    - swsrpchsync (Synchronization mode change command)

    - swsrprevsync (Synchronization reverse command)

    - swsrprecoverres (Resource adjustment command)

9.4.1.4 swsrprecbuffstat (REC buffer information display command)

NAME

swsrprecbuffstat - displays information for the REC buffer

SYNOPSIS

To be executed on the Managed Server

```
/opt/FJSVswsrp/bin/swsrprecbuffstat [-Xaltpath] [-L] Volume-Name
```

To be executed on the Management Server

```
/opt/FJSVswsrp/bin/swsrprecbuffstat -h Server-Name [-Xaltpath] [-L] Volume-Name
```

DESCRIPTION

This command displays information for the REC buffer used by ETERNUS Disk storage system when REC is performed in Consistency mode.
The following information will be displayed:

Title	Description
BoxID	The box identifier for the box connected to the server where the command was executed.
ID	The ID for the REC buffer.
Use	The attribute (usage) for the REC buffer. SEND : REC buffer for sending data RECV : REC buffer for receiving data
Rate	The usage rate of the REC buffer as a percentage. **If the usage is SEND** Displays the usage rate of the REC buffer as a percentage. This will never be 0%, even if REC is not being performed in Consistency mode. **If the usage is RECV** "0%" will always be displayed. "----" will be displayed if the REC buffer cannot be used.
RemoteBoxID	The box identifier for the box performing a remote copy.
DiskBuff	Displays the REC Disk buffer status only when -L option is specified: - --- : REC buffer is used for receiving - active : REC Disk buffer is available for use - invalid : REC Disk buffer is not set - warning(code=xx) : REC Disk buffer is available for use but there is a disk fault - inactive(code=xx) : REC Disk buffer cannot be used - not_support : REC Disk buffer is not supported The following codes are displayed at code=xx in the case of warning or inactive: - 04 : Rebuild or other recovery operation is in progress (REC Disk buffer is available for use) - 05 : RAID that comprises the REC Disk buffer has a fault - 06 : The REC Disk buffer RAID group is being formatted - 07 : Copy destination does not support REC Disk buffer - 08 : A RAID that comprises the REC Disk buffer has no redundancy
DiskRate	Displays the busy rate as a percentage if the REC Disk buffer state is "active" or "warning" only when -L option is specified. If the REC Disk buffer is not being used, 0% is displayed. In cases other than the above, "---" is displayed.

OPTIONS

-h *Server-Name*

Specifies the name of a Managed Server.
This option can be specified only on a Management Server.

-Xaltpath

When replication volume is applying per logic volume by the SDX object, an alternative path is specified as an operand in the form of the following.

/dev/sfdsk/<*class name*>/dsk/<*volume name*>:<*DEVNAM value for sdxinfo*>

or
/dev/sfdsk/<*class name*>/dsk/<*volume name*>:<*by-id name of DEVNAM value for sdxinfo*>

or

/dev/sfdsk/<*class name*>/dsk/<*volume name*>:<*by-path name of DEVNAM value for sdxinfo*>

-L

This option specifies to display in extended format.

For the items displayed by this option, refer to the table in the description above.

Note that additional items corresponding to new functions in future versions might be displayed in future when this option is specified.

OPERANDS

Volume-Name

Specifies the name of a volume that has been registered as a replication volume.

You can only specify volumes that are connected to the Managed Server where the command is executed.

When replication volume is applying per logic volume by the SDX object, it is necessary to use the -Xaltpath option and to specify an alternative path.

Specifies the AdvancedCopy Manager device name as the volume name.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Display information for the REC buffer used by the REC, and replication is performed within the Managed Server "Server1":

```
# /opt/FJSVswsrp/bin/swsrprecbuffstat /dev/sda1@Server1 -L
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use  Rate RemoteBoxID                              DiskBuff DiskRate
1  SEND 25%  BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB active   50%
# /opt/FJSVswsrp/bin/swsrprecbuffstat /dev/sdb1@Server1
BoxID = BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
ID Use  Rate RemoteBoxID                              DiskBuff DiskRate
1  RECV 0%   AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA ---      ---
#
```

- Display information for the REC buffer used by the REC in an environment where replication is being performed between Managed Server "Server1" and Managed Server "Server2":

[Command executed on Managed Server "Server1"]

```
# /opt/FJSVswsrp/bin/swsrprecbuffstat /dev/sda1@Server1
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use  Rate RemoteBoxID                              DiskBuff DiskRate
1  SEND 25%  BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB active   0%
2  RECV 0%   BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB ---      ---
#
```

[Command executed on Managed Server "Server2"]

```
# /opt/FJSVswsrp/bin/swsrprecbuffstat /dev/sdk1@Server2 -L
BoxID = BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
ID Use  Rate RemoteBoxID                              DiskBuff DiskRate
1  RECV 0%   AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA ---      ---
2  SEND 25%  AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA active   0%
#
```

- Display information for the REC buffer used by the REC if a box where the copy source volume is located has been set up to make remote copies to multiple boxes, and replication is being performed between Managed Server "Server1" and Managed Server "Server2":

[Command executed on Managed Server "Server1"]

```
# /opt/FJSVswsrp/bin/swsrprecbuffstat /dev/sda1@Server1 -L
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use  Rate RemoteBoxID                              DiskBuff          DiskRate
1  SEND 25%  BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB warning(code=08)  0%
2  RECV 0%   BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB ---               ---
3  SEND 12%  CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC active            50%
4  RECV 0%   CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC ---               ---
5  SEND 12%  DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD invalid           ---
6  RECV 0%   EEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEE ---               ---
7  SEND 12%  FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF inactive(code=05) ---
8  SEND 12%  GGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGG active            0%
#
```

- Display information for the REC buffer used by the REC, and replication is being performed between Managed Server "Server1" and Managed Server "Server2":
  (The SDX object is in a two-box mirror configuration.)

[Command executed on Managed Server "Server1"]

```
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class1/dsk/volume1:sda1
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use  Rate RemoteBoxID
1  SEND 25%  BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
2  RECV 0%   BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class1/dsk/volume1:sdb1
BoxID = CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
ID Use  Rate RemoteBoxID
1  SEND 12%  DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD
2  RECV 0%   DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD
#
```

[Command executed on Managed Server "Server2"]

```
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class2/dsk/volume1:sdk1
BoxID = BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
ID Use  Rate RemoteBoxID
1  RECV 0%   AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
2  SEND 12%  AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class2/dsk/volume1:sdl1
BoxID = DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD
ID Use  Rate RemoteBoxID
1  RECV 0%   CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
2  SEND 12%  CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
#
```

NOTES

- Information is displayed only for boxes that contain replication volumes for the server where the command was executed.

  If inter-box replication is being performed between servers, then the command must be executed on both servers involved in the replication in order to check information for the REC buffers in both boxes.

- It will take a short time for changes to REC buffer settings to take effect, during which time the buffer may become unusable (the "Rate" field will be displayed as "----"). In this case, wait about one minute and then execute the command again.

- This command cannot be executed while the following command is running:

    - swsrprecoverres (Resource adjustment command)

9.4.1.5 swsrprecbuffset (REC buffer settings change command)

NAME

swsrprecbuffset - changes the settings of the REC buffer

SYNOPSIS

To be executed on the Managed Server

```
/opt/FJSVswsrp/bin/swsrprecbuffset -Xbuffid buffer-id -Xbuffuse {SEND|RECV} [-Xaltpath] Volume-Name
```

To be executed on the Management Server

```
/opt/FJSVswsrp/bin/swsrprecbuffset -h Server-Name -Xbuffid buffer-id -Xbuffuse {SEND|RECV} [-Xaltpath] Volume-Name
```

DESCRIPTION

This command changes the settings of the REC buffer used by ETERNUS Disk storage system when REC is performed in Consistency mode.
For information about the setting method of REC buffer, refer to "User's Guide Setup / Maintenance" of ETERNUS Web GUI.

OPTIONS

-h *Server-Name*

Specifies the name of a Managed Server.
This option can be specified only on a Management Server.

-Xbuffid *buffer-id*

Specifies the ID for the REC buffer.

-Xbuffuse {SEND|RECV}

Specifies the usage for the REC buffer.

SEND : REC buffer for sending data
RECV : REC buffer for receiving data

-Xaltpath

When replication volume is a logical volume in the SDX object, an alternative path is specified as an operand in the following form.

/dev/sfdsk/<*class name*>/dsk/<*volume name*>:<*DEVNAM value for sdxinfo*>

or

/dev/sfdsk/<*class name*>/dsk/<*volume name*>:<*by-id name of DEVNAM value for sdxinfo*>

or

/dev/sfdsk/<*class name*>/dsk/<*volume name*>:<*by-path name of DEVNAM value for sdxinfo*>

OPERANDS

Volume-Name

Specifies the name of a volume that has been registered as a replication volume.
You can only specify volumes that are connected to the Managed Server where the command is executed.

When replication volume is applying per logic volume by the SDX object, it is necessary to use the -Xaltpath option and to specify an alternative path.

Specifies the AdvancedCopy Manager device name as the volume name.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Change the attributes for the REC buffer used by the REC, and replication is being performed within the Managed Server "Server1":

```
# /opt/FJSVswsrp/bin/swsrprecbuffstat /dev/sda1@Server1
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use  Rate RemoteBoxID
1  SEND 25%  BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
# /opt/FJSVswsrp/bin/swsrprecbuffset -Xbuffid 1 -Xbuffuse RECV /dev/sda1@Server1
swsrprecbuffset completed
# /opt/FJSVswsrp/bin/swsrprecbuffstat /dev/sda1@Server1
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use  Rate RemoteBoxID
1  RECV 0%   BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
# /opt/FJSVswsrp/bin/swsrprecbuffstat /dev/sdb1@Server1
BoxID = BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
ID Use  Rate RemoteBoxID
1  RECV 0%   AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
# /opt/FJSVswsrp/bin/swsrprecbuffset -Xbuffid 1 -Xbuffuse SEND /dev/sdb1@Server1
swsrprecbuffset completed
# /opt/FJSVswsrp/bin/swsrprecbuffstat /dev/sdb1@Server1
BoxID = BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
ID Use  Rate RemoteBoxID
1  SEND 12%  AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
#
```

- Change the connected box for the REC buffer of the copy source volume , and replication is being performed between Managed Server "Server1" and Managed Server "Server2":

[Command executed on Managed Server "Server1"]

```
# /opt/FJSVswsrp/bin/swsrprecbuffstat /dev/sda1@Server1
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use  Rate RemoteBoxID
1  SEND 25%  BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
2  RECV 0%   BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
# /opt/FJSVswsrp/bin/swsrprecbuffset -Xbuffid 1 -Xbuffuse SEND /dev/sdb1@Server1
swsrprecbuffset completed
# /opt/FJSVswsrp/bin/swsrprecbuffset -Xbuffid 2 -Xbuffuse RECV /dev/sdb1@Server1
swsrprecbuffset completed
# /opt/FJSVswsrp/bin/swsrprecbuffstat /dev/sdb1@Server1
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use  Rate RemoteBoxID
1  SEND 12%  BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
2  RECV 0%   BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
#
```

- Change the attributes for the REC buffer used by the REC, and replication is being performed between Managed Server "Server1" and Managed Server "Server2":

(The SDX object is in a two-box mirror configuration.)

[Command executed on Managed Server "Server1"]

```
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class1/dsk/volume1:sda1
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use  Rate RemoteBoxID
1  SEND 25%  BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class1/dsk/volume1:sdb1
BoxID = CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
ID Use  Rate RemoteBoxID
1  SEND 12%  DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD
# /opt/FJSVswsrp/bin/swsrprecbuffset -Xbuffid 1 -Xbuffuse RECV -Xaltpath /dev/sfdsk/class1/dsk/
volume1:sda1
swsrprecbuffset completed
# /opt/FJSVswsrp/bin/swsrprecbuffset -Xbuffid 1 -Xbuffuse RECV -Xaltpath /dev/sfdsk/class1/dsk/
volume1:sdb1
swsrprecbuffset completed
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class1/dsk/volume1:sda1
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use  Rate RemoteBoxID
1  RECV 0%   BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class1/dsk/volume1:sdb1
BoxID = CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
ID Use  Rate RemoteBoxID
1  RECV 0%   DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD
```

[Command executed on Managed Server "Server2"]

```
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class1/dsk/volume1:sdk1
BoxID = BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
ID Use  Rate RemoteBoxID
1  RECV 0%   AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class2/dsk/volume1:sdl1
BoxID = DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD
ID Use  Rate RemoteBoxID
1  RECV 0%   CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
# /opt/FJSVswsrp/bin/swsrprecbuffset -Xbuffid 1 -Xbuffuse SEND -Xaltpath /dev/sfdsk/class1/dsk/
volume1:sdk1
swsrprecbuffset completed
# /opt/FJSVswsrp/bin/swsrprecbuffset -Xbuffid 1 -Xbuffuse SEND -Xaltpath /dev/sfdsk/class2/dsk/
volume1:sdl1
swsrprecbuffset completed
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class1/dsk/volume1:sdk1
BoxID = BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
ID Use  Rate RemoteBoxID
1  SEND 12%  AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class2/dsk/volume1:sdl1
BoxID = DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD
ID Use  Rate RemoteBoxID
1  SEND 12%  CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
#
```

NOTES

- The device name on a box to display information can be specified as -Xaltpath.

- Settings for dedicated REC buffers must be made in advance.
  Change the settings while there are no Consistency mode sessions in either the ACTIVE or ERROR SUSPEND state.

- It will take a short time for changes to buffer settings to take effect, during which time the REC buffer will become unusable. Do not perform REC operations during this time.

- This command cannot be executed while any of the following commands are running:

    - swsrpsetvol (Replication volume information setting command)

    - swsrpdelvol (Replication volume information deletion command)

    - swsrpstartsync (Synchronous processing start command)

    - swsrpmake (Replication creation command)

    - swsrpstat (Operation status display command)

    - swsrpcancel (Replication cancellation command)

    - swsrpchsync (Synchronization mode change command)

    - swsrprevsync (Synchronization reverse command)

    - swsrprecoverres (Resource adjustment command)

9.4.2 Operation commands

This section describes operation commands for replication management.

9.4.2.1 swsrpstartsync (Synchronous processing start command)

NAME

swsrpstartsync - starts synchronous processing

SYNOPSIS

To be executed on the Managed Server

```
/opt/FJSVswsrp/bin/swsrpstartsync [[-y [-a | -v]| -k | -i] [-g] [-m] [-Xskip | -Xremain] | -Xgds-
softcopy | -Xgds-selectcopy ] [-t] [-Xda] [-Xgds-slice-skipchk] From-Volume-Name To-Volume-Name
```

To be executed on the Management Server

```
/opt/FJSVswsrp/bin/swsrpstartsync -h Server-Name [[-y [-a | -v]| -k | -i] [-g] [-m] [-Xskip | -
Xremain] | -Xgds-softcopy | -Xgds-selectcopy ] [-t] [-Xda] [-Xgds-slice-skipchk] From-Volume-Name
To-Volume-Name
```

To be executed on the Managed Server: group specification

```
/opt/FJSVswsrp/bin/swsrpstartsync [-y [-a | -v]| -k | -i] [-g] [-m] [-Xskip | -Xremain] [-t] [-Xda]
[-Xgds-slice-skipchk] [-Xreverse] -Xgroup Group-Name
```

To be executed on the Management Server: group specification

```
/opt/FJSVswsrp/bin/swsrpstartsync -h Server-Name [-y [-a | -v]| -k | -i] [-g] [-m] [-Xskip | -
Xremain] [-t] [-Xda] [-Xgds-slice-skipchk] [-Xreverse] -Xgroup Group-Name
```

DESCRIPTION

Start the synchronous replication processing (copying between disks by EC and REC).
It will copy all the data from the source volume to the destination volume at the start of the synchronization process, or if run as a restart of the synchronization process it will just copy updated data to the destination volume. Use this command at the start of operation or when you want to reflect updated data (differences) from the equivalency maintenance state.

For a GDS transaction volume, synchronous processing using a software copy is also possible.

OPTIONS

-h *Server-Name*

    This option is valid only on a Management Server.
    This option can be specified only on a Management Server.

-y

    Specifies inter-box synchronization in synchronous mode. This option is valid only when total copy is started for inter-box replication.
    Specify this option together with the -a option to start the split mode for inter-box synchronization in Automatic Split mode.
    Specify this option together with the -v option to start the mode in Manual Split mode.
    Specifying the -y option without the -a or -v option starts the mode in Automatic Split mode.
    Not specifying the -y, -k, and -i options sets Through mode.
    Synchronization in intra-box replication is always handled in synchronous mode.

-a

    Specifies the start of split mode for inter-box synchronization in Automatic Split mode. This option is valid only when total copy is started for inter-box replication.
    Specify this option together with the -y option.

    You cannot use this option when operating SDX objects in units of logical volumes.

-v

    Specifies the start of split mode for inter-box synchronization in Manual Split mode. This option is valid only when total copy is started for inter-box replication.
    Specify this option together with the -y option.
    You must specify this option when specifying synchronization mode (the -y option) on a box that does not support Automatic Split mode.

    You cannot use this option when operating SDX objects in units of logical volumes.

-k

    Specifies the setting of transfer mode for inter-box synchronization in Stack mode. This option is valid only when total copy is started for inter-box replication.
    Not specifying the -y, -k, or -i option sets Through mode.

    You cannot use this option when operating SDX objects in units of logical volumes.

-i

    Specifies the setting of transfer mode for inter-box synchronization in Consistency mode. This option is valid only when total copy is started for inter-box replication.
    Not specifying the -y, -k, or -i option sets Through mode.

    You cannot use this option when operating SDX objects in units of logical volumes.

-g

    Specifies the setting of recovery mode for inter-box synchronization in Manual Recovery mode. This option is valid only when total copy is started for inter-box replication.
    Not specifying this option sets Automatic Recovery mode.

    You cannot use this option when operating SDX objects in units of logical volumes.

-m

    For server-to-server replication, this option specifies that communication processing with non-operational servers is not performed.
    When this option is specified, volume status check processing and the pre-processing for the replication destination volume of a non-operational server are not performed.
    This option is valid only during server-to-server replication.

-t

    Specifies that pre-processing for the replication destination volume is not performed.
    Use this option if any of the following conditions apply:

- AdvancedCopy Manager does not perform pre-processing or post-processing for the replication destination volume, and independent pre-processes and post-processes are inserted before and after AdvancedCopy Manager commands.

- Replication destination pre-processing and post-processing are judged to be unnecessary (for example, for databases constructed on RAW devices).

- Synchronous processing (REC) between boxes is restarted after a forcible suspend.

-Xgds-softcopy

Specifies the start of synchronous processing of an SDX object and use of the software-copy function instead of the Advanced Copy function.
If you specify this option to start differential copying when the Advanced Copy function is being used, the Advanced Copy processing stops and software-copy processing starts.
You can specify the option when the function is used with GDS Snapshot is being used.
The option is applicable only to SDX objects.
The option cannot be specified together with -y or -Xgds-selectcopy option.

-Xgds-selectcopy

Specifies the start of synchronous processing without a specified copy type, such as copying with the ETERNUS Disk storage system (OPC/EC) or software-copy. GDS selects the type of copying to be started.
You can specify the option when the function in linkage with GDS Snapshot is being used.
The option is applicable only to SDX objects.
The option cannot be specified together with -Xgds-softcopy option.

-Xgroup *Group-Name*

Specifies that synchronous processing must be started for each group.

-Xreverse

Specifies that synchronous processing must be started for each group in the reverse direction (from replica volumes to original volumes).

-Xskip

Specifies that synchronous processing must be suspended using the initial copy skip function. This option is valid only when total copy is started for inter-box replication.

This option is not supported when GDS Snapshot linkage is being used.

-Xremain

Enables data that has been updated to the copy destination during replication established status when synchronous processing is restarted.
This option is valid only when inter-box replication is restarted.
This option is used to reflect to the copy destination volume only those parts of the copy source volume that have been updated, after synchronous processing has been started using the -Xskip option and backup data has been copied from the tape medium to the copy destination volume.

This option is not supported when GDS Snapshot linkage is being used.

-Xgds-slice-skipchk

Skips the SDX object status check processing that is usually performed as part of the copy destination pre-processing if SDX objects are operated in slice units.
If this option is specified for a stopped volume, the -t option must also be specified.

-Xda

Disables update (WRITE) and access (READ) to the replication destination volume from the synchronous processing start until the replication has been completed.
This option is valid only when synchronous processing is started.
If the option is omitted, only update (WRITE) to the replication destination volume is disabled.

You cannot use this option when operating SDX objects in units of logical volumes.

OPERANDS

From-Volume-Name

Specifies a replication source volume name.

Specifies the replication source volume or the replication destination volume set with swsrpsetvol (Replication volume information setting command).

Specifies the AdvancedCopy Manager device name for the volume name.

To-Volume-Name

Specifies a replication destination volume name.

Specifies the replication source volume or the replication destination volume set with swsrpsetvol (Replication volume information setting command).

Specifies the AdvancedCopy Manager device name for the volume name.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Start replication from /dev/sdb3 of a Management Server to /dev/sdb3@TARG-SV of Managed Server SRC-SV:

```
# /opt/FJSVswsrp/bin/swsrpstartsync -h SRC-SV /dev/sdb3 /dev/sdb3@TARG-SV
FROM=/dev/sdb3@SRC-SV,TO=/dev/sdb3@TARG-SV swsrpstartsync completed
#
```

- Start replication from /dev/sdg of a Management Server to /dev/sdg@TARG-SV of a Managed Server SRC-SV with replication destination volume access disabled.

```
# /opt/FJSVswsrp/bin/swsrpstartsync -h SRC-SV -Xda /dev/sdg /dev/sdg@TARG-SV
FROM=/dev/sdg@SRC-SV,TO=/dev/sdg@TARG-SV  swsrpstartsync  completed
#
```

- Start replication for group "GRP1":

```
# /opt/FJSVswsrp/bin/swsrpstartsync -Xgroup GRP1
GROUP=GRP1 swsrpstartsync completed
#
```

NOTES

- Replication cannot be started if:

    - The specified source volume and destination volume have not been set as replication volumes.

    - A replication source volume is specified as the copy volume and a replication destination volume is specified as the copied material volume, and one-way copy is set as the copy attribute of the replication volume.

    - Volume information (box identifier, OLU, EXTENT start position, and EXTENT size) is changed after operation starts. In this case, the operation cannot continue for the corresponding volume. Stop the replication operation for the corresponding volume, and then delete and reset the replication volume information.

    - For server-to-server replication, the Managed Server on which the command is executed is not the operation server of the replication volume.

    - For server-to-server replication, communication with a destination server fails.

    - Replication cannot be started as a result of the preceding Advanced Copy function.

    - The -g option is specified, except for a remote copy.

- - The -Xgds-softcopy or -Xgds-selectcopy option is specified when the function in linkage with GDS Snapshot is not being used.

- - The -a, -v, -k, -i or -g option is specified when the function in linkage with GDS Snapshot is being used.

- - The backup management function is executing a process for a specified source/destination volume

- - Either the copy source volume or the copy destination volume is a Snap Data Volume.

- If the -Xgroup option is specified, replication volumes are processed in order, but processing will be terminated immediately if an error occurs.

- When the replication to be performed satisfies either of the following two conditions, please perform unmount of the copy target volume before executing a command.

    - When a replication is performed between different OSs

    - When the sizes of replication source volume and replication destination volume differ

    - When the file systems of replication source volume and replication destination volume differ

  In the above mentioned cases, if dismount have been performed, post-processing of the copy target volume at the time of duplicate creation results in an error.

- You cannot change synchronization mode when synchronization is resumed.
  When synchronization is resumed, you must specify the -y, -a, -v, -k, -i, and -g options as follows:

    - Specify only the options that match the mode of synchronization to be resumed, or

    - Do not specify any of the options.

- The processing of this command varies according to the status of the synchronization processing of the specified replication source and replication destination volumes.

Table 9.18 Differences in processing corresponding to synchronization status

Status of synchronization processing	Processing
Dismount status	Start of total copy
Total or incremental copy status	Informational message (swsrp2401) is output and the command ends normally
Equivalency maintain status	Informational message (swsrp2401) is output and the command ends normally
Replication established status (i.e., halt status)	Start of incremental copy

- Before starting the replication, run the replication pre-processing script for the replication destination volume. The contents of this script can be customized. For more information about this, refer to "Appendix C Pre-processing and Post-processing of Replication". If you do not want to implement the pre-processing script, use the -t option.

  If replication is started with the -Xgroup option specified, the replication pre-processing script is not executed for the copy destination volume. Therefore, pre-processing must be executed for all copy destination volumes in the group before this command is executed.

- Before starting replication that targets a Logical Unit (disk), execute pre-processing and post-processing of the Logical Unit (disk), refer to "Appendix C Pre-processing and Post-processing of Replication" for details.

- To protect processing that accesses a replication destination volume, set the replication destination volume before starting replication so that other processes cannot access it. To do so, execute a dismount command. If you cannot execute dismount resulting from any of the conditions described below, this command fails and terminates.

    - There is a directory in which other volumes are mounted under the mount point.

    - A file in the volume is being used. In addition, if the replication destination volume is being used on a server other than the server on which this command is entered, the replication-operation administrator must disable access to the replication destination volume by cancellation of the share setting, unmounting, or by some other method.

- When GDS SDX objects are used: Refer to "Configuration and conditions of unsupported SDX objects" and "10.1.5.2 Notes on SDX object operations (replication management)".

- Refer to "10.1.1 General notes" for notes about starting replication.

- An error will occur if the synchronous process is restarted with -Xda option. When changing the access permission for replication destination volume, please first cancel the current replication processing with swsrpcancel (Replication cancellation command) and then restart this command.

- This command will terminate abnormally with the following error message if the -Xda option is set for an ETERNUS Disk storage system that does not support the Destination Access Permission function.

```
swsrp0719 EC cannot be run. Error=17(not support.)
```

- This command cannot be executed while any of the following commands are running:

    - swsrpsetvol (Replication volume information setting command)

    - swsrpdelvol (Replication volume information deletion command)

    - swsrprecoverres (Resource adjustment command)

9.4.2.2 swsrpmake (Replication creation command)

NAME

swsrpmake - creates a replication volume

SYNOPSIS

To be executed on the Managed Server

Synchronized replication

```
/opt/FJSVswsrp/bin/swsrpmake [-m] [-f] [-t] ([-j | -Xconcur]) [-Xgds-slice-skipchk] From-Volume-
Name To-Volume-Name
```

Snapshot replication

```
/opt/FJSVswsrp/bin/swsrpmake [-m] [-f] [-t] ([-T | -C | -P]) [-Xgds-slice-skipchk] From-Volume-
Name To-Volume-Name
```

To be executed on the Management Server

Synchronized replication

```
/opt/FJSVswsrp/bin/swsrpmake -h Server-Name [-m] [-f] [-t] ([-j | -Xconcur]) [-Xgds-slice-
skipchk] From-Volume-Name To-Volume-Name
```

Snapshot replication

```
/opt/FJSVswsrp/bin/swsrpmake -h Server-Name [-m] [-f] [-t] ([-T | -C | -P]) [-Xgds-slice-
skipchk] From-Volume-Name To-Volume-Name
```

To be executed on the Managed Server: group specification

Synchronized replication

```
/opt/FJSVswsrp/bin/swsrpmake [-m] [-f] [-t] ([-j | -Xconcur]) [-Xreverse] -Xgroup Group-Name
```

Snapshot replication

```
/opt/FJSVswsrp/bin/swsrpmake [-m] [-f] [-t] ([-T | -C | -P]) [-Xconcur-opc] [-Xreverse] -Xgroup
Group-Name
```

To be executed on the Management Server: group specification

Synchronized replication

```
/opt/FJSVswsrp/bin/swsrpmake -h Server-Name [-m] [-f] [-t] ([-j | -Xconcur]) [-Xreverse] -Xgroup
Group-Name
```

Snapshot replication

```
/opt/FJSVswsrp/bin/swsrpmake -h Server-Name [-m] [-f] [-t] ([-T | -C | -P]) [-Xconcur-opc] [-
Xreverse] -Xgroup Group-Name
```

DESCRIPTION

When synchronization processing (EC) is not performed, this command will start snapshot processing (OPC/QuickOPC/SnapOPC/SnapOPC+) to create a replication volume.

When synchronization processing is performed, check whether the status of the processing is in the equivalency maintenance state. If so, suspend the synchronization processing and create a replication volume. The status in which synchronization processing is suspended is referred to as the replication established status. After this command is executed, the replication destination volume can be accessed.

The behavior of this command could differ depending on the status of the Advanced Copy during progress and the following should be noted:

- When synchronous processing is not performed, the command starts snapshot processing and creates replication volumes. This form of replication is called snapshot replication.

- When synchronous processing is performed, the command checks the state of synchronous processing and, if the synchronous processing is in the equivalency maintenance state, the command suspends the processing in order to create a replication volume. This form of replication is called synchronous replication. The state in which synchronous processing is suspended is called the copy established state.

In either case, the copy destination volume is made accessible after this command is executed.

This command executes the replication pre-processing and post-processing before a replication is created. For details of these processes, refer to "Appendix C Pre-processing and Post-processing of Replication".

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on a Management Server.

-m

For server-to-server replication, this specifies that communication processing is not performed with the non-operation server. When this option is specified, volume status check processing and pre-processing and post-processing are not performed for the non-operation server volume.
This option is valid only during server-to-server replication.

-f

Specifies that the pre-processing and post-processing are not performed for a replication source volume.
Use this option only when AdvancedCopy Manager does not execute pre-processing and post-processing of the copy source volume and unique pre-processing and post-processing scripts are inserted before and after (respectively) the AdvancedCopy Manager command, or when you judge that copy pre-processing and post-processing are unnecessary (e.g., for a database constructed on a raw device).

-t

Specifies that the pre-processing and post-processing are not performed for a replication destination volume.
Use this option only when AdvancedCopy Manager does not execute pre-processing and post-processing of the copy destination volume and unique pre-processing and post-processing scripts are inserted before and after (respectively) the AdvancedCopy Manager command, or when you judge that copy pre-processing and post-processing are unnecessary (e.g., for a database constructed on a raw device).

-j

Specifies that copy processing will be forcibly suspended even if the execution status of copying is "sync: Total copy or incremental copy is being performed" or "halt: Hard suspend status" in synchronization processing of intra-box replication. This option can only be used for intra-box replication; if attempted for intra-box replication, an error will occur.

When this option is specified, the data in the destination volume is not guaranteed.

Additionally, when this option is specified, pre-processing and post-processing for the source and destination volumes are not performed. When restarting synchronous processing (REC) between boxes, execute swsrpstartsync (Synchronous processing start command) with the -t option specified.

-T

Specifies that differential snapshot processing will be executed.

This option cannot be specified while synchronous processing is performed.

This option is valid only for a replication within the ETERNUS Disk storage system and when the ETERNUS Disk storage system supports the QuickOPC function.

For a snapshot replication without specifying this option, ordinary snapshot processing (OPC without using the QuickOPC function) is started. This option must be specified to perform the replication operation using differential snapshots.

-Xgroup *Group-Name*

Specifies that a replication will be created for each group.

-Xreverse

Specifies that a replication will be created for each group in the reverse direction (i.e., from replica volumes to original volumes).

-Xconcur

Specifies that a replication will be created using the concurrent suspend function.

This option is not supported when GDS Snapshot linkage is being used.

Xconcur-opc

Specifies that a replication will be created using the concurrent OPC function.

This option is not supported when linking with GDS Snapshot.

-C

Specifies that SnapOPC will be started.

This option is only valid for intra-box replications and for ETERNUS Disk storage system that supports the SnapOPC function.

In order to execute SnapOPC, the copy destination volume must be a Snap Data Volume.

-P

Specifies to start SnapOPC+.

This option is only valid with intra-box replications and for ETERNUS Disk storage system that supports the SnapOPC+ function.

In order to execute SnapOPC+, the copy destination volume must be a Snap Data Volume.

An error occurs in the following cases:

- The same SnapOPC+ session exists at the copy source volume and the copy destination volume.

-Xgds-slice-skipchk

Skips the SDX object status check processing that is usually performed as part of the copy source and copy destination pre-processing if SDX objects are operated in slice units.

If this option is specified for a stopped volume, the -f and -t options must also be specified.

OPERANDS

From-Volume-Name

Specifies a replication source volume.

Specifies the replication source volume or the replication destination volume that was set with swsrpsetvol (Replication volume information setting command).

Specifies the AdvancedCopy Manager device name for the volume name.

To-Volume-Name

Specifies a replication destination volume.

Specifies the replication source volume or the replication destination volume that was set with swsrpsetvol (Replication volume

information setting command).

Specifies the AdvancedCopy Manager device name for the volume name.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- The Management Server instructs Managed Server: SRC-SV to create a copy of /dev/sdb3 in /dev/sdb3@TARG-SV:

```
# /opt/FJSVswsrp/bin/swsrpmake -h SRC-SV /dev/sdb3 /dev/sdb3@TARG-SV
FROM=/dev/sdb3@SRC-SV,TO=/dev/sdb3@TARG-SV swsrpmake completed
#
```

- Creates a copy for the group "GRP1":

```
# /opt/FJSVswsrp/bin/swsrpmake -Xgroup GRP1
GROUP=GRP1 swsrpmake completed
#
```

- Creates a copy for the group "GRP1" by using the concurrent OPC:

```
# /opt/FJSVswsrp/bin/swsrpmake -Xgroup GRP1 -Xconcur-opc
GROUP=GRP1 swsrpmake completed
#
```

NOTES

- When the -bind option in the mount command is specified to mount part of the file hierarchy of the copy source or copy destination volume at another mount point, use backup manager's authority to set the copy source or copy destination volume to the unmounted state.

- Do not execute the replication creation command in a state where another volume is mounted at the mount point of the target volume. For such a state, unmount the other volume before executing the replication creation command.

- Replication cannot be created if:

    - The specified replication source and replication destination volumes are not set as replication volumes.

    - The synchronization processing from the specified replication source volume to the replication destination volume is not in the equivalency maintenance status.

    - After operation starts, physical volume information (box identifier, OLU, EXTENT start position, and EXTENT size) is changed. In this case, operation for the corresponding volume cannot continue. Stop the replication operation for the volume, and then delete and reset the replication volume information.

    - A destination volume has multiple mount points.

    - If a replication source volume is an SDX object (SynfinityDISK), the status of the SDX object does not match the following:

        - The mirror volume is ACTIVE or STOP.

        - The mirror partition status is ACTIVE or TEMP.
          When the mirror partition is TEMP, a reproduction former volume is not preprocessed.

        - The SDX disk status is ENABLE.

    - For server-to-server replication, the Managed Server on which the command is executed is not the operation server of the replication volume.

    - For server-to-server replication, communication with a destination server fails.

- If the -j, -T, or -C option is specified when the function is used with GDS Snapshot.

- The backup management function is executing a process for a specified source or destination volume

- From the copy source server, the command was executed with both -Xconcur-opc and -m options specified simultaneously for any of the following ETERNUS Disk storage systems:

  - ETERNUS DX80/DX80 S2 (for firmware version numbers before V10L30)

  - ETERNUS DX90/DX90 S2 (for firmware version numbers before V10L30)

  - ETERNUS DX400/DX400 S2 series (for firmware version numbers before V10L30)

  - ETERNUS DX8000/DX8000 S2 series (for firmware version numbers before V10L30)

  - ETERNUS2000

  - ETERNUS4000

  - ETERNUS8000

- If the -Xgroup option is specified, processing will be terminated without creating a copy if the operation statuses of the replication volumes being processed include both snapshot-type replication and synchronous-type replication. Additionally, if the -Xgroup option is specified, replication volumes are processed in order, but processing will be terminated immediately if an error occurs.

- Note the following when you specify the -Xconcur-opc option:

  - Specify this option together with the -Xgroup option.

  - If an error is detected by the preliminary check for the replication volumes to be processed, processing will be terminated without creating a copy.
    Refer to the following table for details of the preliminary checks:

Type	Details of the preliminary checks
Common	Check if a pair of synchronized replications is not exist in the group.
When the - T option is specified	Check if a pair of volumes which is in the tracking state and a pair of volumes on which the tracking is not executed are not mixed in the group.

  - If this command is executed during the snapshot processing, start or restart the snapshot processing. Refer to the following table for details of copy operation:

OPC/SnapOPC/SnapOPC+

Operation status (Display by swsrpstat -L)	Copy operation
When the operation status of the replication volume in the group is in one of the following: - The concurrent OPC is defined. (Status is "readying".) - Copy is not executed. (Status is "----" and Trk is "----".) - Snapshot processing is in progress and tracking is not executed. (Status is "snap" and Trk is "off".) - Snapshot processing is in progress and tracking is in progress. (Status is "snap" and Trk is "on".) - Tracking is in progress. (Status is "----" and Trk is "on".) - SnapOPC is in progress. (Status is "copy-on-write".)	1. Copying that is in progress is terminated. 2. Define the concurrent OPC. Snapshot processing is newly started with the concurrent OPC specified.t

Operation status (Display by swsrpstat -L)	Copy operation
- SnapOPC+ is in progress. (Status is "copy-on-write(active)" or "copy-on-write(inactive)".) - Error suspend status (Status is "failed".)	

QuickOPC

Operation status (Display by swsrpstat -L)	Copy operation
When the operation status of the replication volume in the group is in one of the following: - The concurrent OPC is defined. (Status is "readying".) - Copy is not executed. (Status is "----" and Trk is "----".) - Snapshot processing is in progress and tracking is not executed. (Status is "snap" and Trk is "off".) - SnapOPC is processing (Status is "copy-on-write") - SnapOPC+ is processing (Status is "copy-on-write(active)" or "copy-on-write(inactive)") - Error suspend status (Status is "failed".)	1. Copying that is in progress is terminated. 2. Define the concurrent OPC. 3. Differential snapshot processing is newly started with the concurrent OPC specified.
When the operation status of the replication volume in the group is in one of the following: - Snapshot processing is in progress and tracking is in progress. (Status is "snap" and Trk is "on".) - Tracking is in progress. (Status is "----" and Trk is "on".)	Differential snapshot processing is restarted with the concurrent OPC specified.

- When the replication to be performed satisfies either of the following conditions, please perform dismount of the copy target volume before executing a command.

    - When a replication is performed between different OSs

    - When the sizes of replication source volume and replication destination volume differ

    - When the file systems of replication source volume and replication destination volume differ

    If dismount is performed in the abovementioned case, post-processing of the copy target volume at the time of duplicate creation produces an error.

- Before a replica is created, the pre-processing for the source volume and destination volume are executed (for synchronized replication, only pre-processing is executed). After the replica has been created, the post-processing for the source volume and destination volume is executed. The contents of the script can be customized. For more information about this refer to "Appendix C Pre-processing and Post-processing of Replication". If the -f or -t option is specified, the pre-processing and post-processing are not performed.

    If a replica is created with the -Xgroup option specified, the replication, pre-processing and post-processing scripts are not executed for the copy source volume/copy destination volume. Therefore, execute pre-processing for all copy source volumes/copy destination

volumes in the group before this command is executed.
(However, for synchronous replication, pre-processing for the copy destination volume is executed when replication starts, so it need not be executed.)
In addition, post-processing must be executed after this command is executed.

- If this command is executed during execution of snapshot processing, stop the copy processing in progress and restart the snapshot processing.

- To protect the data, this command sets the replication source volume before creating a replica so that other processes cannot access the volume. The replication source volume is unmounted in order to disable to access it. If you cannot unmount the source volume because of the situations described below, this command fails and terminates.

    - There is a directory in which other volumes are mounted under the mount point.

    - A file on the volume is being used. In addition, if the replication source volume is being used from a server other than the server on which this command is entered, the backup-operation administrator must disable access to the replication source volume by cancellation of the share setting, unmounting, or by using some other method. If for whatever reason you do not want to remove the share setting of the replication source volume, or execute unmount or use some other method, refer to "Appendix C Pre-processing and Post-processing of Replication" and note the following:

        - For replication pre-processing, execute the UNIX sync command to synchronize volumes and the file system.

        - In replication post-processing, execute the <fsck> command to check the file system at the copy destination.

    - When the replication source volume is being used by backup management, dismount may not possible. In this case, "swsrp2613 An error occurred in the pre-processing script for creating a replica. Error Code = 2" is output.

- Before replication is performed in snapshot-type replication, this command option prevents other processes from accessing the replication destination volume to protect the process that is accessing the volume. The command accomplishes this by locking or dismounting the replication destination volume. In cases where the replication destination volume cannot be locked or dismounted, such as those described below, the command terminates with an error:

    - The mount point has a directory for mounting another volume.

    - The volume has a file that is being used. If the replication destination volume is being used by a server that is not the server where the command was entered, prevent the replication destination volume from being accessed by cancelling share settings or performing dismounting for the replication operation, using the Administrator's authority.

- For Snap Data Volume, only SnapOPC and SnapOPC+ can be executed.

    Additionally, when restoring data from a Snap Data Volume, note that there must be a SnapOPC session between the copy source volume and the copy destination volume (Snap Data Volume).

- If this command is executed during snapshot processing, snapshot processing is restarted.
  Refer to the following table for details of copy operation:

Table 9.19 Copy operation when command is executed during snapshot processing

Operation status (Display by swsrpstat -L)	-T option specification	Copy operation
Copy is not executed. (Status is "----" and Trk also is "----".)	No	Ordinary snapshot processing is started.
Copy is not executed. (Status is "----" and Trk also is "----".)	Yes	Differential snapshot processing is started.
Snapshot processing is in progress and tracking is not executed. (Status is "snap" and Trk is "off".)	No	Copying that is in progress is terminated and ordinary snapshot processing is restarted.
Snapshot processing is in progress and tracking is not executed. (Status is "snap" and Trk is "off".)	Yes	Copying in progress is terminated and differential snapshot processing is restarted.
Snapshot processing is in progress and tracking is in progress. (Status is "snap" and Trk is "on".)	No	Copying in progress is terminated and ordinary snapshot processing is restarted.

Operation status (Display by swsrpstat -L)	-T option specification	Copy operation
Snapshot processing is in progress and tracking is in progress. (Status is "snap" and Trk is "on".)	Yes	Differential snapshot processing is restarted.
Tracking is in progress. (Status is "----" and Trk is "on".)	No	Tracking processing is terminated and ordinary snapshot processing is restarted.
Tracking is in progress. (Status is "----" and Trk is "on".)	Yes	Differential snapshot processing is restarted.

- When GDS SDX objects are used: Refer to "Configuration and conditions of unsupported SDX objects", and "10.1.5.2 Notes on SDX object operations (replication management)".

- Refer to "10.1.1 General notes" for notes on creating replication.

- This command cannot be executed while any of the following commands are running:

  - swsrpsetvol (Replication volume information setting command)

  - swsrpdelvol (Replication volume information deletion command)

  - swsrprecoverres (Resource adjustment command)

9.4.2.3 swsrpstat (Operation status display command)

NAME

swsrpstat - displays the operation status

SYNOPSIS

To be executed on the Managed Server

```
/opt/FJSVswsrp/bin/swsrpstat [ [-L] [-Xdate] [-Xda] | [-Xall] ] [-E] [-H] [-O] [Original-Volume-
Name]
```

To be executed on the Management Server

```
/opt/FJSVswsrp/bin/swsrpstat -h Server-Name [ [-L] [-Xdate] [-Xda] | [-Xall] ] [-E] [-H] [-O]
[Original-Volume-Name]
```

To be executed on the Managed Server: group specification

```
/opt/FJSVswsrp/bin/swsrpstat [ [-L] [-Xdate] [-Xda] | [-Xall] ] [-E] [-H] [-O] -Xgroup Group-Name
```

To be executed on the Management Server: group specification

```
/opt/FJSVswsrp/bin/swsrpstat -h Server-Name [ [-L] [-Xdate] [-Xda] | [-Xall] ] [-E] [-H] [-O] -
Xgroup Group-Name
```

DESCRIPTION

This command displays the operation status of the specified volume. If a volume is not specified, the operation status of each replication volume is displayed.
The following information will be displayed:

Title	Description
Server	Displays a Managed Server name.
Original-Volume	Displays a replication source volume name.
Replica-Volume	Displays a replication destination volume name.

Title	Description
Direction	Displays the copy direction if a copy session exists. regular : Copying from the replication source to the replication destination is being performed. reverse : Copying from the replication destination to the replication source is being performed. ---- : Copying is not being performed.

Status	Displays the execution status.
	
	---- : Copying is not being performed.
	readying : The concurrent OPC has been defined.
	sync : Total copy or incremental copy is being performed.
	equivalent : Equivalency maintained status.
	replicated : Replication established status.
	copy-on-write : SnapOPC processing is being performed.
	copy-on-write(active) : SnapOPC+ processing is being performed (most recent snap generation).
	copy-on-write(inactive) : SnapOPC+ processing is being performed (past snap generation).
	failed : Error suspend status.
	failed(badsector) : There is a bad sector in the copy source.
	failed(overflow) : There is insufficient Snap Data Volume or Snap Data Pool capacity.
	failed(other) : The error suspend status is neither failed(badsector) or failed(overflow).
	halt : Hard suspend status.
	halt(sync) : Hard suspend status (hard suspend status has been entered during a REC total copy or differential copy).
	halt(equivalent) : Hard suspend status (hard suspend status has been entered during an REC equivalency maintenance state).
	halt(use-buffer) : A line fault occurred and the REC buffer is in use. (**NOTE**)
	halt(use-disk-buffer) : A line fault occurred and the REC Disk buffer is in use. (**NOTE**)
	snap : OPC or QuickOPC processing is being performed.
	????? : Failed to get copy status.
	gds-error : Copy using GDS terminated abnormally.
	
	NOTE: "halt(use-buffer)" and "halt(use-disk-buffer)" are output when the following two conditions are fulfilled:
	
	  - The -H option was specified, and
	
	  - In inter-server replication, this command was executed on the replication source server of REC.

Execute	Displays the copy progress rate as a percentage. "----" is displayed during SnapOPC or when no copying is being performed.
Trk	Displays whether tracking processing is in progress. on : Tracking processing is in progress. off : Tracking processing is not in progress. ---- : Snapshot processing is not being performed. Always displays "----" when the GDS Snapshot linkage function is enabled.

Update	Displays the percentage of data that has been updated since the last copy was created, in the cases shown below.
	For SnapOPC+, this displays the percentage of data that has been updated since the previous snap generation.
	
	  - When the Status column is "----" and the Trk column is "on" (that is, tracking is in progress and the physical copy is not being executed).
	
	  - When the Status column is "copy-on-write" (that is, when SnapOPC processing is in progress)
	
	  - When the Status column is "copy-on-write(active)" (that is, when SnapOPC+ (most recent snap generation) is being performed)
	
	  - When the Status column is "copy-on-write(inactive)" (that is, when SnapOPC+ (past snap generation) is being performed)

Title	Description
	Displays "----" during physical copying or when tracking is not in progress. Always displays "----" when the GDS Snapshot linkage function is enabled.
Rcv	Indicates the Recovery mode for Inter-ETERNUS synchronization. auto : Automatic Recovery mode manual : Manual Recovery mode ---- : Inter-ETERNUS synchronous replication (REC) is not performed.
Split	Indicates the Split mode for Inter-ETERNUS synchronization. auto : Automatic Split mode manual : Manual Split mode ---- : Inter-ETERNUS synchronous replication (REC) is not performed using the Synchronous mode.
Xfer	Indicates the Transfer for Inter-ETERNUS synchronization. sync : Synchronous mode async : Through mode stack : Stack mode consist : Consistency mode ---- : Inter-ETERNUS synchronous replication (REC) is not performed.
Snap-Gen	When SnapOPC+ is performed, displays the snap generation number. When SnapOPC+ is not performed, "----" is displayed. 1 is the oldest snap generation, and the snap generation number increases until the most recent generation is reached.
Copy-DA (Copy Destination Access)	Displays the access permission for the replication destination volume. on : Access is not being permitted. off : Access is being permitted. When an ETERNUS Disk storage system does not support the Destination Access Permission function or the SDX objects is specified with the units of logical volumes, 'off' is displayed.

Date	If a copy session exists, a timestamp in the [yyyy/mm/dd hh:mm:ss] format is displayed for the last status transition of the copy session as described below.
	The timestamp is relative to the time zone settings of the specified storage system.
	
	[For the OPC and QuickOPC]
	
	  - Starting the logical copy (Execution the "swsrpmake" or "swsrpmake -T")
	
	  - Changing to error suspend status (The Status is "failed")
	
	[For the SnapOPC and SnapOPC+]
	
	  - Starting the logical copy (Execution the "swsrpmake -C" or "swsrpmake -P")
	
	  - Changing to error suspend status (The Status is "failed")
	
	[For the EC]
	
	  - Suspend (Execution the "swsrpmake")
	
	  - Concurrent suspend (Execution the "swsrpmake -Xconcur")
	
	  - Changing to error suspend status (The Status is "failed")
	
	[For the REC]
	
	  - Starting the REC by initial copy skip function (Execution the "swsrpstartsync -Xskip")
	
	  - Suspend (Execution the "swsrpmake") (**Note**)
	
	  - Concurrent suspend (Execution the "swsrpmake -Xconcur")
	
	  - Changing to error suspend status (The Status is "failed")
	
	  - Changing to hard suspend status (The Status is "halt")

Title	Description

	Note: When a forced suspend is performed on a REC session in hardware suspend condition, the timestamp of the last execution is initialized as "----/--/-- --:--:--". When a forced suspend is performed while copy is executing, its timestamp is not updated.
	
	"----/--/-- --:--:--" is displayed in the following situations:
	
	  - There is no copy session (Status field is "----")
	
	  - Copy status could not be obtained from the storage system (Status field is "?????")
	
	  - Transition between no copy session and start of a synchronous EC/REC initial copy (swsrpstartsync) is started under the condition of no copy session.
	
	  - The REC session is recovered automatically from the hardware suspend status (The Status is "halt")
	
	  - The storage system where last execution timestamp function is not supported
	
	  - The SDX objects being set per logical volume

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on a Management Server.

-L

Specifies that display will be in extension format.
For the items displayed by this option, refer to the table in the description above.

-E

An option which when set will display the operational status in the opposite direction only (from the copy destination to the copy source).
If this option is not set, the operation status will always be displayed in the normal direction (from the copy source to the copy destination) when a bi-directional OPC physical copy is taking place.
In case this option is set, the operational status in the normal direction will not be displayed. Only the operational status in the opposite direction will be displayed.
This option can be used to check the operational status of an OPC session performed in the opposite direction if a QuickOPC/SnapOPC session in the normal direction exists.

-Xgroup *Group-Name*

Specifies that the operation statuses are displayed by the group.
If this option is set, all the operational statuses relating to the specified group will be displayed.

-H

Changes the Status field display method used when a REC enters a hard-suspend (halt) status.
If this option is specified, then "halt (sync)" will be displayed when hard-suspend status occurs during a total copy or a differential copy. "halt (equivalent)" will be displayed when hard-suspend status occurs during an equivalency maintenance state.
If the transfer mode is Consistency mode and a REC Disk buffer is set, "halt(use-disk-buffer)" will be displayed when issue of a hard suspend causes transfer data to be saved to the REC Disk buffer.
If the transfer mode is Consistency mode and a REC buffer is set, "halt(use-buffer)" will be displayed when issue of a hard suspend causes transfer data to be saved to the REC buffer.

Note that hard-suspend status will still be displayed as "halt" even if this option is specified for GDS snapshot linkage.

-O

Changes the error suspend(failed) status Status field display method.
If this option is specified, the display will be as follows:
There is a bad sector: "failed(badsector)"
When there is insufficient Snap Data Volume or Snap Data Pool capacity: "failed(overflow)"
All other cases: "failed(other)"

However, when linked with GDS Snapshot, only "failed" is displayed when an error suspend status occurs even if -O option is specified.

-Xda

Displays the access permission settings for copy destination volume.

-Xdate

Displays the timestamp of the last copy session status transition.

-Xall

Displays the all information.
For the items displayed by this option, refer to the example following.
Note that the items displayed by this option may be subject to change in a future release.

OPERANDS

Original-Volume-Name

Specifies a replication source volume name.
For the volume names of other servers, use the format "*Volume-Name@Managed-Server-Name*".
Specifies the AdvancedCopy Manager device name for the *Volume-Name*.
Displays all operation statuses for the specified replication source volume.
If this operand is omitted, all the operation statuses of the Managed Server on which this command is executed are displayed.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Direct that all the replication operation statuses for Managed Server (SRC-SV) be displayed on the Management Server:

```
# /opt/FJSVswsrp/bin/swsrpstat -h SRC-SV
Server Original-Volume  Replica-Volume    Direction Status               Execute
SRC-SV /dev/sdb3@SRC-SV /dev/sdb3@TARG-SV regular   replicated           ----
SRC-SV /dev/sdc3@SRC-SV /dev/sdc3@TARG-SV regular   ----                 ----
SRC-SV /dev/sdd3@SRC-SV /dev/sdd3@TARG-SV reverse   snap                 45%
SRC-SV /dev/sdb3@SRC-SV /dev/sdb3@TARG-SV regular   copy-on-write(inactive) ----
SRC-SV /dev/sdb3@SRC-SV /dev/sdb3@TARG-SV regular   copy-on-write(inactive) ----
SRC-SV /dev/sdb3@SRC-SV /dev/sdb3@TARG-SV regular   copy-on-write(active)  ----
#
```

- Direct that all the replication operation statuses for Managed Server (SRC-SV) be displayed on the Management Server (with the -L option specified):

```
# /opt/FJSVswsrp/bin/swsrpstat -h SRC-SV -L
Server Original-Volume  Replica-Volume    Direction Status               Execute Trk  Update
Rcv  Split Xfer  Snap-Gen
SRC-SV /dev/sdb3@SRC-SV /dev/sdb3@TARG-SV regular   replicated           ----    ---- ----
auto ----   async ----
SRC-SV /dev/sdc3@SRC-SV /dev/sdb3@TARG-SV regular   ----                 ----    on   6%
---- ----  ----  ----
SRC-SV /dev/sdd3@SRC-SV /dev/sdb3@TARG-SV reverse   snap                 45%     off  ----
---- ----  ----  ----
SRC-SV /dev/sdb3@SRC-SV /dev/sdb3@TARG-SV regular   copy-on-write(inactive) ----    off  0%
---- ----  ----  1
SRC-SV /dev/sdb3@SRC-SV /dev/sdb3@TARG-SV regular   copy-on-write(inactive) ----    off  5%
---- ----  ----  2
SRC-SV /dev/sdb3@SRC-SV /dev/sdb3@TARG-SV regular   copy-on-write(active)  ----    off  6%
```

```
   ---- ----  ----  3
   #
```

- Direct that all the replication operation statuses for Managed Server (SRC-SV) be displayed on the Management Server (with the -Xall option specified):

```
# /opt/FJSVswsrp/bin/swsrpstat -h SRC-SV -Xall
Server Original-Volume Replica-Volume    Direction Status     Execute Trk  Update Rcv  Split Xfer
Snap-Gen Copy-DA Date
SRC-SV /dev/sdb3@SRC-SV /dev/sdb3@TARG-SV regular   replicated ----    ---- ----   auto ---- async
----     on      2010/10/20 22:00:45
SRC-SV /dev/sdc3@SRC-SV /dev/sdc3@TARG-SV ----      ----       ----    on   6%      ---- ---- ----
----     off     2010/10/20 22:10:39
SRC-SV /dev/sdd3@SRC-SV /dev/sdd3@TARG-SV reverse   snap       45%     off  ----    ---- ---- ----
----     on      2010/10/21 22:00:23
#
```

- Direct that all the replication operation statuses for Managed Server (SRC-SV) be displayed on the Management Server (with the -Xdate option specified):

```
# /opt/FJSVswsrp/bin/swsrpstat -h SRC-SV -Xdate
Server Original-Volume Replica-Volume    Direction Status     Execute Date
SRC-SV /dev/sdb3@SRC-SV /dev/sdb3@TARG-SV regular   replicated ----    2010/10/20 22:00:45
SRC-SV /dev/sdc3@SRC-SV /dev/sdc3@TARG-SV ----      ----       ----    2010/10/20 22:10:39
SRC-SV /dev/sdd3@SRC-SV /dev/sdd3@TARG-SV reverse   snap       45%     2010/10/21 22:00:23
#
```

NOTES

- When the execution status (the status displayed in the "Status" column) is "failed", "halt", or "?????", refer to "8.4 Hardware or Other Error during Replication" and take action.

- If the execution status (displayed in the Status field) is "gds-error":
  A GDS error probably occurred. Remove the cause of the GDS error, then use swsrpcancel (Replication cancellation command) to stop the replication operation.

- The operation status cannot be displayed if:

  - The volumes have not been set as replication volumes.

- This command cannot be executed while the following command is running:

  - swsrprecoverres (Resource adjustment command)

9.4.2.4 swsrpcancel (Replication cancellation command)

NAME

swsrpcancel - stops snapshot processing and synchronization processing

SYNOPSIS

To be executed on the Managed Server

```
/opt/FJSVswsrp/bin/swsrpcancel [-c] {([-f] [-t] [-m] [-b])|[-T]|[-Xforce]} [-Xgds-slice-skipchk]
From-Volume-Name To-Volume-Name
```

To be executed on the Management Server

```
/opt/FJSVswsrp/bin/swsrpcancel -h Server-Name [-c] {([-f] [-t] [-m] [-b])|[-T]|[-Xforce]} [-Xgds-
slice-skipchk] From-Volume-Name To-Volume-Name
```

To be executed on the Managed Server: group specification

```
/opt/FJSVswsrp/bin/swsrpcancel [-c] {([-f] [-t] [-m] [-b])|[-T]|[-Xforce]} [-Xreverse] -Xgroup
Group-Name
```

To be executed on the Management Server: group specification

```
/opt/FJSVswsrp/bin/swsrpcancel -h Server-Name [-c] {([-f] [-t] [-m] [-b])|[-T]|[-Xforce]} [-
Xreverse] -Xgroup Group-Name
```

DESCRIPTION

This command stops snapshot processing (OPC/QuickOPC/SnapOPC/SnapOPC+) and synchronization processing (EC) of the specified replication volume. Use this command when an error is detected during the replication operation and when you want to stop replication processing.

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on a Management Server.

-c

Specifies that the copy processing is forcibly stopped from a non-operation server when the operation server cannot be used because of a system failure. If this option is used, the copy processing is stopped without implementing pre-processing and post-processing.

This option cannot be specified on the operation server.

-m

For server-to-server replication, this option specifies that communication processing is not performed with non-operation servers. When this option is specified, volume status check processing and pre-processing and post-processing are not performed for the replication destination volume of a non-operation server.
This option is valid only during server-to-server replication.

-f

When stopping synchronous processing under the equivalency maintained status, specify that pre-processing and post-processing for the copy source volume is not being executed.
Use this option only when AdvancedCopy Manager does not execute pre-processing and post-processing of the copy source volume and unique pre-processing and post-processing scripts are inserted before and after (respectively) the AdvancedCopy Manager command, or when you judge that copy pre-processing and post-processing are unnecessary (e.g., for a database constructed on a raw device).

-t

When stopping the synchronous processing under the equivalency maintained status, this specifies that post-processing for the copy destination volume will not be executed.
Use this option only when AdvancedCopy Manager does not execute pre-processing and post-processing of the copy destination volume and unique pre-processing and post-processing scripts are inserted before and after (respectively) the AdvancedCopy Manager command, or when you judge that copy pre-processing and post-processing are unnecessary (e.g., for a database constructed on a raw device).

-b

Specifies that the inter-box copy operation in the halt status be stopped.
For replication on a single server, the copying is cancelled for both boxes.
For replication between servers, the copying is cancelled only for the box connected to the server on which this command was executed.
Use the option only for a copy operation in the halt status and whose recovery is not expected to be done successfully.
The option is valid only for a copy operation in the HALT status.
The option can be used on the server that is not specified as the operation server.

This option cannot be used for the logical volume of an SDX object.

If this option is specified together with the -Xgroup option, copy processing is only stopped for sessions in the group whose status is halt.

-T

Specifies that differential snapshot processing will be terminated.

This option only terminates tracking processing and does not terminate snapshot processing (OPC physical copying) when it is in progress. This option releases the OPC session if snapshot processing (i.e., OPC physical copying) is complete.

This option can only be used while physical copying or tracking is in progress.

This option cannot be used for the logical volume of an SDX object.

If this option is specified together with the -Xgroup option, only tracking processing within the group is stopped.

-Xgroup *Group-Name*

Specifies that copy processing will be stopped for each group.

-Xreverse

Specifies that copy processing that is being executed in the reverse direction (from replica volumes to original volumes) will be stopped.

-Xgds-slice-skipchk

Skips the SDX object status check processing that is usually performed as part of the copy source and copy destination pre-processing. This takes place if SDX objects are operated in slice units and the synchronization processing for equivalency maintenance state is stopped.

If this option is specified for a stopped volume, the -f and -t options must also be specified.

-Xforce

Specification of this option forcibly stops a SnapOPC+ session.

Forcibly stopping SnapOPC+ stops the specified snap generation and all earlier snap generations.

This option can be specified even if there is only one snap generation.

This option is a SnapOPC+ fixed option.

If this option is specified at the same time as the -Xgroup option, the SnapOPC+ sessions in that group are stopped, and all the snap generations and snap generations before it in the group are also stopped.

OPERANDS

From-Volume-Name

Specifies a copy source volume name.

Specifies the replication source volume or the replication destination volume set with swsrpsetvol (Replication volume information setting command).

Specifies the AdvancedCopy Manager device name for the volume name.

To-Volume-Name

Specifies a copy destination volume name.

Specifies the replication source volume or the replication destination volume set with swsrpsetvol (Replication volume information setting command).

Specifies the AdvancedCopy Manager device name for the volume name.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

- On the Management Server, direct that copying from /dev/sdb3 to /dev/sdb3@TARG-SV be stopped for Managed Server SRC-SV:

```
#/opt/FJSVswsrp/bin/swsrpcancel -h SRC-SV /dev/sdb3 /dev/sdb3@TARG-SV
FROM=/dev/sdb3@SRC-SV,TO=/dev/sdb3@TARG-SV swsrpcancel completed
#
```

- Stops copy processing for group "GRP1":

```
# /opt/FJSVswsrp/bin/swsrpcancel -Xgroup GRP1
GROUP=GRP1 swsrpcancel completed
#
```

NOTES

- Following execution with the -b option specified in replication on a server, the command stops normally when cancellation of the copying succeeds in either of the boxes.

- Copy processing cannot be performed if:

    - The specified replication source volume and replication destination volume have not been set as replication volumes.

    - The copy processing is not performed for the specified replication source and replication destination volumes.

    - For server-to-server replication, the Managed Server on which the command is executed is not the operation server of the replication volume.

    - For server-to-server replication, communication with a destination server fails.

    - The backup management function is executing a process for a specified source or destination volume

- If the -Xgroup option is specified, processing will be performed for each active replication volume, but processing will be terminated without cancelling replication in the following case:

    - If no replication volumes have been copied

    - When you execute -b option specifying a nonexistent copy in the "halt" state in the group.

    - When you execute -T option specifying a nonexistent copy in the tracking process of QuickOPC in the group.

    Additionally, if the -Xgroup option is specified, replication volumes are processed in order, but processing will be terminated immediately if an error occurs.

- When stopping SnapOPC+ sessions, the oldest snap generation must be stopped first.
  Specify the -Xforce option to forcibly stop the second and subsequent snap generations. However, note that this option stops the specified snap generation and all earlier snap generations.
  If the -Xforce and -Xgroup options are specified at the same time, the SnapOPC+ sessions in the group are stopped, and all the snap generations and snap generations before it in the group are also stopped.

- The volume information for the specified volume is not checked (check for matching of the volume information in the management list and the actual physical information). The operation is cancelled unconditionally using the volume information in the management list.

- The following processing is performed depending on the operation status of the specified replication source volume and replication destination volume.

Table 9.20 Processing differences in accordance with operating status

Operation status	Processing
Reverse copy direction	Suspends the processing.
Not operating yet	Suspends the processing.
Total or incremental copy is performed	Suspends synchronization processing. The replication destination volume cannot be used.

Operation status	Processing
Equivalency maintain status	Suspends synchronization processing. The replication destination volume can be used as a copy
Replication established status	Suspends synchronization processing. The replication destination volume can be used as a copy
Snapshot processing is being performed	Stops the snapshot processing. The replication destination volume cannot be used.

- Pre-processing and post-processing for a replication source and replication destination volumes are performed only when the status of synchronization processing is equivalency maintenance status.

- Synchronization processing cannot be stopped in equivalency maintenance status in the following situations:

  - When a replication source volume is an SDX object (SynfinityDISK) and the status of the SDX object is not one of the following:

    - The mirror volume is ACTIVE or STOP.

    - The mirror partition status is ACTIVE or TEMP.
      When the mirror partition is TEMP, a reproduction former volume is not preprocessed.

    - The SDX disk status is ENABLE.

- When the -T option is not specified
  command processing varies depending on the operation status of the source and destination volumes.

Table 9.21 Processing when the -T option is not specified

Operation status (Display by swsrpstat -L)	Processing
Reverse copy	Processing is interrupted.
Not in operation (Status is "----" and Trk also is "----".)	Processing is interrupted.
Total or difference copying is in progress. (Status is "sync".)	Synchronous processing is terminated. The destination volume cannot be used.
Equivalency maintained status (Status is "equivalent".)	Synchronous processing is terminated. The destination volume can be used as a replica.
Replica created status (Status is "replicated".)	Synchronous processing is terminated. The destination volume can be used as a replica.
Snapshot processing is in progress. (Status is "snap")	Synchronous processing is terminated. The destination volume cannot be used. Tracking processing, if in progress, is also terminated.
Snapshot processing complete status and tracking processing in progress. (Status is "----" and Trk also is "on".)	Tracking processing is terminated. The destination volume can be used as a replica.

- When the -T option is specified
  command processing varies depending on the operation status of the source and destination volumes.

Table 9.22 Processing when the -T option is specified

Operation status (Display by swsrpstat -L)	Processing
Reverse copy	Processing is interrupted.
Not in operation (Status is "----" and Trk also is "----".)	Processing is interrupted.
Total or difference copying is in progress. (Status is "sync".)	Processing is interrupted.

Operation status (Display by swsrpstat -L)	Processing
Equivalency maintained status (Status is "equivalent".)	Processing is interrupted.
Replica created status (Status is "replicated".)	Processing is interrupted.
Snapshot processing is in progress and tracking processing is not in progress. (Status is "snap" and Trk is "off")	Processing is interrupted.
Snapshot processing is in progress and tracking processing is progress. (Status is "snap" and Trk is "on".)	Only tracking processing is terminated. Snapshot processing (ie, physical copying) is continued. The destination volume can be used as a replica.
Snapshot processing complete status and tracking processing in progress. (Status is "----" and Trk also is "on".)	Tracking processing is terminated. The destination volume can be used as a replica.

- When GDS SDX objects are used: Refer to "Configuration and conditions of unsupported SDX objects" and "10.1.5.2 Notes on SDX object operations (replication management)".

- This command cannot be executed while any of the following commands are running:

    - swsrpsetvol (Replication volume information setting command)

    - swsrpdelvol (Replication volume information deletion command)

    - swsrprecoverres (Resource adjustment command)

9.4.2.5 swsrpchsync (Synchronization mode change command)

NAME

swsrpchsync - changes the operation mode

SYNOPSIS

To be executed on the Managed Server

```
/opt/FJSVswsrp/bin/swsrpchsync [-F {Sync|Async|Stack|Consist}] [-R {Auto|Manual}] [-S {Auto|
Manual}] From-Volume-Name To-Volume-Name
```

To be executed on the Management Server

```
/opt/FJSVswsrp/bin/swsrpchsync -h Server-Name [-F {Sync|Async|Stack|Consist}] [-R{Auto|Manual}] [-
S {Auto|Manual}] From-Volume-Name To-Volume-Name
```

To be executed on the Managed Server: group specification

```
/opt/FJSVswsrp/bin/swsrpchsync [-F {Sync|Async|Stack|Consist}] [-R {Auto|Manual}] [-S {Auto|
Manual}] [-Xreverse] -Xgroup Group-Name
```

To be executed on the Management Server: group specification

```
/opt/FJSVswsrp/bin/swsrpchsync -h Server-Name [-F {Sync|Async|Stack|Consist}] [-R {Auto|Manual}] [-
S {Auto|Manual}] [-Xreverse] -Xgroup Group-Name
```

DESCRIPTION

This command changes the operation mode of inter-box synchronization (transfer mode, recovery mode, or split mode).

OPTIONS

-h *Server-Name*

> Specifies the name of a Managed Server.
> This option can be specified only on a Management Server.

-F {Sync|Async|Stack|Consist}

> Changes the transfer mode.
> You can specify one of the following parameters in the operand:
>
> Sync : Changes to Synchronous mode.
> Async : Changes to Through mode.
> Stack : Changes to Stack mode.
> Consist : Changes to Consistency mode.
>
> Changing the transfer mode requires this option.
> You can only specify the -S option together with the -F Sync option.
> Specifying the -F Sync option to change the transfer mode from a non-synchronous mode to the synchronous mode without specifying the -S option sets the automatic split mode.

-R {Auto|Manual}

> Changes the recovery mode.
> You can specify either of the following parameters in the operand:
>
> Auto : Changes to Automatic Recovery mode.
> Manual : Changes to Manual Recovery mode.
>
> Changing the recovery mode requires this option.

-S {Auto|Manual}

> Changes the split mode.
> You can specify either of the following parameters in the operand:
>
> Auto : Changes to Automatic Split mode.
> Manual : Changes to Manual Split mode.
>
> Changing the split mode requires this option.
> To specify this option, you must also specify -F Sync, or the transfer mode of the current synchronization must be synchronous mode.

-Xgroup *Group-Name*

> Changes the operation mode of inter-box synchronous processing for each group (transfer mode, recovery mode, or split mode).

-Xreverse

> Specifies that the operation mode of synchronous processing that is being executed in reverse direction (from replica volumes to original volumes) will be changed.

OPERANDS

From-Volume-Name

> Specifies a copy source volume name.
> Specifies the replication source volume or the replication destination volume set with swsrpsetvol (Replication volume information setting command).
> Specifies the AdvancedCopy Manager device name for the volume name.

To-Volume-Name

> Specifies a copy destination volume name.
> Specifies the replication source volume or the replication destination volume set with swsrpsetvol (Replication volume information setting command).
> Specifies the AdvancedCopy Manager device name for the volume name.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- From Management Server to Managed Server: Instruct SRC-SV to change the transfer mode for inter-box synchronization between /dev/sdb3@SRC-SV and /dev/sdb3@TARG-SV from Consistency mode to Through mode:

```
#/opt/FJSVswsrp/bin/swsrpstat -h SRC-SV -L /dev/sdb3@SRC-SV
Server Original-Volume  Replica-Volume    Direction Status Execute Trk  Rcv  Split Xfer Snap-Gen
SRC-SV /dev/sdb3@SRC-SV /dev/sdb3@TARG-SV regular   sync   99%     ---- auto ---- consist ----
:
#/opt/FJSVswsrp/bin/swsrpchsync -h SRC-SV -F Async /dev/sdb3@SRC-SV /dev/sdb3@TARG-SV
FROM=/dev/sdb3@SRC-SV, TO=/dev/sdb3@TARG-SV swsrpchsync completed
:
#/opt/FJSVswsrp/bin/swsrpstat -h SRC-SV -L /dev/sdb3@SRC-SV
Server Original-Volume  Replica-Volume    Direction Status    Execute Trk  Rcv  Split Xfer Snap-Gen
SRC-SV /dev/sdb3@SRC-SV /dev/sdb3@TARG-SV regular   equivalent 100%    ---- auto ---- async ----
#
```

- Specify that the transfer mode of inter-box synchronous processing for group "GRP1" be changed from Consistency mode to Through mode:

```
# /opt/FJSVswsrp/bin/swsrpchsync -F Async -Xgroup GRP1
GROUP=GRP1 swsrpchsync completed
#
```

NOTES

- In any of the following cases, you cannot change the operation mode of inter-box synchronization:

    - The specified original/replica volume has not been configured as a replica volume.

    - None of the -F, -R and -S options are specified.

    - In inter-server replication, the Managed Server on which you execute the command is not the operation server for the replica volume.

    - Synchronization from the specified source volume to the destination volume is not secured.

    - Synchronization from the specified source volume to the destination volume is in error-suspended status (with "failed" indicated as the Status by swsrpstat (Operation status display command) or hard-suspended status (with "halt" indicated as the Status by swsrpstat (Operation status display command).

    - The ETERNUS Disk storage system in which the specified original/replica volume is placed do not support the operation mode change function.

    - The specified volume is being used with GDS Snapshot.

- If the -Xgroup option is specified, the synchronous processing for all replication volumes being processed is changed to the same operation mode. However, the operation mode cannot be changed if there is at least one replication volume that meets the conditions above.
  Additionally, if the -Xgroup option is specified, replication volumes are processed in order, but processing will be terminated immediately if an error occurs.

- If the REC operating mode is changed while in equivalency maintenance status, after the operating mode is changed, users must wait until the equivalency maintenance status is attained.

- This command cannot be executed while any of the following commands are running:

    - swsrpsetvol (Replication volume information setting command)

9.4.2.6 swsrprevsync (Synchronization reverse command)

NAME

swsrprevsync - reverses the copying direction of synchronization

SYNOPSIS

To be executed on the Managed Server

```
/opt/FJSVswsrp/bin/swsrprevsync From-Volume-Name To-Volume-Name
```

To be executed on the Management Server

```
/opt/FJSVswsrp/bin/swsrprevsync -h Server-Name From-Volume-Name To-Volume-Name
```

To be executed on the Managed Server: group specification

```
/opt/FJSVswsrp/bin/swsrprevsync [-Xreverse] -Xgroup Group-Name
```

To be executed on the Management Server: group specification

```
/opt/FJSVswsrp/bin/swsrprevsync -h Server-Name [-Xreverse] -Xgroup Group-Name
```

DESCRIPTION

This command reverses the copying direction of synchronization when in suspended status.

Note that this command cannot be used to handle a logical volume of an SDX object.

OPTIONS

-h *Server-Name*

Specifies the name of a Managed Server.
This option can be specified only on a Management Server.

-Xgroup *Group-Name*

Changes the direction of synchronous processing for each group.

-Xreverse

Specifies that the direction of synchronous processing that is being executed in reverse direction (i.e., from replica volumes to original volumes)will be inverted.

OPERANDS

From-Volume-Name

Specifies a copy source volume name. (as the destination in the reverse direction)
Specifies the replication source volume or the replication destination volume set with swsrpsetvol (Replication volume information setting command).
Specifies the AdvancedCopy Manager device name for the volume name.

To-Volume-Name

Specifies a copy destination volume name. (as the source in the reverse direction)
Specifies the replication source volume or the replication destination volume set with swsrpsetvol (Replication volume information

setting command).

Specifies the AdvancedCopy Manager device name for the volume name.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- From Management Server to Managed Server: Instruct SRC-SV to reverse the copying direction of synchronization between /dev/sdb3@SRC-SV and /dev/sdb3@TARG-SV:

```
#/opt/FJSVswsrp/bin/swsrpstat -h SRC-SV /dev/sdb3@SRC-SV
Server Original-Volume  Replica-Volume   Direction  Status     Execute
SRC-SV /dev/sdb3@SRC-SV /dev/sdb3@TARG-SV reverse    replicated ----
#/opt/FJSVswsrp/bin/swsrprevsync /dev/sdb3@TARG-SV /dev/sdb3@SRC-SV
FROM=/dev/sdb3@SRC-SV,TO=/dev/sdb3@TARG-SV swsrprevsync completed
#/opt/FJSVswsrp/bin/swsrpstat /dev/sdb3@SRC-SV
Server Original-Volume  Replica-Volume   Direction  Status     Execute
SRC-SV /dev/sdb3@SRC-SV /dev/sdb3@TARG-SV regular    replicated ----
#
```

- Change the direction of synchronous processing for group "GRP1":

```
# /opt/FJSVswsrp/bin/swsrprevsync -Xgroup GRP1
GROUP=GRP1 swsrprevsync completed
#
```

NOTES

- In any of the following cases, you cannot reverse the direction of synchronization:

  - The specified original/replica volume has not been configured as a replica volume.

  - Unidirectional copying has been specified in the copying attribute for the replica volume.

  - In inter-server replication, the Managed Server on which you execute the command is not the operation server for the replica volume.

  - Replication has not been established in synchronization from the specified source volume to the destination volume.

  - The ETERNUS Disk storage system in which the specified original/replica volume is placed do not support the reverse function.

  - The specified volume is being used with GDS Snapshot.

- If the -Xgroup option is specified, the direction of synchronous processing for all replication volumes being processed is changed. However, the direction of synchronous processing cannot be changed if there is at least one replication volume that meets the conditions above.
  Additionally, if the -Xgroup option is specified, replication volumes are processed in order, but processing will be terminated immediately if an error occurs.

- If the -Xgroup option is specified, processing is not performed for those sessions in the group where the direction has already been changed.

- This command cannot be executed while any of the following commands are running:

  - swsrpsetvol (Replication volume information setting command)

  - swsrpdelvol (Replication volume information deletion command)

  - swsrprecoverres (Resource adjustment command)

9.4.3 Maintenance commands

This section describes the maintenance commands for replication management.

9.4.3.1 swsrprecoverres (Resource adjustment command)

NAME

swsrprecoverres - executes recovery to make replication management information

SYNOPSIS

To be executed on the Managed Server

```
/opt/FJSVswsrp/bin/swsrprecoverres [-r] [-p]
```

To be executed on the Management Server

```
/opt/FJSVswsrp/bin/swsrprecoverres -h Server-Name [-r] [-p]
```

DESCRIPTION

This command executes recovery to make replication management information consistent if a system failure or other problem has made the replication management information inconsistent.

OPTIONS

-h *Server-Name*

Specifies a Managed Server name.
This option can be specified only on a Management Server.

-r

Specify this option when an incorrect status exists in the replication management files. E.g., when swsrpsetvol (Replication volume information setting command) or swsrpdelvol (Replication volume information deletion command) that update the replication management files terminate suddenly because of a forced stop, a system failure or the replication management files have been deleted by mistake.
When this option is specified, the following processing is performed to recreate the replication management files.
Replication management information in the Management Server repository is downloaded to the Managed Server, and the replication volume information is restored.
The ETERNUS Disk storage system is queried, and the execution information for the synchronization processing is restored.

-p

Specifies to not perform any post-processing that has not already been performed. This option can be used when swsrpmake (Replication creation command) processing terminates suddenly because of a forced stop or system failure.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Perform resource adjustment:

```
#/opt/FJSVswsrp/bin/swsrprecoverres
swsrprecoverres completed
#
```

NOTES

- The resource adjustment processing cannot be performed if:

    - Communication with a Management Server fails and the -r option is specified.

    - The -r option is not specified when this command is executed after processing by swsrpsetvol (Replication volume information setting command) and swsrpdelvol (Replication volume information deletion command) has been forcibly terminated.

- The consistency of resource information is restored only for the Managed Server on which this command is executed. When a resource inconsistency occurs during processing related to replication between servers, this command must be executed both on the Managed Server on which the processing was performed and on the Managed Server with which communication is being made.

- If this command is executed for the copy source volume and the copy destination volume in the status where swsrpmake (Replication creation command) is interrupted without completing its process and the remaining processing must be completed, the replication post-processing is executed.

- This command cannot be executed while any of the following commands are running:

    - swsrpsetvol (Replication volume information setting command)

    - swsrpvolinfo (Replication volume information display command)

    - swsrpdelvol (Replication volume information deletion command)

    - swsrpstartsync (Synchronous processing start command)

    - swsrpmake (Replication creation command)

    - swsrpstat (Operation status display command)

    - swsrpcancel (Replication cancellation command)

    - swsrpchsync (Synchronization mode change command)

    - swsrprevsync (Synchronization reverse command)

    - swsrprecoverres (Resource adjustment command)

9.5 Start/Stop Daemon Commands

This section describes how to start and stop AdvancedCopy Manager daemons.

 Note
. . . .

Notes when using commands:

- Confirm that the environment variable SWSTGNODE is not defined.

- Execute this command using the same environment variable LANG that of the code system used by the system.

- Only a root user can execute this command.

- For cluster operation, refer to "Starting and Stopping Daemons".

. . . .

9.5.1 Starting and stopping the communication daemon

Starting the communication daemon

Enter the following command to stop the communication daemon:

```
# /opt/FJSVswstf/bin/stgfwcom start
```

Stopping the communication daemon

Enter the following command to stop the communication daemon:

```
# /opt/FJSVswstf/bin/stgfwcom stop
```

📘 Point
. . . .

In order to verify if the communication daemon has been launched normally, please check if the message stgxfws0060 has been generated in the syslog.

. . . .

9.5.2 Starting and stopping the RMI daemon

The authentication daemon runs only on the Management Server.

Starting the RMI daemon

Enter the following command to start the RMI daemon:

```
# /opt/FJSVswstf/bin/stgrmicom start
```

Stopping the RMI daemon

Enter the following command to stop the RMI daemon:

```
# /opt/FJSVswstf/bin/stgrmicom stop
```

9.5.3 Starting and stopping the RDB daemon

The authentication daemon runs only on the Management Server.

Starting the RDB daemon

Enter the following command to start the RDB daemon:

```
# /opt/FJSVswstf/bin/stgdbcom start
```

Stopping the RDB daemon

Enter the following command to stop the RDB daemon:

```
# /opt/FJSVswstf/bin/stgdbcom stop
```

9.5.4 Starting and stopping AdvancedCopy Manager daemons

The daemons listed below are started and stopped together:

Table 9.23 Daemons started and stopped as a package

Server type	Daemon name
Managed Server	Communication daemon

Server type	Daemon name
Management Server	Communication daemon RMI daemon RDB daemon

Starting a AdvancedCopy Manager daemon

Enter the following command to start a AdvancedCopy Manager daemon:

```
# /opt/swstorage/bin/startacm
```

Stopping a AdvancedCopy Manager daemon

Enter the following command to stop a AdvancedCopy Manager daemon:

```
# /opt/swstorage/bin/stopacm
```

9.6 SnapOPC/SnapOPC+ Support Commands

This section explains the commands can be used to execute SnapOPC/SnapOPC+.

Note

Operating environment for command execution

The table below lists the requirements for the operating environment in which to execute the respective commands.

Table 9.24 Operating environment when executing SnapOPC/SnapOPC+ support commands

Command	Operating environment of AdvancedCopy Manager
swstestupdate	-
swstsdv	-

Note

In cluster operation

In cluster operation, execute a command by performing the following steps:

1. If necessary, specify the logical node name to the transaction in the environment variable SWSTGNODE. Specify the environment variable as follows:

```
# SWSTGNODE=logical-node-name
# export SWSTGNODE
```

2. Execute the command.

The table below lists the commands available in cluster operation, whether environment variables must be configured, and the nodes on which commands are executable.

Table 9.25 Commands when cluster operation is used

Command	Environment variable	Command execution node	Remarks
swstestupdate	Not required	Transaction operation node	-
swstsdv	Not required	Transaction operating node	-

9.6.1 swstestupdate (Update volume measurement command)

NAME

swstestupdate - measures the physical capacity of a Snap Data Volume

SYNOPSIS

When starting update volume measurements (setting up a pseudo session)

```
/opt/FJSVswsts/bin/swstestupdate start Volume-Name
```

When looking up the update volume (displaying the status of the pseudo session)

```
/opt/FJSVswsts/bin/swstestupdate status Volume-Name
```

When finishing update volume measurements (cancelling the pseudo session)

```
/opt/FJSVswsts/bin/swstestupdate stop Volume-Name
```

DESCRIPTION

This command measures the physical capacity of a Snap Data Volume using the pseudo session setting function.

When starting update volume measurements (setting up a pseudo session)

Specify "start" as the subcommand to start measure of the physical capacity of a Snap Data Volume.

When looking up the update volume (displaying the status of the pseudo session)

Specify "status" as the subcommand to display the physical capacity of a Snap Data Volume.
The following information will be displayed:

Title	Description
Volume-Name	The volume name.
Update	The amount of data that has been updated since measurements started in the number of blocks (1 block = 512 bytes). Displays "----" if there is no pseudo session. Displays "failed" if the pseudo session is error-suspended.

When finishing update volume measurements (cancelling the pseudo session)

Specify "stop" as the subcommand to stop measure of the physical capacity of a Snap Data Volume

OPERANDS

Volume-Name

- Specifies the volume name.
  Specify a copy source volume that is used with SnapOPC/SnapOPC+ operations

  GDS logical volumes cannot be specified.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Measure the update volume:

```
(Start taking measurements by starting a pseudo session)
#/opt/FJSVswsts/bin/swstestupdate start /dev/sda1
/dev/sda1 swstestupdate completed

(First measurement)
#/opt/FJSVswsts/bin/swstestupdate status /dev/sda1
Volume-Name Update
/dev/sda1   644333

(Second measurement)
#/opt/FJSVswsts/bin/swstestupdate status /dev/sda1
Volume-Name Update
/dev/sda1   1045000

(Stop taking measurements by cancelling the pseudo session)
# /opt/FJSVswsts/bin/swstestupdate stop /dev/sda1
/dev/sda1 swstestupdate completed
#
```

NOTES

- The amount of data updated cannot be measured in the following cases:

    - When the specified volume does not exist.

    - When the ETERNUS Disk storage system device does not support the SnapOPC/SnapOPC+ function.

    - When the pseudo session has already been set up (i.e., when the "start" option is specified).

    - When the pseudo session does not exist (i.e., when the "stop" option is specified).

9.6.2 swstsdv (Snap Data Volume operation/reference command)

NAME

swstsdv - Initialize or display the Snap Data Volume

SYNOPSIS

Snap Data Volume initialization

```
/opt/FJSVswsts/bin/swstsdv init Volume-Name
/opt/FJSVswsts/bin/swstsdv init -G Access-Volume-Name LU-Number
```

Snap Data Volume status display

```
/opt/FJSVswsts/bin/swstsdv stat Volume-Name
/opt/FJSVswsts/bin/swstsdv stat -G Access-Volume-Name LU-Number
```

Snap Data Pool status display

```
/opt/FJSVswsts/bin/swstsdv poolstat -G Volume-Name
```

DESCRIPTION

Snap Data Volume initialization/status display and Snap Data Pool status display are all possible. Specify one of the following subcommands according to the contents you want to operate on, and then execute the subcommand.

Snap Data Volume initialization

Specify "init" as the subcommand to initialize the Snap Data Volume.

Initialize the Snap Data Volume when SnapOPC/SnapOPC+ is not running.

All the data in the Snap Data Volume is destroyed in the initialization of the Snap Data Volume. For this reason, the following procedure is recommended.

(Creation of the access volume and collection of the LU-Number must be completed before logical volume manager operations.)

1. If there is necessary data in the Snap Data Volume, back it up in advance using the relevant OS command.

2. Initialize the disk management information.
   Set the target Snap Data Volume (physical disk) outside logical volume manager management when there is a logical volume manager operation.

3. Initialize the Snap Data Volume using this command.
   Initialize the information using the access volume when there is a logical volume manager operation.

Snap Data Volume status display

Specify "stat" as the subcommand to display Snap Data Volume logical capacity, physical capacity, used space capacity, and unused space capacity.

The contents that are displayed are as follows:

Title	Description
BoxID	Displays the Box ID.
LUN	Displays the logical unit number on the ETERNUS Disk storage system. The format used for the display is "\<decimal\>(\<hexadecimal\>)".
Rate	Displays the Snap Data Volume usage as a percentage (%).
Logical	Displays the Snap Data Volume logical capacity. The unit used for display is 'sector'.
Physical	Displays the Snap Data Volume physical capacity. The unit used for display is 'sector'.
Used	Displays the amount of Snap Data Volume used space. The unit used for display is 'sector'. The capacity for both Copy usage (**Note**) and Host usage (**Note**) is displayed.
Copy	For the space used in the Snap Data Volume, Copy usage (**Note**) is displayed. The unit used for display is 'sector'.
Host	For the space used in the Snap Data Volume, Host usage (**Note**) is displayed. The unit used for display is 'sector'.
Pool	Displays the Snap Data Pool usage allocated to the Snap Data Volume. The unit used for display is 'sector'.

Note:
Copy usage: Usage in which the update occurs while SnapOPC/SnapOPC+ is running.
Host usage: Update of Snap Data Volume while SnapOPC/SnapOPC+ is not running.

Snap Data Pool status display

Specify "poolstat" as the subcommand to display the Snap Data Pool total amount, used space capacity, and unused space capacity.

The contents that are displayed are as follows:

Title	Description
BoxID	Displays the Box ID.
Pool-Type	Displays the Snap Data Pool type. Normal Encrypted

Title	Description
Rate	Displays the Snap Data Pool usage rate (%).
Total	Displays the Snap Data Pool total amount. The unit used for display is 'sector'.
Used	Displays the Snap Data Pool used space capacity. The unit used for display is 'sector'. The capacity for both Copy usage (**Note**) and Host usage (**Note**) is displayed.
Copy	For the space used in the Snap Data Pool, Copy usage (**Note**) is displayed. The unit used for display is 'sector'.
Host	For the space used in the Snap Data Pool, Host usage (**Note**) is displayed. The unit used for display is 'sector'.
Free	Displays the Snap Data Pool unused space capacity. The unit used for display is 'sector'.

Note:

Copy usage: Usage in which the update occurs while SnapOPC/SnapOPC+ is running.

Host usage: Update of Snap Data Volume while SnapOPC/SnapOPC+ is not running.

OPTIONS

Specific options of Snap Data Volume initialization(init)

-G *Access-Volume-Name*

Specifies the access volume.
For details on access volume, refer to "Creating an access volume".

Specific options of Snap Data Volume status display(stat)

-G *Access-Volume-Name*

Specifies the access volume.
For details on access volume, refer to "Creating an access volume".

Specific options of Snap Data Pool status display(poolstat)

-G *Volume-Name*

Specifies the ETERNUS Disk storage system volume.
If a Snap Data Volume is specified with insufficient capacity however, an error occurs in the command. In this case, specify an access volume for which there is sufficient capacity.
For details on access volume, refer to "Creating an access volume".

OPERANDS

Specific operands of Snap Data Volume initialization(init)

Volume-Name

Specifies the volume name (AdvancedCopy Manager device name) of the Snap Data Volume to be initialized.
If a partition (slice) has been specified, the Logical Unit (disk) on which the partition (slice) exists is initialized.

LU-Number

Specifies the logical unit number in the ETERNUS Disk storage system.
This value is specified as a decimal.

Specific operands of Snap Data Volume status display(stat)

Volume-Name

Specifies the Snap Data Volume volume name (AdvancedCopy Manager device name).
If a partition (slice) has been specified, the information about the Logical Unit (disk) on which the partition (slice) exists is displayed.

Specifies the logical unit number in the ETERNUS Disk storage system.
This value is specified as a decimal.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Snap Data Volume is initialized:

```
# /opt/FJSVswsts/bin/swstsdv init /dev/sdz
swstsdv init completed
#
```

- The Snap Data Volume is initialized using the access volume following insufficient capacity:

```
# /opt/FJSVswsts/bin/swstsdv init -G /dev/sdt 110
swstsdv init completed
#
```

- The Snap Data Volume status is displayed:

```
# /opt/FJSVswsts/bin/swstsdv stat /dev/sdz
BoxID = 00E4000M3#####E450S20A####KD4030639004##
LUN = 117 (0x75)
Rate Logical(sector) Physical(sector) Used(sector) Copy(sector) Host(sector) Pool(sector)
0%   8388608         1045876          0            0            0            0
#
```

- The LU-Number (110) status is displayed using the access volume:

```
# /opt/FJSVswsts/bin/swstsdv stat -G /dev/sdt 110
BoxID = 00e4000m3#####E450S20A####KD4030639004##
LUN = 110 (0x6E)
Rate Logical(sector) Physical(sector) Used(sector) Copy(sector) Host(sector) Pool(sector)
0%   8388608         1048576          0            0            0            0
#
```

- The Snap Data Pool status is displayed by specifying the access volume:

```
# /opt/FJSVswsts/bin/swstsdv poolstat -G /dev/sdt
BoxID = 00E4000M3#####E450S20A####KD40306390004##
Pool-Type Rate Total(sector) Used(sector) Copy(sector) Host(sector) Free(sector)
Normal    10%  20971520      2097152      0            2097152      18874368
Encrypted 0%   20971520      0            0            0            20971520
#
```

NOTES

- When there is insufficient Snap Data Volume capacity, operations are possible using the access volume and logical unit number (LU-Number) in the ETERNUS Disk storage system.
  When there is insufficient Snap Data Volume capacity, however, and "stat" is specified in the subcommand and an attempt is made to check the LU-Number of the Snap Data Volume with insufficient Snap Data Volume capacity, an access error occurs when the attempt to access the Snap Data Volume with insufficient Snap Data Volume capacity is made. In this case, an error occurs in the

command and it is not possible to check the LU-Number.
For this reason, you must check the LU-Number in advance and record it.

9.7 Troubleshooting data collection command

This section explains the command used to collect troubleshooting data when an AdvancedCopy Manager fault occurs.

9.7.1 esfsnap (AdvancedCopy Manager's manager troubleshooting data collection command)

NAME

esfsnap - Collection of AdvancedCopy Manager's manager troubleshooting data

SYNOPSIS

```
/opt/FJSVesfcm/bin/esfsnap -dir dirname [-all]
```

DESCRIPTION

This command collects AdvancedCopy Manager troubleshooting data. Data is collected for AdvancedCopy Manager's manager installed on the Management Server where this command is executed.

Only the root user can execute this command.

Please send the collected data to a Fujitsu system engineer.

OPERANDS

-dir *dirname*

Specify a directory where the troubleshooting data is stored. It is necessary to specify the full path name for *dirname*.

The collected troubleshooting data is stored in the directory which its name begins with "esfsnap_". Before executing this command, please check that *dirname* directory has the following free space.

Collected information	Required free space
Initial investigation information	more than 40MB
All troubleshooting information	more than "80 + (2 * number of registered device)" MB

-all

Collects all troubleshooting data on servers. Data includes lots of information, which will increase the size of the data. This option is not necessary for initial investigation.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Collect data to the /tmp directory:

  ```
  # /opt/FJSVesfcm/bin/esfsnap -dir /tmp
  ```

NOTES

- Refer to NOTES under "esfsnap (CCM Server troubleshooting data collection command)" in the "ETERNUS SF Operation Guide for Copy Control Module".

9.7.2 acmcapture (AdvancedCopy Manager's agent troubleshooting data collection command)

NAME

acmcapture - Collection of AdvancedCopy Manager's agent troubleshooting data

SYNOPSIS

```
/opt/FJSVswstf/bin/acmcapture dirname
```

DESCRIPTION

This command collects AdvancedCopy Manager's agent troubleshooting data. Data is collected for AdvancedCopy Manager's agent installed on the Managed Server where this command is executed.

Only the root user can execute this command.

Troubleshooting data file name

The collected data is stored in a file with the "acmcapture_HOSTNAME_YYYYMMDD_ttmmss.tar.gz" format under the directory specified for "dirname".

The name of the host where the command is executed is used for HOSTNAME. If characters other than alphanumerics and hyphens are used in the host name, then "NOHOSTNAME" is used as the host name. Values are truncated to 15 characters.

- For YYYYMMDD, YYYY is the year, MM is the month, and DD is the day.

- For ttmmss, tt is the hour, mm is minutes, and ss is seconds.

OPERANDS

dirname

Directory where the troubleshooting data is stored.
Specify a directory with plenty free space (500MB or more) is recommended. Regard 500MB as a rough guide, since the size of the troubleshooting data will vary depending on the system environment and conditions. If there is not enough free space, then this command will terminate abnormally without collecting troubleshooting data.
Specify an existing directory, otherwise this command will terminate abnormally without collecting troubleshooting data.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Collect data to the /tmp directory:

```
# /opt/FJSVswstf/bin/acmcapture /tmp
Executing.......................................
Completed.
OUTPUT=/tmp/acmcapture_srv01_20090323_094614.tar.gz
```

NOTES

- In a clustered system, execute this command at all nodes. To gather information on the shared disk, mount the shared disk on a node and execute this command on that node.

- Execute this command when other AdvancedCopy Manager commands are in the stopped state. Data collection may not be performed correctly if other AdvancedCopy Manager commands are running.

Chapter 10 Notes on Operations

This chapter provides notes on the operations of AdvancedCopy Manager.

10.1 Notes on Backup and Replication Operation

This section provides notes on backup and replication.

10.1.1 General notes

This section provides general notes on backup and replication.

10.1.1.1 Size of a Partition

A transaction volume and a backup volume must have the same size.

However, not all of the size specified by the fdisk command will be usable for the following disks because the OS control space is automatically allocated.

- Primary partition in the beginning of a disk (the leftmost on the fdisk command)

- All the extended logical disks

Figure 10.1 Available Size of a Partition

Create partitions for the transaction and backup volumes in the same size and combination as shown below.

Figure 10.2 Combination of Transaction and Backup Volume

Figure 10.3 Example of a Disk with Partitions of Different Sizes

If you are creating the source and destination volumes in the same size for replication, create partitions in the combination as shown above.

 Point
. . . .

To prevent an operational error, you are recommended to define as extended logical disks all the logical disks on the ETERNUS Disk storage system to be managed by AdvancedCopy Manager.

If you cannot define all the logical disks as extended logical disks, you are recommended to use the same configuration for a partition, including a device defined as a transaction volume and a partition with a device defined as a backup volume.

. . . .

📖 Note

. . . .

To back up a PRIMECLUSTER GDS volume

To backup a GDS volume, specify the sector size to set a backup volume size that matches the GDS volume.
Create the backup volume as follows:

1. Find out the accurate size and number of the GDS volumes to be backed up.
    The accurate size is obtained by multiplying the number of blocks output by the following command by 512 bytes.

    ```
    sdxinfo -o <volume-name>
    ```

    In the following example, 425984 is output as the number of blocks. Therefore, the size is 218103808 bytes (425984 blocks x 512).

    ```
    # sdxinfo -o v1
    OBJ     NAME    CLASS   GROUP   SKIP JRM 1STBLK   LASTBLK  BLOCKS   STATUS
    ------ ------- ------- ------- ---- --- -------- -------- -------- --------
    volume v1      c1      g1      off  on     65536   491519   425984 ACTIVE
    ```

2. Use the parted command to check that the disk format of the backup volume is msdos. If the format is not msdos, execute "/sbin/parted <device name> mklabel msdos" to format as msdos.

3. Find out the number of bytes in each sector of the disk where the backup volume is to be created.

    In the following example, "sectors of 512 bytes" indicates that the size of each sector is 512 bytes.

    ```
    # /sbin/sfdisk -uS -l /dev/sdp
    Disk /dev/sdp: 1017 cylinders, 133 heads, 62 sectors/track
    Units = sectors of 512 bytes, counting from 0

    Device Boot    Start       End   #sectors  Id  System
    /dev/sdp1           0         -          0   0  Empty
    /dev/sdp2           0         -          0   0  Empty
    /dev/sdp3           0         -          0   0  Empty
    /dev/sdp4           0         -          0   0  Empty
    ```

    The following error is returned if the disk does not have an msdos disk label:

    ```
    "sfdisk: ERROR: sector 0 does not have an msdos signature"
    ```

    Use the fdisk command or the parted command in advance to create an msdos disk label on the targeted disk.

    Example: Use the # parted command to create an msdos disk label.

    ```
    # /sbin/parted /dev/sdp mklabel msdos
    ```

4. Calculate the number of sectors in the backup volume being created.

    ```
    218103808 bytes / 512 bytes = 425984 sectors
    ```

    However, the partition descriptor is located in the first sector of a logical volume, so this sector cannot be used. Therefore, add one sector to the number of sectors obtained in this calculation.

    ```
    425984 sectors + 1 = 425985 sectors
    ```

Thus, the number of sectors to create is 425985.

5. Use the sfdisk command to create the files to be input.
   Since files need to be created as logical partitions in extension partitions in a backup volume, files like the following are created:

```
# cat ./sdp-partition.dat
,,E
;
;
;
,128,L
,425985,L
,425985,L
,425985,L
```

The first logical partition (/dev/sdp5) cannot be made the same size as the other logical partitions for reasons related to partition descriptors.
Therefore, the first logical partition is created to be the appropriate size (in this case, 128 sectors) and is not used as part of the backup volume.

6. Use the sfdisk command to create the partitions.

```
#  /sbin/sfdisk -uS --force /dev/sdp < ./sdp-partition.dat
Checking that no-one is using this disk right now ...
OK

Disk /dev/sdp: 1017 cylinders, 133 heads, 62 sectors/track
Old situation:
Units = sectors of 512 bytes, counting from 0

   Device Boot    Start        End   #sectors  Id  System
/dev/sdp1              0         -          0   0  Empty
/dev/sdp2              0         -          0   0  Empty
/dev/sdp3              0         -          0   0  Empty
/dev/sdp4              0         -          0   0  Empty
New situation:
Units = sectors of 512 bytes, counting from 0

   Device Boot    Start        End   #sectors  Id  System
/dev/sdp1              1   8386181    8386181   5  Extended
/dev/sdp2              0         -          0   0  Empty
/dev/sdp3              0         -          0   0  Empty
/dev/sdp4              0         -          0   0  Empty
/dev/sdp5              2       129        128  83  Linux
/dev/sdp6            131    426114     425984  83  Linux
/dev/sdp7         426116    852099     425984  83  Linux
/dev/sdp8         852101   1278084     425984  83  Linux
Warning: partition 5 does not end at a cylinder boundary
Successfully wrote the new partition table

Re-reading the partition table ...

If you created or changed a DOS partition, /dev/foo7, say, then use dd(1)
to zero the first 512 bytes:  dd if=/dev/zero of=/dev/foo7 bs=512 count=1
(See fdisk(8).)
```

Use the created /dev/sdp6, /dev/sdp7, and /dev/sdp8 as the backup volume.

10.1.1.2 Backup and Replication Exclusions

The following devices must not be backed up or replicated:

- The device on which AdvancedCopy Manager has been installed

- The device on which the management list of AdvancedCopy Manager resides

- Logical volume of LVM(Logical Volume Manager) or volume group

10.1.1.3 Consistency of data on volumes

Table 10.1 Consistency of data on volumes

Data in transaction volume	Maintenance of data integrity	Operation
Oracle database space that is currently in use	Linkage with Oracle is required.	Use a sample script designed for linkage with Oracle. For details, refer to "Appendix B Sample Script for Backup Operations".
Symfoware database space that is currently in use	AdvancedCopy Manager commands link with Symfoware to ensure data consistency.	Refer to "Chapter 4 Backup Operations (Symfoware)".
File system	An AdvancedCopy Manager command unmounts the file system to maintain the integrity.	Refer to "Appendix A Pre-processing and Post-processing of Backup and Restoration" and "Appendix C Pre-processing and Post-processing of Replication".
Other than the above	The data integrity must be maintained in operations.	Take appropriate action, such as stopping transactions during execution of backup or replication.

10.1.1.4 Fetching Device Information

Before starting either the backup or replication operation, "3.4.4 Fetching device information on a Managed Server" that is managed by all Managed Servers using the Web Console. The time taken to complete this operation is proportional to the total number of devices defined on the selected Managed Servers. If many devices are involved, perform this operation while the CPU or I/O load on the system is low.

As a guide, about 0.5 seconds are required per device (i.e., partition) when there is no load on the system.

10.1.1.5 Maximum Number of Simultaneously Executing Operations

An upper limit (maximum value) applies to the number of Advanced Copy sessions in one Logical Unit (LU) that can be executed simultaneously. Ensure the number of volumes in one LU is the same or less than the maximum number of copy sessions per LU.

When you create more volumes in one LU than the maximum number of copy sessions per LU, for the volumes in that LU, an operational configuration is required that does not implement Advanced Copy simultaneously in excess of the maximum number of copy sessions per LU.

Refer to the ETERNUS Disk storage system SA/SE Handbook for the maximum number of copy sessions per LU.

10.1.1.6 Mounting of disks from other servers

Before backup/restoration or replication of a volume that can be mounted from multiple servers, unmount and cancel any pending mount on the other servers.

For disks that need not be mounted from other servers, prevent multiple servers from detecting or accessing the same partition by setting up hardware, such as devices in the ETERNUS Disk storage system and Fibre Channel switches, accordingly.

10.1.1.7 Notes on copy operations targeted at file systems

For a copy operation targeted at a file system, unmount the relevant volume to prevent data access and maintain the data integrity.
If backup/restoration and replication are performed without unmounting the volume, an error occurs since a volume in use cannot be unmounted.

Note the following about processing where unmounting is enabled:

- Another application is not using the volume. If another application is using it, suspend the application temporarily.

- Users are not using the volume. If users are using it, prohibit them from using it temporarily.

- Another volume is not mounted on the volume. If it is mounted, unmount it temporarily.

- The volume is not shared with NFS sharing. If it is shared, unshare it temporarily.

Unmounting is required only during command execution. Normal operations can be resumed after command execution.

10.1.1.8 Mirroring in GDS

For box-to-box mirroring, both systems of the mirroring must be backed up if restoration is required by OPC in the event of a box failure.
In these cases, the size of the backup volume must be a physical volume size, not logical volume size.

10.1.1.9 Notes on executing Logical Unit (disk) unit copy

Logical Unit (disk unit) copy can be used for replication operations.
Note the following points if Logical Unit (disk unit) copy is executed:

Inter-server copy

Both servers must run Version 13.1 or later.

Copying between Logical Units (disks) of different sizes

When copying between Logical Units (disks) of different sizes, the physical copy matches the size of the smaller Logical Unit (disk). If the copy direction is from the smaller Logical Unit (disk) to the larger Logical Unit (disk), part of the area on the larger disk is not physically copied to, and data from before the copy operation remains. If this remaining area contains the Logical Unit (disk) management information, partitioning tools may identify the disk format incorrectly. Therefore, format the copy destination Logical Unit (disk) before copying to it.

About GPT disks

To copy the GPT disks by each Logical Unit (disk), create the source and destination volumes in the same size for replication in the unit of bytes.
When copying between volumes of different sizes, the GPT label of the destination volume may be changed. If you restore the destination volume whose GPT label is changed, you cannot use the destination volume.

Notes when executing copy

- If the replication source/destination volume areas for replication management, and the transaction/backup volume areas for backup management duplicate each other, replication management and backup management operations cannot be performed simultaneously for volumes that have duplicate areas.

- Execute copy operations when the partitions included in the Logical Unit (disk) are in the unmounted state. If copy is executed when the partitions are in the mounted state, the partitions included in the copy destination disk may not be able to be made consistent as a file system.

SDX objects

Logical Unit (disk unit) copying cannot be used for SDX objects.

10.1.1.10 Notes on executing copy between different operating systems

Copies can be made between different operating systems by using replication operations.
Use the following combinations to execute copy between different operating systems:

Table 10.2 Combinations for executing copy between different operating systems

Copy source	Copy destination	
	slices	Logical Unit (disk)
Partition	A	N/A
Logical Unit (disk)	A	N/A

(A=Available, N/A=Not Available)

Note

If SnapOPC+ is to be used for executing copy operations between different operating systems, perform copy between operating systems that can use SnapOPC+.

10.1.1.11 Supported disk formats

The MBR disk format (recommended) and GPT disk format are supported.

10.1.1.12 Operations that use the device names (by-id and by-path) generated by the udev facility

Note the following points if AdvancedCopy Manager operations use device names generated by the udev facility.
The device names used when a file system is mounted must be the same as the device names used by AdvancedCopy Manager.

- If AdvancedCopy Manager operations use by-id, use by-id to mount the file system.

- If AdvancedCopy Manager operations use by-path, use by-path to mount the file system.

- If the udev name is used for replication between servers, the AdvancedCopy Manager version must be the same as on the other server.

- If the other server has the following operating systems, the udev name cannot be used for replication between servers:

    - HP-UX

    - AIX

10.1.1.13 Notes when using the -m option

When replication occurs between servers, the -m option can be specified in the following commands in order to avoid communication processing to non-operating servers:

- swsrpstartsync (Synchronous processing start command)

- swsrpmake (Replication creation command)

- swsrpcancel (Replication cancellation command)

When the -m option is specified, any one of the following conditions must apply in order to ensure that problems do not arise as a result of not performing volume pre-processing and post-processing at the non-operating server:

- The non-operating server has been shut down.

- Manual pre-processing and post-processing is performed for the volumes on the non-operating server.

- The target volume is a RAW device (Oracle, etc., that does not need pre-processing and post-processing), rather than a file system.

10.1.1.14 Copy processing performance of Advanced Copy

The Advanced Copy processing is performed by the ETERNUS Disk storage system. Therefore, direct enquiries concerning copy processing performance to the support department for the ETERNUS Disk storage system.

10.1.1.15 Mount by Label

In the case where a source volume has a filesystem label and where the following conditions hold, Advanced Copy can cause the status where multiple volumes with the same label exist in one system.

a. Backup/restore or replication within a server is employed.

b. Replication between servers is employed and a volume in the destination server has the same label as a source volume.

When there are multiple volumes with the same label, processing such as mount may not work correctly.
To avoid this phenomenon, specify not a label but a device (e.g. /dev/sda1) for /etc/grub.conf and /etc/fstab.

10.1.2 Notes on snapshot high-speed backup

This section provides notes on snapshot backup operation.

Number of Required Backup Volumes

Snapshot high-speed backup requires as many backup volumes as the number of specified backup generations. Accordingly, snapshot backup is disabled if:

- All the specified number of generations have been backed up, and

- There is no new volume that can be used as a backup volume.

Simultaneous Backups

Performing snapshot high-speed backup simultaneously with the same transaction volume specified will perform two backups in parallel. However, performing snapshot high-speed backup simultaneously more often than the number of preservation generations will cancel the backup operations starting from the oldest one. In other words, you cannot perform more simultaneous backups than the specified number of generations.

Notes on combining cascade copy and Concurrent OPC in operation

If you specify a volume used in cascade copy as the copy source and then execute a Concurrent OPC by swsrpmake (Replication creation command), it may not be possible to preserve the sequence of data updates across multiple volumes.
In this case, an error message will be displayed and the command terminates in an error.

If the error cannot be resolved by responding to the message, we recommend you change the operating method of Advanced Copy or reconfigure the system.

If you cannot change the operating method of Advanced Copy or reconfigure the system, create a file (check.ini) that sets the sequence of data updates across multiple volumes. Then, disable sequence preservation before you retry Concurrent OPC. The sequence of data updates will not be preserved across multiple volumes.

Create check.ini in the following path of the server on which to execute swsrpmake (Replication creation command):

- In a non-cluster environment

```
/etc/opt/FJSVswsrp/data/DEFAULT
```

- In a cluster environment

```
/etc/opt/FJSVswsrp/<logical-node-name>/data/DEFAULT
```

Table 10.3 File format of the check.ini

Section name	Key name	Value
[check]	ConcurOPCCascade	Set preservation of the sequence of data updates across multiple volumes. By setting "No", preservation of the sequence will be disabled. By setting any value other than "No", preservation of the sequence will be enabled.

[Example for disabling sequence preservation]

```
[check]
ConcurOPCCascade=No
```

P Point

For any of the below ETERNUS Disk storage systems, Concurrent OPC can be executed regardless of sequence preservation of data updates. However, the sequence of data updates will not be preserved across multiple volumes.

- ETERNUS DX80/DX80 S2 (for firmware version numbers before V10L30)

- ETERNUS DX90/DX90 S2 (for firmware version numbers before V10L30)

- ETERNUS DX400/DX400 S2 series (for firmware version numbers before V10L30)

- ETERNUS DX8000/DX8000 S2 series (for firmware version numbers before V10L30)

10.1.3 Notes on synchronous high-speed backup

This section provides notes on synchronous high-speed backup.

The processing of the backup

The swstbackup (Backup execution command) cannot be executed until the transaction and backup volumes are in an equivalent state.

Backup policy settings

A backup policy can still be set up if the number of registered volumes is enough for snapshot type high-speed backup, even if the number is not enough for synchronous high-speed backup. In this case, synchronous high-speed backup cannot be performed.

10.1.4 Notes on restoration

This section provides notes on restoring data.

Timing of Restoration

If data on a transaction volume has been updated since the latest backup, then if a restore is performed from the backup volume, recently updated data will be overwritten.

Figure 10.4 Notes on restoration of AdvancedCopy Manager

TD: Transaction volume
BD: Backup volume

10.1.5 Notes on SDX object operations

10.1.5.1 Notes on SDX object operations (backup management)

Note the following about the backup operation of SDX objects.

swstdevinfoset (Device information setting command)

Note the following about operation with an SDX object in units of logical volumes:

- An SDX object cannot be set as a transaction volume for operation of an SP-series product.

- A volume in a disk group to which the transaction volume belongs cannot be registered as a backup volume.

- A volume in a single disk to which the transaction volume belongs cannot be registered as a backup volume.

- A volume in a disk group to which an associated backup volume belongs cannot be registered as the transaction volume.

- A volume in a single disk to which an associated backup volume belongs cannot be registered as the transaction volume.

swststartsync (Backup synchronous processing start command)

Note the following about operation with an SDX object in units of logical volumes.

Backup synchronous processing of a GDS transaction volume cannot be started in the following circumstances:
In the case of any of the above, you should take appropriate action according to the [System administrator's action] for the error message as displayed.

- The volume structure of the transaction volume is changed.

- The transaction volume is set as a proxy object.

- A backup volume cannot be stopped.

- The slices that make up the transaction volume include a temporarily separated slice or a slice being copied.

Backup synchronous processing with EC cannot be started in the following circumstances:
In these cases, only software-copy with GDS is available. To use the software-copy function of GDS, specify the -Xgds-softcopy option.

- An SDX object that does not support the EC function is specified.

- All available backup volumes are in a mirroring structure.

- The transaction volume belongs to a mirror group connected to any subordinate group.

- An SDX object is in the root class.

If a backup volume satisfies any of the following conditions, it is not selected as the backup destination because it cannot be used as a proxy object of the transaction volume. Even if it is specified as the backup destination, it cannot be used.

- The backup volume is not the same class as the transaction volume.

- The slices that make up the backup volume include a temporarily separated slice or a slice that is being copied.

- The backup volume is set as the master object associated with another proxy of GDS.

- The backup volume is mounted.

- The backup volume is a logical volume whose size is different from that of the transaction volume.

- The backup volume is in a mirroring structure (excluding cases where -Xgds-softcopy or -Xgds-selectcopy option is specified).

- No disk is connected directly to the backup volume (excluding cases where -Xgds-softcopy or -Xgds-selectcopy option is specified).

swstcancelsync (Backup synchronous processing cancel command)

Note the following points about use with an SDX object in units of logical volumes.

The backup volume enters the invalid data state (INVALID) if synchronous processing is cancelled in the following circumstances during the backup of an SDX object:

- When the synchronous processing is not in the equivalency maintenance state or copy established state

Restart synchronous processing or snapshot processing again, and the volume should enter the stop state (STOP), which will automatically clear the data invalid state.

To use an alternative method to restore the state, refer to Section F.1.3, "Errors Related to Volume States", in the "PRIMECLUSTER(TM) Global Disk Services User's Guide (Linux)" and perform appropriate action.

swstbackup (Backup execution command)

Note the following points about use with an SDX object in units of logical volumes.

Backup of a GDS transaction volume cannot be performed in the following circumstances:
In the case of any of the above, you should take appropriate action according to the [System administrator's action] for the error message as displayed.

- The volume structure of the transaction volume is changed.

- The transaction volume is set as a proxy object.

- A backup volume cannot be stopped.

- The slices that make up the transaction volume include a temporarily separated slice or a slice being copied.

- An SDX object that does not support the OPC function is specified.

- All available backup volumes are in a mirroring structure.

- The transaction volume belongs to a mirror group connected to any subordinate group.

- An SDX object is in the root class.

If a backup volume satisfies any of the following conditions, it is not selected as the backup destination because it cannot be used as a proxy object of the transaction volume even if it is specified as the backup destination.

- The backup volume is not the same class as the transaction volume.

- The slices that make up the backup volume include a temporarily separated slice or a slice being copied.

- The backup volume is set as the master object associated with another proxy of GDS.

- The backup volume is mounted.

- The backup volume is a logical volume whose size is different from that of the transaction volume.

- The backup volume is in a mirroring structure.

- No disk is connected directly to the backup volume.

Backup with the output destination backup volume specified cannot be performed in the following circumstances:

- An ordinary volume is specified as the transaction volume and a GDS logical volume as the backup volume.

- A GDS logical volume is specified as the transaction volume and an ordinary volume as the backup volume.

Note the following about operation with an SDX object in units of slices.

Backup cannot be performed in the following circumstances:
In the case of any of the above, you should perform appropriate action according to the [System administrator response] for the error message as displayed.

- An SDX object does not satisfy the following conditions:

- The state of the mirror volume is ACTIVE.

- The state of the SDX disk is ENABLE.

- The state of the mirror slice is ACTIVE or TEMP.

Pre-processing and post-processing scripts cannot be executed in the following circumstances

- The state of the SDX object slice is TEMP.

swstrestore (Restore execution command)

Note the following about operation with an SDX object in units of logical volumes.

A GDS transaction volume cannot be restored in the following circumstances:
In the case of any of the above, you should take appropriate action according to the [System administrator response] for the error message as displayed.

- In restoration to a transaction volume, the backup volume is not the proxy object of the transaction volume.

- In restoration to a transaction volume, the transaction volume cannot be stopped.

- In restoration to a transaction volume, the target backup volume cannot be stopped.

Restoration with OPC cannot be performed for a GDS transaction volume in the following circumstances:
In these cases, only software-copy is available. To use the software-copy function of GDS, specify the -Xgds-softcopy option.

- A GDS logical volume that does not support the OPC function is specified.

- The backup volume belongs to a mirror group connected to any subordinate group.

- An SDX object is in the root class.

If swstrestore (Restore execution command) using the software-copy function of GDS is suspended, such as for a forced termination, the started software-copy processing continues. In this case, wait for the software-copy processing to finish, and then execute swstsrsemtch (Resource match command).

Note the following about operation with an SDX object in units of slices.

If the transaction volume is a slice of an SDX object and it is not a single volume, restoration using this command cannot be performed. This is because advanced copy cannot be performed on a mirror volume.
For restoration to the transaction volume of an SDX object, follow the procedure below.

1. Confirm the backup volume by using the backup history list view swsthistdisp (History information display command).

2. For restoration, use the "dd" command to copy data to the volume.

10.1.5.2 Notes on SDX object operations (replication management)

Note the following about the replication operation of SDX objects.

swsrpsetvol (Replication volume information setting command)

Note the following about operation with an SDX object in units of logical volumes.

Replication volume information for the replication operation of an SDX object cannot be set in the following circumstances:

- An SDX object and slice, or an SDX object and a volume of another OS are specified.

- A logical volume and the n option are specified.

- The original volume and replica volume belong to different classes.

- The original volume and replica volume belong to the same group; alternatively, in operation of single volumes, they belong to the same single disk.

swsrpdelvol (Replication volume information deletion command)

Note the following about operation with an SDX object in units of logical volumes.

If the original volume and replica volume are bound, or if they are separated when a copy session exists, the replication volume information cannot be deleted.

swsrpstartsync (Synchronous processing start command)

Note the following about operation with an SDX object in units of logical volumes.

If any of the following is true, swsrpstartsync (Synchronous processing start command) in the replication operation of an SDX object cannot be executed:

- The copy source volume and copy destination volume are in the master-proxy relationship when synchronous processing is started.

- The copy source volume and copy destination volume have not been separated from their master-proxy relationship when synchronous processing is resumed.

- The copy source volume and copy destination volume belong to different classes when synchronous processing is started.

- The copy source volume and copy destination volume belong to the same disk group or single disk.

- The copy source volume satisfies any of the following conditions:

    - The volume contains any temporarily separated slice or slice that is being copied.

    - The copy source volume is a proxy object when synchronous processing is started.

    - Copy processing with another volume is in progress.

- The copy destination volume satisfies any of the following conditions:

    - The volume is in a mirroring structure and the -Xgds-softcopy or -Xgds-selectcopy option is not specified when synchronous processing is started. The volume cannot be stopped.

    - The volume is a master volume when synchronous processing is started.

    - The volume contains any temporarily separated slice or slice that is being copied.

    - The volume is a proxy of another master.

- The hardware copy function is not available (excluding cases when the -Xgds-softcopy or -Xgds-selectcopy option is specified). The typical causes are:

    - EC/REC of the ETERNUS Disk storage system is not properly set.

    - The copy source volume belongs to a mirror group connected to any subordinate group.

    - The copy destination volume is in a mirroring structure.

    - The volume is in the root class.

In the replication operation of an SDX object registered with the replication function, do not use the "sdxproxy" command directly to change the state of the SDX object. Otherwise, the "sdxproxy" command may create an inconsistency between the management information of the replication function and the actual object state. In this case, use swsrpcancel (Replication cancellation command) to cancel the replication operation.

Note the following about operation with an SDX object in units of slices.

Note the following for the replication operation of an SDX object:
Replication cannot be started in the following cases:

- The copy destination volume is an SDX object and the status of the SDX object is not as follows:

    - Mirror volume is either ACTIVE or STOP

    - Mirror slice status is either ACTIVE or STOP

- If the copy destination volume is an SDX object that is not a single volume, replication using this command cannot be performed. Therefore, if either the original volume or replica volume is an SDX object that is not a single volume, you can perform replication only from an SDX object to a non-SDX object.

- To copy data to a volume that is an SDX object but not a single volume, use the dd command.

swsrpmake (Replication creation command)

Note the following about operation with an SDX object in units of logical volumes.

If any of the following is true, replication in the replication operation of an SDX object cannot be created:

- The copy source volume and copy destination volume are in the master-proxy relationship in snapshot replication.

- The copy source volume and copy destination volume belong to different classes in snapshot replication.

- The copy source volume and copy destination volume belong to the same disk group or a single disk in snapshot replication.

- The copy source volume satisfies any of the following conditions:

    - The volume contains any temporarily separated slice or slice that is being copied, in snapshot replication.

    - The copy source volume is a proxy object in snapshot replication.

    - Copy processing with another volume is in progress.

- The copy destination volume satisfies any of the following conditions:

    - The volume is in a mirroring structure in snapshot replication.

    - The volume could not be stopped in snapshot replication.

    - The volume is a master volume in snapshot replication.

    - The volume contains any temporarily separated slice or slice that is being copied, in snapshot replication.

    - The volume is a proxy of another master in snapshot replication.

- The OPC function is not available for snapshot replication. The typical causes are:

    - OPC of the disk array unit is not properly set.

    - EC/REC in the equivalency maintenance state or suspended state is between the specified copy source volume and another proxy. However, advanced copy can be set in units of mirror slices, which makes it possible to set advanced copy for as many copy destination volumes as the number of mirror slices of the copy source volume at the same time. Therefore, this error occurs as the result of an attempt to set advanced copy for a greater number of copy source volumes than the number of mirror slices.

    - The copy source volume belongs to a mirror group connected only to some subordinate groups but not directly to a disk. The copy source volume belongs to a mirror group connected to any subordinate group.

    - The copy destination volume is in a mirroring structure.

    - The volume is in the root class.

In the replication operation of an SDX object registered with the replication function, do not use the "sdxproxy" command directly to change the state of the SDX object. Otherwise, the "sdxproxy" command may create an inconsistency between the management information of the replication function and the actual object state. In this case, use swsrpcancel (Replication cancellation command) to cancel the replication operation.

Note the following about operation with an SDX object in units of slices.

Replication cannot be performed in the following circumstances:

- The copy source volume is an SDX object for replication of an SDX object, and the SDX object is not in the following states:

    - The state of the mirror volume is either ACTIVE or STOP.

    - The state of the mirror slice is either ACTIVE or TEMP.

    - The state of the SDX disk is ENABLE

- The copy destination volume is an SDX object, and the SDX object is not in the following states:

    - The state of the mirror volume is either ACTIVE or STOP.

    - The state of the mirror slice is either ACTIVE or STOP.

- If the copy destination volume is an SDX object that is not a single volume for the replication operation for an SDX object, replication using this command cannot be performed. Therefore, if either the original volume or replica volume is an SDX object that is not a single volume, you can perform replication only from an SDX object to a non-SDX object.

- To copy data to a volume that is an SDX object but not a single volume, use the dd command.

Preprocessing of the copy source volume is not performed in the following circumstances.

- The state of the slice of the SDX object is TEMP.

swsrpcancel (Replication cancellation command)

Note the following about operation with an SDX object in units of logical volumes.

If any of the following is true, replication in the replication operation of an SDX object cannot be cancelled:

- The pair of the target replication volume used for replication cancellation is in the equivalency maintenance state, and copy processing between the copy source volume and another volume is in progress.

If replication cancellation is executed at either of the following times during the replication operation of an SDX object, the copy destination volume enters the invalid data state (INVALID):

- When snapshot replication is neither in the equivalency maintenance state nor copy established state

- When data copy processing of snapshot processing is in progress

Restart synchronous processing or snapshot processing, and the volume will enter the stop state (STOP), which will automatically clear the data invalid state.

To restore the state using an alternative method, refer to Section F.1.3, "Errors Related to Volume Statuses", in the "PRIMECLUSTER(TM) Global Disk Services User's Guide (Linux)", and perform appropriate action.

Note the following about operation with an SDX object in units of slices.

Synchronous processing in the equivalency maintenance state cannot be stopped in the following circumstances:

- The copy source volume is an SDX object, and the SDX object is in any of the following states:

  - The state of the mirror volume is either ACTIVE or STOP.

  - The state of the mirror slice is either ACTIVE or TEMP.

  - The state of the SDX disk is ENABLE.

    Note: If the mirror slice is in the TEMP state, pre-processing of the original volume is not performed.

10.1.6 Notes on cluster operation

Note the following general points on cluster operation:

- Do not set a SWSTGNODE registration that configures an environment variable to the entire system.

- AdvancedCopy Manager does not monitor daemons of an active server. If the daemon of a transaction stops for some reason, AdvancedCopy Manager does not automatically restart it, so you must restart it manually. Moreover, the stopping of an active server daemon of AdvancedCopy Manager does not cause the cluster to fail over.
  For details on the daemons of AdvancedCopy Manager active servers and how to start these daemons, refer to "Chapter 2 Starting and Stopping Daemons".

For a Managed Server daemon (or Management Server daemon) operating together with a daemon in a scalable operation, note the following points that do not apply to regular operation or cluster operation:

- For a Management Server transaction, note that the logical IP address you specify in the cluster setup command must be the logical IP address of the Management Server transaction. Specifying the physical IP address or logical IP address of a server in scalable operation will result in incorrect setting.

- Do not register any cluster resource of a transaction in scalable operation as a cluster resource of a Managed Server transaction (Management Server transaction).

- When specifying the Managed Server's IP address for the execution of AdvancedCopy Manager commands, specify the logical IP address for the Managed Server transactions. Specifying the physical IP address or logical IP address of a server in scalable operation will result in an incorrect setting.

- When specifying the Managed Server's port number for the execution of AdvancedCopy Manager commands, specify the port number for the Managed Server transaction communication daemon (the name of the stgxfws logical node). Specifying a port number besides those applicable to server transaction communication daemons will result in an incorrect setting.

- Please add backup-use volume with the transaction volume for scalable operation transactions which operate backup and replication.

- You can make a backup or replication on a node operating a Managed Server transaction (Management Server transaction) if it is the active node. You cannot make a backup or replication on any standby node of a Managed Server transaction (Management Server transaction).

- If a transaction (in scalable operation) running on a node operating a Managed Server transaction (Management Server transaction) stops for some reason, you must manually switch the Managed Server transaction (Management Server transaction) over to a node running in scalable operation. The Managed Server transactions (Management Server transactions) are not linked with scalable

operation transactions, thus stopping scalable operation transactions does not cause a Managed Server transaction (or Management Server transaction) to fail over.

10.1.7 Notes on backup in a clustered system

PRIMECLUSTER

For backup operation in a clustered system, note the following points that do not apply to regular operation:

- If failover occurs while running a backup command, you must secure consistency by using swstsrsemtch (Resource match command) as a recovery measure.

- You must perform backup operation on the active server of the userApplication because AdvancedCopy Manager is placed in the userApplication and operates as part of said userApplication.
  You cannot make a backup from a standby node or another userApplication.

PRIMECLUSTER 4.1A20

Note the following points if a file system has been created on a transaction volume and the file system has been registered as a cluster resource:

- Change the backup pre-processing so that the transaction volume is not unmounted. Refer to "A.2 Pre-processing and Post-processing of Backup" for information about changing backup pre-processing.

- It is not possible to restore data using swstrestore (Restore execution command) during transaction operations. To restore data on the volume, either copy the data using the dd command, or mount the backup volume and copy the necessary data. Alternatively, perform the procedure described in "10.1.7.1 Backup and restoration during the stoppage of a cluster".

PRIMECLUSTER 4.1A30 or later

If a file system has been created on a transaction volume, and the file system has been registered as a cluster resource, use the following procedure to back up and restore data.

Follow the procedure below to perform backup/restore.

Snapshot-type backup/restore procedure using AdvancedCopy Manager in a PRIMECLUSTER environment

1. Switch userApplication to Maintenance mode.

   ```
   # hvutil -m on userApplication
   ```

2. Unmount the transaction volume (this is mandatory to backup or restore a transaction volume in the file system).

3. If the transaction volume was unmounted, then make sure that the cluster resources of the transaction volume are Offline (wait until the return value of the command is 0):

   ```
   # hvassert -s <transaction_volume> <cluster_resource_name> Offline
   ```

4. Perform backup or restore

   ```
   # /opt/FJSVswsts/bin/swstbackup Device-Name
   Device-Name swstbackup completed
   ```

   or

   ```
   # /opt/FJSVswsts/bin/swstrestore Device-Name
   Device-Name swstrestore completed
   ```

5. If the transaction volume was unmounted, then mount it back.

6. Make sure that the cluster resources of the transaction volume are Online (wait until the return value of the command is 0):

```
# hvassert -s <transaction_volume> <cluster_resource_name> Online
```

7. Stop the Maintenance mode for the userApplication.

```
# hvutil -m off userApplication
```

Synchronous backup procedure using AdvancedCopy Manager in a PRIMECLUSTER environment

1. Start synchronization processing.

```
# /opt/FJSVswsts/bin/swststartsync Device-Name
Device-Name swststartsync completed
```

2. Check the equivalency maintenance status.

3. Swith userApplication to Maintenance mode.

```
# hvutil -m on userApplication
```

4. Unmount the transaction volume (this is mandatory to backup/restore a transaction volume in the file system).

5. If the transaction volume was unmounted , then make sure that the cluster resources of the transaction volume are Offline (wait until the return value of the command is 0):

```
# hvassert -s <transaction_volume> <cluster_resource_name> Offline
```

6. Perform backup.

```
# /opt/FJSVswsts/bin/swstbackup Device-Name
Device-Name swstbackup completed
```

7. If the transaction volume was unmounted, then mount it back.

8. Make sure that the cluster resources of the transaction volume are Online (wait until the return value of the command is 0):

```
# hvassert -s <transaction_volume_cluster_resource_name> Online
```

9. Stop Maintenance mode for userApplication.

```
# hvutil -m off userApplication
```

📒 Note

- In either procedure, if the transaction volume is unmounted, then the pre-processing and post-processing scripts will not be executed (they are only executed when the transaction or replication source volumes are mounted).
  Take this into account if the applicable parts of the pre-processing and post-processing scripts are being customized.

- The cluster detector detects the status of cluster resources either periodically or asynchronously. Therefore, it may take a few moments for the cluster resource status to be reflected, depending on the system load, and on the timing of mounting/unmounting the transaction volume. This means that the status may not be able to be checked using the hvassert command.
  Due to this fact, note the following about the timing of using hvassert to check the cluster resource status:

  1. Once the mounting/unmounting of the file system of a cluster FSystem resource has finished, only use hvassert after the time elapsed is equal to or greater than the cluster resource monitoring interval.
     You can specify the maximum time to wait until the cluster reaches the desired status by using hvassert with the -w option:

```
# hvassert -w <replication_vol> Offline|Online <max_wait_time(secs)>
```

2. When using hvassert with the -w option, the system load and cluster resource reflection processing time must also be taken into account.
   Due to that, make sure to specify a value that suits the system operating environment, purpose, and requirements.

In PRIMECLUSTER 4.2, the default monitoring interval for the PRIMECLUSTER Fsystem resource is 10 seconds - check the default value for your version.
For details on the hvassert command, refer to the "PRIMECLUSTER Installation and Administration Guide", section "hvassertt(1M)".

- -

VERITAS Cluster Server

For backup operation in a clustered system, note the following points that do not apply to regular operation:

- If failover occurs while running a backup command, you must ensure consistency by using swstsrsemtch (Resource match command) as a recovery measure.

- To use synchronous high-speed backup, do not register a backup to any cluster transaction.

- You must perform backups on the active server of the service group because AdvancedCopy Manager is placed into the service group and operates as a part of it.
  You cannot make a backup from a standby node or another service group.

- To perform backup/restoration using AdvancedCopy Manager, you usually must unmount transaction volumes beforehand and mount the volumes later.
  If the mount point of transaction volumes have been registered to a cluster transaction, set the mount point resource online or offline from the cluster transaction instead of mounting or unmounting the volumes. Alternatively, modify the backup/restoration pre-processing and post-processing scripts so that the transaction volumes are set offline or online instead of being unmounted or mounted. Moreover, the Volume is immediately placed offline/online for a mount point resource whereas there will be a time delay when mount/unmount is used. Therefore, please add into the pre-processing script after the statement where the success or failure value of offline/online is returned, the df or sleep command to function until mount/unmount is actually complete.

10.1.7.1 Backup and restoration during the stoppage of a cluster

PRIMECLUSTER

When a userApplicationto which AdvancedCopy Manager belongs is operating, backup operation is enabled only on the active node. The standby node environment is not ready (to provide necessary resources) for backup operation, thus disabling backup operation. Likewise, if the userApplicationhas been stopped, the environment is not ready and disables backup operation.
However, only when the userApplicationhas been stopped is backup operation enabled by temporarily arranging the minimum essentials.

📖 Note

- -

- When transactions in scalable operation are handled together, only stop userApplicationin scalable operation; do not stop userApplicationfor Managed Server (or Management Server) transactions.
  Enabling backup operation disks (transaction volume and backup volume to operate) allows you to perform backup operation in the same way as for regular operations.
  If a backup operation disk cannot be enabled, you cannot perform backup operation while user applications in scalable operation are stopped.

- If any of the following types of shared disks cannot be enabled (i.e., set online or mounted in a system), backup operation is disabled:

  - Shared disk for AdvancedCopy Manager shared data

  - Backup operation disk (transaction volume/backup volume to operate)

- The following operations are not allowed in backup operation:

  - Backup/restoration through Web Console

  - Manipulation for backup operation from a Management Server (host specification with the -h option)

- Addition or deletion of a transaction volume/backup volume, or change in policy

- To execute a display command for backup management to a Management Server transaction that also functions as a Managed Server transaction, you must specify options for the command. For details on the display command for backup management, refer to the "9.2 Backup Management Commands".

. . . .

Perform backup operation while a userApplicationis stopped by taking the following steps:

1. Confirm that the userApplicationhas been stopped on both nodes.
   For details on how to stop a userApplication, refer to the PRIMECLUSTER manual.

2. Log in to the node on which you want to make a backup through telnet, for example.
   You cannot use the logical IP address. Directly use a node that is operable by using the physical IP address.

3. Enable the shared disk.
   Start (set online) the shared disk.
   Mount the shared disk for AdvancedCopy Manager shared data.
   When the transaction volume is a file system, mount it.

📄 Note
. . . .

Be sure to enable the shared disk from either node. Do not enable the shared disk from both nodes.

. . . .

4. In the SWSTGNODE environment variable, set the logical node name for the relevant task.
   Set the environment variable as shown below.

   <Execution example>

   ```
   # SWSTGNODE=logical-node-name
   # export SWSTGNODE
   ```

5. Perform backup operation.
   You can perform backup operation in the same way as for regular operations for cluster operation.
   You can use swstrestore (Restore execution command) to perform restoration.

6. Release all the shared disks you enabled in Step 3.
   Unmount the file systems you mounted.
   Stop (set offline) the shared disk.

7. Start (set online) the userApplication.
   Start the userApplication, if necessary.
   For details on how to start a userApplication, refer to the PRIMECLUSTER manual.

VERITAS Cluster Server

When a service group to which AdvancedCopy Manager belongs is operating, backup is enabled only on the active node. The standby node environment is not ready (to provide necessary resources) for backup operation, therefore backup is disabled. Likewise, if the service group has been stopped, the environment is not ready and disables backup operation.
However, backup is temporarily enabled when the service group is stopped.

📄 Note
. . . .

- If any of the following types of shared disks cannot be enabled (i.e., set online or mounted in a system), backup is disabled:

   - Shared disk for AdvancedCopy Manager shared data

   - Backup operation disk (transaction volume/backup volume to operate)

- The following operations are not allowed in backup:

    - Backup/restoration through Web Console

    - Manipulation for backup operation from a Management Server (host specification with the -h option)

    - Addition or deletion of a transaction volume/backup volume, or change in policy

- To execute a display command for backup management to a Management Server transaction that also functions as a Managed Server transaction, you must specify options for the command. For details on the display command for backup management, refer to "9.2 Backup Management Commands".

. . . .

If the service group has been stopped, perform a backup by taking the following steps:

1. Confirm that the service group has been stopped on both nodes.
   For details on how to stop a service group, refer to the VERITAS Cluster Server manual.

2. Log in (for example, using Telnet) to the node on which you want to make a backup.
   You cannot use the logical IP address. Use a node that is operable via the physical IP address.

3. Enable the shared disk.
   Start (set online) the shared disk.
   Mount the shared disk for AdvancedCopy Manager shared data.
   When the transaction volume is a file system, mount it.

Note

. . . .

Be sure to enable the shared disk from either node. Do not enable the shared disk from both nodes.

. . . .

4. In the environment variable SWSTGNODE, set the logical node name for the relevant task.
   Set the environment variable as shown below.

   <Execution example>

   ```
   # SWSTGNODE=logical-node-name
   # export SWSTGNODE
   ```

5. Perform the backup.
   You can perform the backup in the same way as for regular operations in a cluster environment.
   You can use swstrestore (Restore execution command) to perform restoration.

6. Release all the shared disks you enabled in Step 3.
   unmount the file systems you mounted.
   Stop (set offline) the shared disk.

7. Start (set online) the service group.
   Start the service group, if necessary.
   For details on how to start a service group, refer to the VERITAS Cluster Server manual.

10.1.8 Notes on replication in a clustered system

PRIMECLUSTER

For replication operation in a clustered system, note the following points that do not apply to regular operation:

- If failover occurs while running a replication command, you must secure consistency by using swsrprecoverres (Resource adjustment command) as a recovery measure.

- You must perform replication operation on the active system of the userApplicationbecause AdvancedCopy Manager is placed in the userApplicationand operates as part of said userApplication.
  You cannot perform replication from a standby node or another userApplication.

PRIMECLUSTER 4.1A20

Note the following points if a file system has been created on a copy source volume and the file system has been registered as a clustered system resource.

- Change the pre-processing and post-processing for replication so that the copy source volume is not unmounted.

- Set up a RAW partition on the copy destination volume.

- Data cannot be restored from the copy destination volume to the copy source volume using swsrpmake (Replication creation command) during transaction operations. To restore data on the volume, either copy the data using the dd command, or mount the copy destination volume and copy the necessary data. Alternatively, perform the procedure described in "10.1.8.1 Replication during the stoppage of a cluster".

PRIMECLUSTER 4.1A30 or later

Note the following points if a file system has been created on a copy source volume and the file system has been registered as a clustered system resource.

- Set up a RAW partition on the copy destination volume.

Follow the procedure below to perform replication.

Snapshot-type replication procedure using AdvancedCopy Manager in a PRIMECLUSTER environment

1. Switch userApplication to Maintenance mode.

   ```
   # hvutil -m on userApplication
   ```

2. Unmount the replication source volume (this is mandatory to replicate a replication source volume in the file system).

3. If the replication source volume was unmounted, then make sure that the cluster resources of the volume are Offline (wait until the return value of the command is 0):

   ```
   # hvassert -s <replication_source_volume> <cluster_resource_name> Offline
   ```

4. Perform replication.

   ```
   # /opt/FJSVswsrp/bin/swsrpmake <from_volume_name> <to_volume_name>
   FROM=<from_volume_name> TO=<to_volume_name> swsrpmake completed
   ```

5. If the replication source volume was unmounted, then mount it back.

6. Make sure that the cluster resources of the replication source volume are Online (wait until the return value of the command is 0):

   ```
   # hvassert -s <replication_source_volume> <cluster_resource_name> Online
   ```

7. Stop Maintenance mode for userApplication.

   ```
   # hvutil -m off userApplication
   ```

Synchronous replication procedure using AdvancedCopy Manager in a PRIMECLUSTER environment

1. Start synchronization processing.

```
# /opt/FJSVswsrp/bin/swsrpstartsync <from_volume_name> <to_volume_name>
FROM=<from_volume_name> TO=<to_volume_name> swsrpstartsync completed
```

2. Check the equivalency maintenance status.

3. Switch userApplication to Maintenance mode.

   ```
   # hvutil -m on userApplication
   ```

4. Unmount the replication source volume (this is mandatory to replicate a replication source volume in the file system).

5. If the replication source volume was unmounted, then make sure that its cluster resources are Offline (wait until the return value of the command is 0):

   ```
   # hvassert -s <replication_source_volume> <cluster_resource_name> Offline
   ```

6. Perform replication.

   ```
   # /opt/FJSVswsrp/bin/swsrpmake <from_volume_name> <to_volume_name>
   FROM=<from_volume_name> TO=<to_volume_name> swsrpmake completed
   ```

7. If the replication source volume was unmounted, then mount it back.

8. Make sure that the cluster resources of the replication source volume are Online (wait until the return value of the command is 0):

   ```
   # hvassert -s <replication_source_volume> <cluster_resource_name> Online
   ```

9. Stop Maintenance mode for userApplication.

   ```
   # hvutil -m off userApplication
   ```

📝 Note

- In either procedure, if the replication source volume was unmounted, then the pre- and post-processing scripts will not be executed (they are only executed when the transaction or replication source volumes are mounted).
  If the applicable parts of the pre- and post-processing scripts are being customized, take this point into account.

- The cluster detector detects the status of cluster resources either periodically or asynchronously. Therefore, it may take a few moments for the cluster resource status to be reflected, depending on the system load, and on the timing of mounting/unmounting the replication source volume. This means that the status may not be able to be checked using the hvassert command.
  Due to this fact, note the following about the timing of using hvassert to check the cluster resource status:

  1. Once the mounting/unmounting of the file system of a cluster FSystem resource has finished, only use hvassert after the time elapsed is equal to or greater than the cluster resource monitoring interval.

     You can specify the maximum time to wait until the cluster reaches the desired status by using hvassert with the -w option:

     ```
     # hvassert -w <replic_src_vol> <cluster_res_name> Offline|Online
     <max_wait_time(secs)>
     ```

  2. When using hvassert with the -w option, the system load and cluster resource reflection processing time must also be taken into account.
     Due to that, make sure to specify a value that suits the system operating environment, purpose, and requirements.

In PRIMECLUSTER 4.2, the default monitoring interval for the PRIMECLUSTER Fsystem resource is 10 seconds - check the default value for your version.
For details on the hvassert command, refer to the "PRIMECLUSTER Installation and Administration Guide", section "hvassertt(1M)".

VERITAS Cluster Server

For replication in a clustered system, note the following points that do not apply to regular operation:

- If failover occurs while running a replication command, you must ensure consistency by using swsrprecoverres (Resource adjustment command) as a recovery measure.

- When you use synchronous replication, do not register any destination volumes to a cluster transaction.

- You must perform replication on the active system of the service groups because AdvancedCopy Manager is placed in the service groups and operates as a part of it.
  You cannot perform replication from a standby node or another service group.

- To perform replication, you must unmount original/source volumes beforehand and mount the volumes later for some commands.
  If the mount points of original/source volumes have been registered to a cluster transaction, set the mount point resources online or offline from the cluster transaction instead of mounting or unmounting the volumes. Alternatively, modify the unmount and mount processes in the replication pre-processing and post-processing scripts to offline and online processes, respectively.
  Moreover, the Volume is immediately placed offline/online for a mount point resource whereas there will be a time delay when mount/unmount is used. Therefore, please add into the pre-processing script after the statement where the success or failure value of offline/online is returned, the df or sleep command to function until mount/unmount is actually complete.

10.1.8.1 Replication during the stoppage of a cluster

PRIMECLUSTER

When a userApplicationto which AdvancedCopy Manager belongs is operating, replication operation is enabled only on the active node.
The standby node environment is not ready (to provide necessary resources) for replication operation, thus disabling replication operation.
Likewise, if the userApplicationhas been stopped, the environment is not ready and disables replication operation.
However, only when the userApplicationhas been stopped is replication operation enabled by temporarily arranging the minimum essentials.

Note

- When transactions in scalable operation are handled together, only stop userApplicationsin scalable operation; do not stop userApplicationsfor Management Server transaction and Managed Server transactions.
  Enabling replication operation disks (source volume/destination volume to operate) allows you to perform replication operation in the same way as for regular operations.
  If a replication operation disk cannot be enabled, you cannot perform replication operation while a userApplicationsin scalable operation is stopped.

- If any of the following types of shared disks cannot be enabled (i.e., set online or mounted in a system), replication operation is disabled.

  - Shared disk for AdvancedCopy Manager shared data

  - Replication operation disks (source volume/destination volume to operate)

- The following operations are not allowed in replication operation:

  - Manipulation for replication operation from a Management Server (host specification with the -h option)

  - Addition or deletion of an original/replica volume

  - Inter-server replication without the -m option specified

- To execute a display command for replication management to a Management Server transaction that also functions as a Managed Server transaction, you must specify options for the command. For details on the display command for replication management, refer to the "9.4 Replication Management Commands".

Perform replication operation during stoppage of a userApplicationby taking the following steps:

1. Confirm that the userApplicationhas been stopped on both nodes.
   For details on how to stop a userApplication, refer to the PRIMECLUSTER manual.

2. Log in to the node on which you want to perform replication operation through telnet, for example.
   You cannot use the logical IP address. Directly use a node that is operable by using the physical IP address.

3. Enable the shared disk.
   Start (set online) the shared disk.
   Mount the shared disk for AdvancedCopy Manager shared data.
   When the original volume and replica volume have been mounted as file systems for operation, mount the volumes.

Note

Be sure to enable the shared disk from either node. Do not enable the shared disk from both nodes.

4. In the environment variable SWSTGNODE, set the logical node name for the relevant task.
   Set the environment variable as shown below.

   <Execution example>

   ```
   # SWSTGNODE=logical-node-name
   # export SWSTGNODE
   ```

5. Perform replication operation.
   You can perform replication operation in the same way as for regular operations for cluster operation.
   You can also perform restoration from the replica volume to the original volume.

6. Release all the shared disks you enabled in Step 3.
   Unmount the file systems you mounted.
   Stop (set offline) the shared disk.

7. Start (set online) the userApplication.
   Start the userApplication, if necessary.
   For details on how to start a userApplication, refer to the PRIMECLUSTER manual.

VERITAS Cluster Server

When a service group to which AdvancedCopy Manager belongs is operating, replication is enabled only on the active node. The standby node environment is not ready (to provide necessary resources) for replication, and replication is, therefore, disabled. Likewise, if the service group has been stopped, the environment is not ready and disables replication.
However, replication is temporarily enabled when the service group has been stopped.

Note

- If any of the following types of shared disks cannot be enabled (i.e., set online or mounted in a system), replication is disabled.

  - Shared disk for AdvancedCopy Manager shared data

  - Replication operation disks (source volume/destination volume to operate)

- The following operations are not allowed in replication:

  - Manipulation for replication from a Management Server (host specification with the -h option)

  - Addition or deletion of an original/replica volume

  - Inter-server replication without the -m option specified

- To execute a display command for replication management to a Management Server transaction that also functions as a Managed Server transaction, you must specify options for the command. For details on the display command for replication management, refer to "9.4 Replication Management Commands".

Perform replication during stoppage of the service group by taking the steps below.

1. Confirm that the service group has been stopped on both nodes.
    For details on how to stop a service group, refer to the VERITAS Cluster Server manual.

2. Log in (for example, through Telnet) to the node on which you want to perform replication.
    You cannot use the logical IP address. Use a node that is operable via the physical IP address.

3. Enable the shared disk.
    Start (set online) the shared disk.
    Mount the shared disk for AdvancedCopy Manager shared data.
    When the original volume and replica volume have been mounted as file systems for operation, mount the volumes.

![Note icon] **Note**

Be sure to enable the shared disk from either node. Do not enable the shared disk from both nodes.

4. In the SWSTGNODE environment variable, set the logical node name for the relevant task.
    Set the environment variable as shown below.

    <Execution example>

    ```
    # SWSTGNODE=logical-node-name
    # export SWSTGNODE
    ```

5. Perform replication.
    You can perform replication in the same way as for regular operations in a cluster environment.
    You can also perform restoration from the replica volume to the original volume.

6. Release all the shared disks you enabled in Step 3.
    Unmount the file systems you mounted.
    Stop (set offline) the shared disk.

7. Start (set online) the service group.
    Start the service group, if necessary.
    For details on how to start a service group, refer to the VERITAS Cluster Server manual.

Appendix A Pre-processing and Post-processing of Backup and Restoration

This appendix provides information on pre-processing and post-processing of backup and restoration.

A.1 Overview

Pre-processing and post-processing of backup or restoration are started before and after backup or restoration when a backup or restore command is executed.

This processing is required by AdvancedCopy Manager to back up or restore a transaction volume.
This appendix describes the content and setup of pre-processing and post-processing.

A.2 Pre-processing and Post-processing of Backup

Backup using AdvancedCopy Manager must be performed basically while access to a transaction volume from other processes is inhibited.

Normally, the pre-processing acquires the mount status of a transaction volume using the transaction volume name and then performs the following processing:

Table A.1 Pre-processing for transaction volume

Transaction volume status	Preprocessing
Mounted	Unmount a transaction volume.(*1)
Unmounted	Take no action.

*1 If the transaction volume cannot be unmounted, however, customize the shell script for pre-processing is customized accordingly. For information on this customization, refer to the section on "When you do not want to unmount a transaction volume".

Tasks performed in the post-processing are generally determined by what has been done in the pre-processing.

Table A.2 Backup pre-processing and post-processing

Preprocessing	Post-processing
A transaction volume was unmounted.	Remount the transaction volume.
Take no action.	Take no action.

If no transaction volume has been mounted since the start of system operation (e.g., a transaction volume used as a database), neither pre- nor post-processing is performed.

If special pre-processing or post-processing is required, you need to add the additional instructions to the shell script.

When customizing a script, strictly observe the following rules regarding error codes:

Table A.3 Error code rules

Error code	Usage
0-99	Unusable (these are reserved for AdvancedCopy Manager)
100-255	Usable

Note

If pre-processing failed, execute swstsrsemtch (Resource match command) because the consistency of resource information may be incomplete.

A.2.1 Pre-processing of backup

The name of a script file for pre-processing of a backup is as follows.

In the case of non-cluster operation

```
/etc/opt/FJSVswsts/sh/OpcBackup.pre
```

In the case of cluster operation

```
/etc/opt/FJSVswsts/<logic node name>/sh/OpcBackup.pre
```

A.2.1.1 Pre-processing script for backup

```
 1: #!/bin/sh
 2:
 3: # AdvancedCopy Manager
 4: # Copyright FUJITSU LIMITED, 2002-2012
 5:
 6: #
 7: #   Preprocessing of backup processing
 8: #
 9: #       Argument: $1 Device name of transaction disk
10: #                $2 Mount point of transaction disk
11: #                $3 Device name of backup disk
12: #
13: #  Error number
14: #           2: Argument error
15: #          10: umount error
16: #          13: Illegal mount type (bind/stack mount)
17:
18:
19: # Argument check
20: case $# in
21: 1)
22:    ;;
23: 2)
24:    ;;
25: 3)
26:    ;;
27: *)
28:    exit 2
29:    ;;
30: esac
31:
32: device="`echo $1`"
33: mount_point="`echo $2`"
34: bk_device="`echo $3`"
35:
36: # Determination postprocessing file name
37:
38: if [ "$SWSTGNODE" != "" ]
39: then
40: swstg_node="/`echo $SWSTGNODE`"
41: else
42: swstg_node=""
43: fi
44:
45: err_log_path="/var/opt/FJSVswsts"$swstg_node"/log"
46:
```

```
 47: if [ "`echo $device | /bin/grep "/dev/sd"`" != "" ]
 48: then
 49:     # /dev/sd? -> sd?
 50:     dev="`echo $device | /bin/sed "s/\/dev\///"`"
 51: elif [ "`echo $device | /bin/grep "/dev/vd"`" != "" ]
 52: then
 53:     # /dev/vd? -> vd?
 54:     dev="`echo $device | /bin/sed "s/\/dev\///"`"
 55: elif [ "`echo $device | /bin/grep "/dev/FJSV"`" != "" ]
 56: then
 57:     # /dev/FJSVmphd/dsk/mplb?s? -> mplb?s?
 58:     # /dev/FJSVmphd/dsk/mphd?s? -> mphd?s?
 59:     dev="`echo $device | /bin/cut -d/ -f5`"
 60: elif [ "`echo $device | /bin/grep "/dev/sfdsk/"`" != "" ]
 61: then
 62:     if [ "`echo $device | /bin/grep ":"`" != ""    ]
 63:     then
 64:             devnam="`echo $device | /bin/cut -d: -f2-`"
 65:             # /dev/sfdsk/class/dsk/volume:sd? -> class_volume_sd?
 66:             dev="`echo $device | /bin/sed "s/\/dev\/sfdsk\///"`"
 67:             dev="`echo $dev | /bin/sed "s/\/dsk\//_/"`"
 68:             dev="`echo $dev | /bin/sed "s/:/_/"`"
 69:             device="`echo $device | /bin/cut -d: -f1`"
 70:             if [ "`echo $devnam | /bin/grep "/dev/disk/by-id/"`" != "" ]
 71:             then
 72:                     # /dev/sfdsk/class/dsk/volume:/dev/disk/by-id/<device> ->
class_volume__by_id_<device>
 73:                     dev="`echo $dev | /bin/sed "s/\/dev\/disk\/by-id\//_by-id_/"`"
 74:             elif [ "`echo $devnam | /bin/grep "/dev/disk/by-path/"`" != "" ]
 75:             then
 76:                     # /dev/sfdsk/class/dsk/volume:/dev/disk/by-path/<device> ->
class_volume__by_path_<device>
 77:                     dev="`echo $dev | /bin/sed "s/\/dev\/disk\/by-path\//_by-path_/"`"
 78:             fi
 79:     else
 80:             # /dev/sfdsk/class/dsk/volume -> _gds_class_volume
 81:             dev="_gds_`echo $device | /bin/sed "s/\/dev\/sfdsk\///"`"
 82:             dev="`echo $dev | /bin/sed "s/\/dsk\//_/"`"
 83:     fi
 84: elif [ "`echo $device | /bin/grep "/dev/vx/dmp/"`" != "" ]
 85: then
 86:     # /dev/vx/dmp/device -> _vx_pv_device
 87:     dev="_vx_pv_`echo $device | /bin/sed "s/\/dev\/vx\/dmp\///"`"
 88: elif [ "`echo $device | /bin/grep "/dev/mapper/"`" != "" ]
 89: then
 90:     # "/dev/mapper/<device>" -> "_mapper_<device>"
 91:     dev="_mapper_`echo $device | /bin/sed "s/\/dev\/mapper\///"`"
 92: elif [ "`echo $device | /bin/grep "/dev/disk/by-id/"`" != "" ]
 93: then
 94:     # "/dev/disk/by-id/<device>" -> "_by-id_<device>"
 95:     dev="_by-id_`echo $device | /bin/sed "s/\/dev\/disk\/by-id\///"`"
 96: elif [ "`echo $device | /bin/grep "/dev/disk/by-path/"`" != "" ]
 97: then
 98:     # "/dev/disk/by-path/<device>" -> "_by-path_<device>"
 99:     dev="_by-path_`echo $device | /bin/sed "s/\/dev\/disk\/by-path\///"`"
100: else
101:     exit 0
102: fi
103: post_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".pre"
104: fstype_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".fstype"
105: bd_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".bd"
106:
107: if [ "$mount_point" != "" ]
```

```
108: then
109:
110: # When device cannot be unmounted
111: #
112: #   if [ "$device" = "/dev/sdX" -o "$device" = "/dev/sdY" ]
113: #   then
114: #              if [ "$bk_device" != "" ]
115: #              then
116: #                      echo $bk_device > $bd_file
117: #              fi
118: #              /bin/awk "\$2==\"$mount_point\" {print \$3}" /proc/mounts > $fstype_file
119: #              /bin/sync
120: #              /bin/sync
121: #              /bin/sync
122: #              echo "fsck" > $post_file
123:
124: # When device can be unmounted
125: #
126: #   else
127:              if [ `/bin/cat /proc/mounts |/bin/cut -d' ' -f 2|/bin/grep "^$mount_point\$"|/usr/bin/
wc -w` != 1 ]; then
128:                      # stack mount (multi device on $mount_point)
129:                      /bin/mount  > $err_log_path/$dev.umount 2>&1
130:                      exit 13
131:              fi
132:              if [ `/bin/cat /proc/mounts |/bin/cut -d' ' -f 1|/bin/grep "^$device\$"|/usr/bin/wc -
w` != 1 ]; then
133:                      cdevice="/dev/`/usr/bin/readlink $device | /bin/sed "s/..\/..\///"`"
134:                 if [ `/bin/cat /proc/mounts |/bin/cut -d' ' -f 1|/bin/grep "^$cdevice\$"|/usr/
bin/wc -w` != 1 ]; then
135:                          # bind mount (device has multi mount point)
136:                          /bin/mount  > $err_log_path/$dev.umount 2>&1
137:                          exit 13
138:                  fi
139:              fi
140:
141:              /bin/awk "\$2==\"$mount_point\" {print \$3}" /proc/mounts > $fstype_file
142:              /bin/umount $mount_point 2>/dev/null
143:              if [ $? != 0 ]
144:              then
145:                      retry_count=3
146:                      sleep_time=1
147:                      result_flag=1
148:
149:                      while [ $retry_count -gt 0 ]
150:                      do
151:                              /bin/umount $mount_point > $err_log_path/$dev.umount 2>&1
152:                              if [ $? != 0 ]
153:                              then
154:                                      retry_count=`expr $retry_count - 1`
155:                                      /bin/sleep $sleep_time
156:                              else
157:                                      /bin/rm -f $err_log_path/$dev.umount
158:                                      result_flag=0
159:                                      break
160:                              fi
161:                      done
162:
163:                      if [ $result_flag != 0 ]
164:                      then
165:                              /sbin/fuser -vum $mount_point> $err_log_path/$dev.fuser 2>&1
166:                              /bin/ps -ef > $err_log_path/$dev.ps 2>&1
167:
```

```
168:                              exit 10
169:                    fi
170:                    fi
171:                    echo "mount" > $post_file
172:
173: #  fi
174:
175: # When device was not mounted
176: #
177: else
178:   echo "none" > $post_file
179: fi
180:
181: exit 0
```

A.2.2 Post-processing of backup

The name of a script file for post-processing of a backup is as follows.

In the case of non-cluster operation

   `/etc/opt/FJSVswsts/sh/OpcBackup.post`

In the case of cluster operation

   `/etc/opt/FJSVswsts/<logic node name>/sh/OpcBackup.post`

A.2.2.1 Post-processing script for backup

```
 1: #!/bin/sh
 2:
 3: # AdvancedCopy Manager
 4: # Copyright FUJITSU LIMITED, 2002-2012
 5:
 6: #
 7: #   Postprocessing of backup processing
 8: #
 9: #          Argument: $1 Device name of transaction disk
10: #                    $2 Mount point of transaction disk
11: #
12: #  Error number
13: #           2: Argument error
14: #          11: mount error
15: #          12: fsck error
16:
17: # Argument check
18: case $# in
19: 1)
20:    ;;
21: 2)
22:    ;;
23: *)
24:    exit 2
25:    ;;
26: esac
27:
28: device="`echo $1`"
29: mount_point="`echo $2`"
30:
31: # Determination of postprocessing file name
```

```
32:
33: if [ "$SWSTGNODE" != "" ]
34: then
35: swstg_node="/`echo $SWSTGNODE`"
36: else
37: swstg_node=""
38: fi
39:
40: err_log_path="/var/opt/FJSVswsts"$swstg_node"/log"
41:
42: if [ "`echo $device | /bin/grep "/dev/sd"`" != "" ]
43: then
44:     # /dev/sd? -> sd?
45:     dev="`echo $device | /bin/sed "s/\/dev\///"`"
46: elif [ "`echo $device | /bin/grep "/dev/vd"`" != "" ]
47: then
48:     # /dev/vd? -> vd?
49:     dev="`echo $device | /bin/sed "s/\/dev\///"`"
50: elif [ "`echo $device | /bin/grep "/dev/FJSV"`" != "" ]
51: then
52:     # /dev/FJSVmphd/dsk/mplb?s? -> mplb?s?
53:     # /dev/FJSVmphd/dsk/mphd?s? -> mphd?s?
54:     dev="`echo $device | /bin/cut -d/ -f5`"
55: elif [ "`echo $device | /bin/grep "/dev/sfdsk/"`" != "" ]
56: then
57:     if [ "`echo $device | /bin/grep ":"`" != ""    ]
58:     then
59:             devnam="`echo $device | /bin/cut -d: -f2-`"
60:             # /dev/sfdsk/class/dsk/volume:sd? -> class_volume_sd?
61:             dev="`echo $device | /bin/sed "s/\/dev\/sfdsk\///"`"
62:             dev="`echo $dev | /bin/sed "s/\/dsk\//_/"`"
63:             dev="`echo $dev | /bin/sed "s/:/_/"`"
64:             device="`echo $device | /bin/cut -d: -f1`"
65:             if [ "`echo $devnam | /bin/grep "/dev/disk/by-id/"`" != "" ]
66:             then
67:                     # /dev/sfdsk/class/dsk/volume:/dev/disk/by-id/<device> ->
class_volume__by_id_<device>
68:                     dev="`echo $dev | /bin/sed "s/\/dev\/disk\/by-id\//_by-id_/"`"
69:             elif [ "`echo $devnam | /bin/grep "/dev/disk/by-path/"`" != "" ]
70:             then
71:                     # /dev/sfdsk/class/dsk/volume:/dev/disk/by-path/<device> ->
class_volume__by_path_<device>
72:                     dev="`echo $dev | /bin/sed "s/\/dev\/disk\/by-path\//_by-path_/"`"
73:             fi
74:     else
75:             # /dev/sfdsk/class/dsk/volume -> _gds_class_volume
76:             dev="_gds_`echo $device | /bin/sed "s/\/dev\/sfdsk\///"`"
77:             dev="`echo $dev | /bin/sed "s/\/dsk\//_/"`"
78:     fi
79: elif [ "`echo $device | /bin/grep "/dev/vx/dmp/"`" != "" ]
80: then
81:     # /dev/vx/dmp/device -> _vx_pv_device
82:     dev="_vx_pv_`echo $device | /bin/sed "s/\/dev\/vx\/dmp\///"`"
83: elif [ "`echo $device | /bin/grep "/dev/mapper/"`" != "" ]
84: then
85:     # "/dev/mapper/<device>" -> "_mapper_<device>"
86:     dev="_mapper_`echo $device | /bin/sed "s/\/dev\/mapper\///"`"
87: elif [ "`echo $device | /bin/grep "/dev/disk/by-id/"`" != "" ]
88: then
89:     # "/dev/disk/by-id/<device>" -> "_by-id_<device>"
90:     dev="_by-id_`echo $device | /bin/sed "s/\/dev\/disk\/by-id\///"`"
91: elif [ "`echo $device | /bin/grep "/dev/disk/by-path/"`" != "" ]
92: then
```

```
 93:     # "/dev/disk/by-path/<device>" -> "_by-path_<device>"
 94:     dev="_by-path_`echo $device | /bin/sed "s/\/dev\/disk\/by-path\///"`"
 95: else
 96:     exit 0
 97: fi
 98: post_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".pre"
 99: fstype_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".fstype"
100: bd_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".bd"
101:
102: # Confirmation of postprocessing
103: if [ ! -r $post_file ]
104: then
105:     exit 0
106: fi
107: post="`/bin/cat $post_file`"
108:
109: # Confirmation of FStype
110: if [ ! -r $fstype_file ]
111: then
112:     fs=""
113: else
114:     fs="`/bin/cat $fstype_file`"
115: fi
116:
117: # No processing
118: if [ "$post" = "none" ]
119: then
120:     /bin/rm -rf $post_file 2> /dev/null
121:     /bin/rm -rf $fstype_file 2> /dev/null
122:     exit 0
123: fi
124:
125: # mount processing
126: if [ "$post" = "mount" ]
127: then
128:     if [ "`echo $device | /bin/grep "/dev/disk/by-id/"`" != "" \
129:             -o "`echo $device | /bin/grep "/dev/disk/by-path/"`" != "" ]
130:     then
131:             cdevice="/dev/`/usr/bin/readlink $device | /bin/sed "s/..\/..\///"`"
132:             Result="`/bin/df -l | /bin/grep "$cdevice " | /bin/awk 'END {print NR}'`"
133:     else
134:             Result="`/bin/df -l | /bin/grep "$device " | /bin/awk 'END {print NR}'`"
135:     fi
136:       if [ "$Result" != "1" ]
137:     then
138:             if [ ! -r $fstype_file ]
139:             then
140:                     /bin/mount $device $mount_point 2> /dev/null
141:             else
142:                     Result1="`echo $fs | /bin/awk 'END {print NR}'`"
143:                         if [ "$Result1" != "1" ]
144:                     then
145:                             /bin/mount $device $mount_point 2> /dev/null
146:                     else
147:                             /bin/mount -t $fs $device $mount_point 2> /dev/null
148:                     fi
149:             fi
150:             if [ $? != 0 ]
151:             then
152:                     retry_count=3
153:                     sleep_time=1
154:                     result_flag=1
155:
```

```
156:                    while [ $retry_count -gt 0 ]
157:                    do
158:                        if [ ! -r $fstype_file ]
159:                        then
160:                            /bin/mount $device $mount_point > $err_log_path/$dev.mount 2>&1
161:                        else
162:                            Result1="`echo $fs | /bin/awk 'END {print NR}'`"
163:                            if [ "$Result1" != "1" ]
164:                            then
165:                                /bin/mount $device $mount_point > $err_log_path/
$dev.mount 2>&1
166:                            else
167:                                /bin/mount -t $fs $device $mount_point > $err_log_path/
$dev.mount 2>&1
168:                            fi
169:                        fi
170:                        if [ $? != 0 ]
171:                        then
172:                            retry_count=`expr $retry_count - 1`
173:                            /bin/sleep $sleep_time
174:                        else
175:                            /bin/rm -f $err_log_path/$dev.mount
176:                            result_flag=0
177:                            break
178:                        fi
179:                    done
180:
181:                    if [ $result_flag != 0 ]
182:                    then
183:                        exit 11
184:                    fi
185:            fi
186:    fi
187:    /bin/rm -rf $post_file 2> /dev/null
188:    /bin/rm -rf $fstype_file 2> /dev/null
189:    exit 0
190: fi
191:
192: # fsck processing
193: if [ "$post" = "fsck" ]
194: then
195:    if [ -r $bd_file ]
196:    then
197:        bk_device="`/bin/cat $bd_file`"
198:        fsck_dev="`echo $bk_device `"
199:        if [ ! -r $fstype_file ]
200:        then
201:            /sbin/fsck -c $fsck_dev > /dev/null 2>&1
202:        else
203:            if [ "$fs" = "" ]
204:            then
205:                /sbin/fsck -c $fsck_dev > /dev/null 2>&1
206:            else
207:                if [ "$fs" = "sfxfs" ]
208:                then
209:                    /sbin/sfxadm $fsck_dev > /dev/null 2>&1
210:                    /sbin/fsck -p -t $fs $fsck_dev > /dev/null 2>&1
211:                else
212:                    /sbin/fsck -p -t $fs $fsck_dev > /dev/null 2>&1
213:                fi
214:            fi
215:        fi
216:        if [ $? != 0 ]
```

```
217:                then
218:                        if [ "$fs" = "" ]
219:                        then
220:                                result="`/sbin/fsck -p $fsck_dev `"
221:                        else
222:                                result="`/sbin/fsck -p -t $fs $fsck_dev `"
223:                        fi
224:                        fsck_rc=$?
225:                        if [ $fsck_rc != 0 ] && [ $fsck_rc != 1 ]
226:                        then
227:                                echo "$result" > $err_log_path/$dev.fsck
228:                                exit 12
229:                        fi
230:                fi
231:        fi
232:        /bin/rm -rf $post_file 2> /dev/null
233:        /bin/rm -rf $fstype_file 2> /dev/null
234:        /bin/rm -rf $bd_file 2> /dev/null
235:        exit 0
236: fi
237:
238: exit 0
```

A.2.3 Customizing the shell script

The method of customizing the shell script for backup pre- and post-processing is explained below.

When you do not want to unmount a transaction volume

Edit the script to delete the comment characters ("#") on lines 112-122, 126, and 173 for a transaction volume that you do not want to unmount, and specify the target device in the if statement on line 112.
Make sure that any special editing of this script has not changed the line numbers quoted above or label the lines to be modified in these circumstances clearly.
Do not update the file system of the transaction volume in the period from execution of the sync command on line 119 to backup post-processing. Otherwise, the backed-up file system may be incomplete or execution of the fsck command in the backup post-processing may result in an error.

This operation is possible only when the file system is one of the following two:

- ext2

- ext3

The pre-processing by the customized scripts is as follows:

Table A.4 Pre-processing by the customized script

State of transaction volume	preprocessing
It is mounted.	All of data buffered in memory are written to disk.

The post-processing by the customized scripts is as follows:

Table A.5 Post-processing by the customized script

preprocessing	post-processing
If all of data buffered in memory was written to disk.	Matching of the backup volume is checked.

Note

. . . .

If a SDX object is operated in slice units, do not specify slice name (/dev/sfdsk/<*class name*>/dsk/<*volume name*>:<*sdxinfo DEVNAM value*>) in line 112. Instead, specify logical volume name (/dev/sfdsk/<*class name*>/dsk/<*volume name*>).

. . . .

A.3   Pre-processing and Post-processing of Restoration

In AdvancedCopy Manager, restore processing must be performed while transaction volumes are unmounted.

Thus, in pre-processing, the mount state of the transaction volume is determined using its name and the following processing is then performed:

Table A.6 Pre-processing for a transaction volume

Transaction volume status	Preprocessing
Mounted	Unmount the transaction volume.
Unmounted	Take no action.

Post-processing tasks will be determined by what has occurred in pre-processing

Table A.7 Pre-processing and post-processing for restoration

Preprocessing	Post-processing
A transaction volume was unmounted.	Remount the transaction volume.
Take no action.	Take no action.

If special pre-processing or post-processing is required, you need to add the instructions as appropriate to the shell script.

When customizing a script, strictly observe the following rules regarding error codes:

Table A.8 Error code rules

Error code	Usage
0-99	Unusable (reserved for AdvancedCopy Manager)
100-255	Usable

Note

. . . .

If the post-processing fails, the resource information may have inconsistencies. Execute swstsrsemtch (Resource match command).

. . . .

A.3.1   Pre-processing of restoration

The name of a script file for pre-processing of a restoration is as follows.

In the case of non-cluster operation

```
/etc/opt/FJSVswsts/sh/OpcRestore.pre
```

In the case of cluster operation

```
/etc/opt/FJSVswsts/<logic node name>/sh/OpcRestore.pre
```

A.3.1.1 Pre-processing script for restoration

```
 1: #!/bin/sh
 2:
 3: # AdvancedCopy Manager
 4: # Copyright FUJITSU LIMITED, 2002-2012
 5:
 6: #
 7: #   Preprocessing of restoration processing
 8: #
 9: #            Argument: $1 Device name of transaction disk
10: #                      $2 Mount point of transaction disk
11: #
12: #  Error number
13: #             2: Argument error
14: #            10: umount error
15: #            13: Illegal mount type (bind/stack mount)
16:
17:
18: # Argument check
19: case $# in
20: 1)
21:    ;;
22: 2)
23:    ;;
24: *)
25:    exit 2
26:    ;;
27: esac
28:
29: device="`echo $1`"
30: mount_point="`echo $2`"
31:
32: # Determination of postprocessing file name
33:
34: if [ "$SWSTGNODE" != "" ]
35: then
36: swstg_node="/`echo $SWSTGNODE`"
37: else
38: swstg_node=""
39: fi
40:
41: err_log_path="/var/opt/FJSVswsts"$swstg_node"/log"
42:
43: if [ "`echo $device | /bin/grep "/dev/sd"`" != "" ]
44: then
45:    # /dev/sd? -> sd?
46:    dev="`echo $device | /bin/sed "s/\/dev\///"`"
47: elif [ "`echo $device | /bin/grep "/dev/vd"`" != "" ]
48: then
49:    # /dev/vd? -> vd?
50:    dev="`echo $device | /bin/sed "s/\/dev\///"`"
51: elif [ "`echo $device | /bin/grep "/dev/FJSV"`" != "" ]
52: then
53:    # /dev/FJSVmphd/dsk/mplb?s? -> mplb?s?
54:    # /dev/FJSVmphd/dsk/mphd?s? -> mphd?s?
55:    dev="`echo $device | /bin/cut -d/ -f5`"
56: elif [ "`echo $device | /bin/grep "/dev/sfdsk/"`" != "" ]
57: then
58:    if [ "`echo $device | /bin/grep ":"`" != ""   ]
59:    then
```

```
 60:            devnam="`echo $device | /bin/cut -d: -f2-`"
 61:            # /dev/sfdsk/class/dsk/volume:sd? -> class_volume_sd?
 62:            dev="`echo $device | /bin/sed "s/\/dev\/sfdsk\///"`"
 63:            dev="`echo $dev | /bin/sed "s/\/dsk\//_/"`"
 64:            dev="`echo $dev | /bin/sed "s/:/_/"`"
 65:            device="`echo $device | /bin/cut -d: -f1`"
 66:            if [ "`echo $devnam | /bin/grep "/dev/disk/by-id/"`" != "" ]
 67:            then
 68:                    # /dev/sfdsk/class/dsk/volume:/dev/disk/by-id/<device> ->
class_volume__by_id_<device>
 69:                    dev="`echo $dev | /bin/sed "s/\/dev\/disk\/by-id\//_by-id_/"`"
 70:            elif [ "`echo $devnam | /bin/grep "/dev/disk/by-path/"`" != "" ]
 71:            then
 72:                    # /dev/sfdsk/class/dsk/volume:/dev/disk/by-path/<device> ->
class_volume__by_path_<device>
 73:                    dev="`echo $dev | /bin/sed "s/\/dev\/disk\/by-path\//_by-path_/"`"
 74:            fi
 75:    else
 76:            # /dev/sfdsk/class/dsk/volume -> _gds_class_volume
 77:            dev="_gds_`echo $device | /bin/sed "s/\/dev\/sfdsk\///"`"
 78:            dev="`echo $dev | /bin/sed "s/\/dsk\//_/"`"
 79:    fi
 80: elif [ "`echo $device | /bin/grep "/dev/vx/dmp/"`" != "" ]
 81: then
 82:    # /dev/vx/dmp/device -> _vx_pv_device
 83:    dev="_vx_pv_`echo $device | /bin/sed "s/\/dev\/vx\/dmp\///"`"
 84: elif [ "`echo $device | /bin/grep "/dev/mapper/"`" != "" ]
 85: then
 86:    # "/dev/mapper/<device>" -> "_mapper_<device>"
 87:    dev="_mapper_`echo $device | /bin/sed "s/\/dev\/mapper\///"`"
 88: elif [ "`echo $device | /bin/grep "/dev/disk/by-id/"`" != "" ]
 89: then
 90:    # "/dev/disk/by-id/<device>" -> "_by-id_<device>"
 91:    dev="_by-id_`echo $device | /bin/sed "s/\/dev\/disk\/by-id\///"`"
 92: elif [ "`echo $device | /bin/grep "/dev/disk/by-path/"`" != "" ]
 93: then
 94:    # "/dev/disk/by-path/<device>" -> "_by-path_<device>"
 95:    dev="_by-path_`echo $device | /bin/sed "s/\/dev\/disk\/by-path\///"`"
 96: else
 97:    exit 0
 98: fi
 99: post_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".pre"
100: fstype_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".fstype"
101:
102: if [ "$mount_point" != "" ]
103: then
104:         if [ `/bin/cat /proc/mounts |/bin/cut -d' ' -f 2|/bin/grep "^$mount_point\$"|/usr/bin/wc
-w` != 1 ]; then
105:                    # stack mount (multi device on $mount_point)
106:                    /bin/mount  > $err_log_path/$dev.umount 2>&1
107:                exit 13
108:         fi
109:         if [ `/bin/cat /proc/mounts |/bin/cut -d' ' -f 1|/bin/grep "^$device\$"|/usr/bin/wc -w` !
= 1 ]; then
110:                    cdevice="/dev/`/usr/bin/readlink $device | /bin/sed "s/..\/..\///"`"
111:                    if [ `/bin/cat /proc/mounts |/bin/cut -d' ' -f 1|/bin/grep "^$cdevice\$"|/usr/
bin/wc -w` != 1 ]; then
112:                            # bind mount (device has multi mount point)
113:                            /bin/mount  > $err_log_path/$dev.umount 2>&1
114:                            exit 13
115:                    fi
116:         fi
117:
```

```
118:
119: # When device can be unmounted
120: #
121: #    df -ln $mount_point | cut -f2 -d: | cut -f2 -d' ' > $fstype_file
122:     /bin/awk "\$2==\"$mount_point\" {print \$3}" /proc/mounts > $fstype_file
123:     /bin/umount $mount_point 2> /dev/null
124:     if [ $? != 0 ]
125:     then
126:             retry_count=3
127:             sleep_time=1
128:             result_flag=1
129:
130:             while [ $retry_count -gt 0 ]
131:             do
132:                     /bin/umount $mount_point > $err_log_path/$dev.umount 2>&1
133:                     if [ $? != 0 ]
134:                     then
135:                             retry_count=`expr $retry_count - 1`
136:                             /bin/sleep $sleep_time
137:                     else
138:                             /bin/rm -f $err_log_path/$dev.umount
139:                             result_flag=0
140:                             break
141:                     fi
142:             done
143:
144:             if [ $result_flag != 0 ]
145:             then
146:                     /sbin/fuser -vum $mount_point> $err_log_path/$dev.fuser 2>&1
147:                     /bin/ps -ef > $err_log_path/$dev.ps 2>&1
148:
149:                     exit 10
150:             fi
151:     fi
152:     echo "mount" > $post_file
153:
154: # When device was not mounted
155: #
156: else
157:     echo "none" > $post_file
158: fi
159:
160: exit 0
```

Restore processing is not possible on a mounted transaction volume that cannot be unmounted. Specify a device at the restore destination.

A.3.2 Post-processing of restoration

The name of a script file for post-processing of a restoration is as follows.

In the case of non-cluster operation

```
/etc/opt/FJSVswsts/sh/OpcRestore.post
```

In the case of cluster operation

```
/etc/opt/FJSVswsts/<logic node name>/sh/OpcRestore.post
```

A.3.2.1 Post-processing script for restoration

```
 1: #!/bin/sh
 2:
 3: # AdvancedCopy Manager
 4: # Copyright FUJITSU LIMITED, 2002-2012
 5: #
 6: #    Postprocessing of restoration processing
 7: #
 8: #            Argument: $1 Device name of transaction disk
 9: #                      $2 Mount point of transaction disk
10: #
11: #  Error number
12: #            2: Argument error
13: #           11: mount error
14:
15: # Argument check
16: case $# in
17: 1)
18:    ;;
19: 2)
20:    ;;
21: *)
22:    exit 2
23:    ;;
24: esac
25:
26: device="`echo $1`"
27: mount_point="`echo $2`"
28:
29: # Determination of postprocessing file name
30:
31: if [ "$SWSTGNODE" != "" ]
32: then
33: swstg_node="/`echo $SWSTGNODE`"
34: else
35: swstg_node=""
36: fi
37:
38: err_log_path="/var/opt/FJSVswsts"$swstg_node"/log"
39:
40: if [ "`echo $device | /bin/grep "/dev/sd"`" != "" ]
41: then
42:    # /dev/sd? -> sd?
43:    dev="`echo $device | /bin/sed "s/\/dev\///"`"
44: elif [ "`echo $device | /bin/grep "/dev/vd"`" != "" ]
45: then
46:    # /dev/vd? -> vd?
47:    dev="`echo $device | /bin/sed "s/\/dev\///"`"
48: elif [ "`echo $device | /bin/grep "/dev/FJSV"`" != "" ]
49: then
50:    # /dev/FJSVmphd/dsk/mplb?s? -> mplb?s?
51:    # /dev/FJSVmphd/dsk/mphd?s? -> mphd?s?
52:    dev="`echo $device | /bin/cut -d/ -f5`"
53: elif [ "`echo $device | /bin/grep "/dev/sfdsk/"`" != "" ]
54: then
55:    if [ "`echo $device | /bin/grep ":"`" != ""   ]
56:    then
57:            devnam="`echo $device | /bin/cut -d: -f2-`"
58:            # /dev/sfdsk/class/dsk/volume:sd? -> class_volume_sd?
59:            dev="`echo $device | /bin/sed "s/\/dev\/sfdsk\///"`"
60:            dev="`echo $dev | /bin/sed "s/\/dsk\//_/"`"
61:            dev="`echo $dev | /bin/sed "s/:/_/"`"
```

```
 62:            device="`echo $device | /bin/cut -d: -f1`"
 63:            if [ "`echo $devnam | /bin/grep "/dev/disk/by-id/"`" != "" ]
 64:            then
 65:                    # /dev/sfdsk/class/dsk/volume:/dev/disk/by-id/<device> ->
class_volume__by_id_<device>
 66:                    dev="`echo $dev | /bin/sed "s/\/dev\/disk\/by-id\//_by-id_/"`"
 67:            elif [ "`echo $devnam | /bin/grep "/dev/disk/by-path/"`" != "" ]
 68:            then
 69:                    # /dev/sfdsk/class/dsk/volume:/dev/disk/by-path/<device> ->
class_volume__by_path_<device>
 70:                    dev="`echo $dev | /bin/sed "s/\/dev\/disk\/by-path\//_by-path_/"`"
 71:            fi
 72:    else
 73:            # /dev/sfdsk/class/dsk/volume -> _gds_class_volume
 74:            dev="_gds_`echo $device | /bin/sed "s/\/dev\/sfdsk\///"`"
 75:            dev="`echo $dev | /bin/sed "s/\/dsk\//_/"`"
 76:    fi
 77: elif [ "`echo $device | /bin/grep "/dev/vx/dmp/"`" != "" ]
 78: then
 79:    # /dev/vx/dmp/device -> _vx_pv_device
 80:    dev="_vx_pv_`echo $device | /bin/sed "s/\/dev\/vx\/dmp\///"`"
 81: elif [ "`echo $device | /bin/grep "/dev/mapper/"`" != "" ]
 82: then
 83:    # "/dev/mapper/<device>" -> "_mapper_<device>"
 84:    dev="_mapper_`echo $device | /bin/sed "s/\/dev\/mapper\///"`"
 85: elif [ "`echo $device | /bin/grep "/dev/disk/by-id/"`" != "" ]
 86: then
 87:    # "/dev/disk/by-id/<device>" -> "_by-id_<device>"
 88:    dev="_by-id_`echo $device | /bin/sed "s/\/dev\/disk\/by-id\///"`"
 89: elif [ "`echo $device | /bin/grep "/dev/disk/by-path/"`" != "" ]
 90: then
 91:    # "/dev/disk/by-path/<device>" -> "_by-path_<device>"
 92:    dev="_by-path_`echo $device | /bin/sed "s/\/dev\/disk\/by-path\///"`"
 93: else
 94:    exit 0
 95: fi
 96: post_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".pre"
 97: fstype_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".fstype"
 98:
 99: # Confirmation of postprocessing
100: if [ ! -r $post_file ]
101: then
102:    exit 0
103: fi
104: post="`/bin/cat $post_file`"
105:
106: # Confirmation of FStype
107: if [ ! -r $fstype_file ]
108: then
109:    fs=""
110: else
111:    fs="`/bin/cat $fstype_file`"
112: fi
113:
114: # No processing
115: if [ "$post" = "none" ]
116: then
117:    /bin/rm -rf $post_file 2> /dev/null
118:    /bin/rm -rf $fstype_file 2> /dev/null
119:    exit 0
120: fi
121:
122: # mount processing
```

```
123: if [ "$post" = "mount" ]
124: then
125:     if [ "`echo $device | /bin/grep "/dev/disk/by-id/"`" != "" \
126:             -o "`echo $device | /bin/grep "/dev/disk/by-path/"`" != "" ]
127:     then
128:             cdevice="/dev/`/usr/bin/readlink $device | /bin/sed "s/..\/..\///"`"
129:             Result="`/bin/df -l | /bin/grep "$cdevice " | /bin/awk 'END {print NR}'`"
130:     else
131:             Result="`/bin/df -l | /bin/grep "$device " | /bin/awk 'END {print NR}'`"
132:     fi
133:        if [ "$Result" != "1" ]
134:     then
135:             if [ ! -r $fstype_file ]
136:             then
137:                     /bin/mount $device $mount_point 2> /dev/null
138:             else
139:                     Result1="`echo $fs | /bin/awk 'END {print NR}'`"
140:                         if [ "$Result1" != "1" ]
141:                     then
142:                             /bin/mount $device $mount_point 2> /dev/null
143:                     else
144:                             /bin/mount -t $fs $device $mount_point 2> /dev/null
145:                     fi
146:             fi
147:             if [ $? != 0 ]
148:             then
149:                     retry_count=3
150:                     sleep_time=1
151:                     result_flag=1
152:
153:                     while [ $retry_count -gt 0 ]
154:                     do
155:                             if [ ! -r $fstype_file ]
156:                             then
157:                                     /bin/mount $device $mount_point > $err_log_path/$dev.mount 2>&1
158:                             else
159:                                     Result1="`echo $fs | /bin/awk 'END {print NR}'`"
160:                                     if [ "$Result1" != "1" ]
161:                                     then
162:                                             /bin/mount $device $mount_point > $err_log_path/
$dev.mount 2>&1
163:                                     else
164:                                             /bin/mount -t $fs $device $mount_point > $err_log_path/
$dev.mount 2>&1
165:                                     fi
166:                             fi
167:                             if [ $? != 0 ]
168:                             then
169:                                     retry_count=`expr $retry_count - 1`
170:                                     /bin/sleep $sleep_time
171:                             else
172:                                     /bin/rm -f $err_log_path/$dev.mount
173:                                     result_flag=0
174:                                     break
175:                             fi
176:                     done
177:
178:                     if [ $result_flag != 0 ]
179:                     then
180:                             exit 11
181:                     fi
182:             fi
183:     fi
```

```
184:     /bin/rm -rf $post_file 2> /dev/null
185:     /bin/rm -rf $fstype_file 2> /dev/null
186:     exit 0
187: fi
188:
189: exit 0
```

Appendix B Sample Script for Backup Operations

This appendix provides information on sample scripts for AdvancedCopy Manager.
A sample script allows you to back up data on a transaction volume with AdvancedCopy Manager.

P Point

This section explains how the sample script is created for the backup management function. By using it as a reference when creating similar scripts for the replication management function, you can copy Oracle databases.

B.1 Overview

Use AdvancedCopy Manager to back up a transaction volume.
AdvancedCopy Manager provides a sample scripts for backup. The manager of the backup operation should customize the sample scripts according to their specific requirements and apply them to the backup.

B.2 Backup and Restoration of a Transaction volume

B.2.1 Backup procedure

The following is an overview of the backup procedure:

1. Check the work operation status
   Check whether processing is in operation or not.

2. Declare the start of online backup
   If processing was found to be in operation in Step 1, declare the start of online backup for the database to stop updating the data files. This will make the partitions that store transaction data ready for online backup.

3. Backup (copy) partitions
   Backup (copy) the partitions that store transaction data using the AdvancedCopy Manager backup function.

4. Declare the end of online backup
   If processing was found to be in operation in Step 1, declare the end of online backup for the database to resume updating the data files.

Figure B.1 Backup
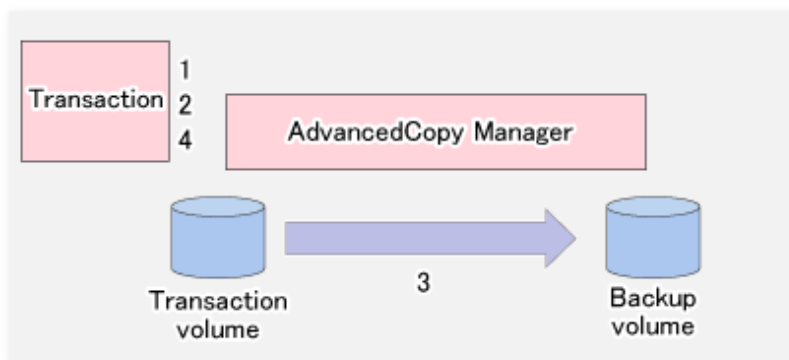
B.2.2 Restoration procedure

The following provides an overview of how to perform restoration from a backup volume.

1. Checking the processing stop status
   Check that processing has stopped. If processing is in operation, stop the processing completely.

2. Performing restoration from a backup volume
   Using the restoration function of AdvancedCopy Manager, restore data from a backup volume to the transaction volume.

Figure B.2 Performing restoration from a backup volume

B.3 Sample Scripts for Oracle

B.3.1 Overview of sample scripts

The sample shell scripts provided are summarized below. These scripts are stored in the /etc/opt/FJSVswsts/samp directory.

Table B.1 Sample shell script for Oracle

No.	Type	Script name	Associated transaction	Associated resource type
1	Collection of DB information	swst_or_iv	Oracle transaction	-
2	Backup	swst_or_fs_bk	Oracle transaction, general file transaction	File system
3	Backup	swst_or_rd_bk	Oracle transaction	Raw device

B.3.1.1 Collecting DB information

swst_or_iv outputs information about an Oracle database. The output file names and their contents are:

- spdata.lst : List of table space names and storage destinations

- spdata.srt : Results of spdata.lst, with storage destination as the sort key

- spdata.bgn : Command in the following format:

  ```
  Alter tablespace  Table-space-name  begin backup;
  ```

- spdata.end : Command in the following format:

```
 Alter tablespace  Table-space-name  end backup;
```

- spcont.lst : Oracle database control file information

- splog.lst : Oracle database log file information

Use the database information above as a reference for creating shell scripts for backup processing.

Format

```
swst_or_iv
```

Usage example

```
# swst_or_iv
swst_or_iv completed.
```

Result example

```
spdata.lst
TABLESPACE-NAME              DEVICE-NAME
------------------- --------------------------------------------------
ORA10                /ora1/ora0.dbf
ORA20                /ora2/ora0.dbf
RBS                  /oracle/ora/oradata/rbs01.dbf
SYSTEM               /oracle/ora/oradata/system01.dbf
TEMP                 /oracle/ora/oradata/temp01.dbf
TOOLS                /oracle/ora/oradata/tools01.dbf
USERS                /oracle/ora/oradata/users01.dbf
spdata.bgn
alter tablespace ORA10 begin backup;
alter tablespace ORA20 begin backup;
alter tablespace RBS begin backup;
alter tablespace SYSTEM begin backup;
alter tablespace TEMP begin backup;
alter tablespace TOOLS begin backup;
alter tablespace USERS begin backup;
spdata.end
alter tablespace ORA10 end backup;
alter tablespace ORA20 end backup;
alter tablespace RBS end backup;
alter tablespace SYSTEM end backup;
alter tablespace TEMP end backup;
alter tablespace TOOLS end backup;
alter tablespace USERS end backup;
spcont.lst
NAME                                TYPE    VALUE
--------------------------------- ------- ------------------------------
control_files                      string  /oracle/ora/oradata/control
splog.lst
NAME                                TYPE    VALUE
--------------------------------- ------- ------------------------------
log_archive_dest                   string  /oracle/ora/oradata/archive
```

Processing outline

```
#!/bin/sh
reterr() --- An error message is output before an error is returned.
sqldata()
{
sqlplus AAA/BBB <<!
SELECT --- Acquire Oracle table space information.
}
```

```
sqllog()
{
sqlplus AAA/BBB <<!
show --- Acquire Oracle log information.
}
sqlcont()
{
sqlplus AAA/BBB <<!
show --- Acquire Oracle control information.
}

# main()

# 0) Environmental variable setting
PATH=$PATH:/usr/bin:/usr/sbin:/opt/FJSVswsts/bin

# 1) Acquisition of Oracle table space information
sqldata() invocation --- Acquire Oracle table space information

# 2) Acquisition of Oracle log information
sqllog() invocation --- Acquire Oracle log information

# 3) Acquisition of Oracle control information
sqlcont() invocation --- Acquire Oracle control information

exit 0
```

B.3.1.2 File system backup

swst_or_fs_bk, run using AdvancedCopy Manager, backs up processing data stored in an Oracle database on a file system.
Processing data is copied to the backup volumes of AdvancedCopy Manager. These scripts correspond to steps 2 to 5 in "B.3.2.1 DB
information collection".

Format

```
swst_or_fs_bk {-a | -i} <mount-point-name>
```

Options

-a

Specify when an Oracle database is online. The start of backup processing is declared to the Oracle server.

-i

Specify when an Oracle database is stopped or the file is a general file. The start of backup processing is not declared to the Oracle
server.

Operands

mount-point-name

Specify the name of the mount point targeted for backup.

Example using swst_or_fs_bk

While processing is active, back up /gyoumu_db in which an Oracle database is located.

```
# swst_or_fs_bk -a /gyoumu_db
swst_or_fs_bk completed. ( /gyoumu_db -> /dev/sdb2 at 23:00 on 1999.11.01 )
```

Processing outline of swst_or_fs_bk

```
#!/bin/sh
usage() --- Command syntax is output before an error is returned.
reterr() --- An error message is output before an error is returned.
chkmp() --- Check for a mount point
get_mp_blk() --- Conversion from a mount point to a device
getfs() --- Conversion from a mount point to a file system type
sqlbegin()
{
sqlplus AAA/BBB <<!
alter system switch logfile;
alter system flush shared_pool;
alter tablespace CCC begin backup; --- Notify Oracle of the start of backup.
}
sqlend()
{
sqlplus AAA/BBB <<!
alter tablespace CCC end backup; --- Notify Oracle of the stopping of backup.
}
getbd() --- Fetch the name of the backup (copy) destination device of AdvancedCopy Manager.
getdate() --- Fetch the backup (copy) date of AdvancedCopy Manager.
gettime() --- Fetch the backup (copy) time of AdvancedCopy Manager.

# main()
Parameter analysis and mount point check

# 0) Environmental variable setting
PATH=$PATH:/usr/bin:/usr/sbin:/opt/FJSVswsts/bin

# 1) Notification issued to Oracle of the start of backup
if [ "-a" operand specification ]
then
sqlbegin() invocation --- Notify Oracle of the start of backup.
Error handling
fi

# 2) Disk backup (copy) by AdvancedCopy Manager command
swstbackup $UDBLK --- Copy the disk.
Error handling

# 3) Notification issued to Oracle of the stopping of backup
if [ "-a" operand specification ]
then
sqlend() invocation --- Notify Oracle of the stopping of backup.
Error handling
fi

# 4) Acquisition of the backup (copy) destination of AdvancedCopy Manager
BDBLK="`getbd $UDBLK`" --- Fetch a backup volume as a character device.
BDDATE="`getdate $UDBLK`"--- Fetch the backup date.
BDTIME="`gettime $UDBLK`"--- Fetch the backup time.

exit 0
```

B.3.1.3 Raw device backup

swst_or_rd_bk is run using AdvancedCopy Manager to back up live data using an Oracle database on a raw device.
Live data is copied to work disks prepared by users.
These scripts apply to steps 2 to 5 in "B.3.2.1 DB information collection".

Format

```
swst_or_rd_bk {-a | -i} <raw device name> <block device name>
```

Options

  -a

  Specify when an Oracle database is online. The start of backup processing is declared to the Oracle server.

  -i

  Specify when an Oracle database is stopped or the file is a general file. The start of backup processing is not declared to the Oracle server.

Operands

  raw device name

  Specify the name of the raw device targeted for backup.

  block device name

  Specify the name of the block device targeted for backup.

Example using swst_or_rd_bk

  While business is active, back up /dev/sdb0 in which an Oracle database is located.

```
# swst_or_rd_bk -a /dev/raw/raw1 /dev/sdb6
swst_or_rd_bk completed. ( /dev/sdb6 -> /dev/sdb2 at 23:01 on 1999.11.01 )
```

Processing outline of swst_or_rd_bk

```
#!/bin/sh

usage() --- Command syntax is output before an error is returned.
reterr() --- An error message is output before an error is returned.
sqlbegin()
{
sqlplus AAA/BBB <<!
alter system switch logfile;
alter system flush shared_pool;
alter tablespace CCC begin backup; --- Notify Oracle of the start of backup.
}
sqlend()
{
sqlplus AAA/BBB <<!
alter tablespace CCC end backup; --- Notify Oracle of the stopping of backup.
}
getbd() --- Fetch the name of the backup (copy) destination device of AdvancedCopy Manager.
getdate() --- Fetch the backup (copy) date of AdvancedCopy Manager.
gettime() --- Fetch the backup (copy) time of AdvancedCopy Manager.

# main()
Parameter analysis and raw device check

# 0) Environmental variable setting
PATH=$PATH:/usr/bin:/usr/sbin:/opt/FJSVswsts/bin

# 1) Notification issued to Oracle of the start of backup
if [ "-a" operand specification ]
then
sqlbegin() invocation --- Notify Oracle of the start of backup.
Error handling
fi

# 2) Disk backup (copy) by AdvancedCopy Manager command
```

```
swstbackup $SRC --- Copy the disk.
Error handling

# 3) Notification issued to Oracle of the stopping of backup
if [ "-a" operand specification ]
then
sqlend() invocation --- Notify Oracle of the stopping of backup.
Error handling
fi

# 4) Acquisition of the backup (copy) destination of AdvancedCopy Manager
DSTBLK="`getbd $SRC`" --- Fetch a backup volume.
BKDATE="`getdate $SRC`"
BKTIME="`gettime $SRC`"

exit 0
```

B.3.2 Applying the sample shell scripts

B.3.2.1 DB information collection

Sample shell script update

Update the procedures below of swst_or_iv according to actual operations. The following numbers correspond to those in the scripts:

- 0) Environmental variable setting

- 1) Acquisition of Oracle table space information

- 2) Acquisition of Oracle log information

- 3) Acquisition of Oracle control information

Procedures 1) to 3) change Oracle account information.
These scripts can operate only on business servers of AdvancedCopy Manager.

B.3.2.2 File system backup

AdvancedCopy Manager setup

Set up backup volumes so that the mount point targeted for backup is backed up correctly in AdvancedCopy Manager.
The generations of the backup volumes to be managed are arbitrary.

Sample shell script update

Update the procedures below of swst_or_fs_bk to adjust them for actual operations. The following numbers correspond to those in the scripts.

- 0) Environmental variable setting

- 1) Notification issued to Oracle of the start of backup

- 3) Notification issued to Oracle of the stopping of backup

Procedures 1) and 3) refer to the spdata.bgn file and spdata.end file created by swst_or_iv.
These scripts can operate only on business servers of AdvancedCopy Manager.

B.3.2.3 Raw device backup

Set up backup volumes so that the raw device targeted for backup is backed up correctly in AdvancedCopy Manager.
The generations of the backup volumes to be managed are arbitrary.

Sample shell script update

Update the procedures below of swst_or_rd_bk to adjust them for actual operations. The following numbers correspond to those in the scripts.

- 0) Environmental variable setting

- 1) Notification issued to Oracle of the start of backup

- 3) Notification issued to Oracle of the stopping of backup

Procedures 1) and 3) refer to spdata.lst, spdata.srt, spdata.bgn, spdata.end, splog.lst and spcont.lst, created by swst_or_iv.
These scripts can operate only on business servers of AdvancedCopy Manager.

Recording and managing the backup volume name and processing time

After executing a shell script, a backup volume name and the processing time are output in a message. This information is needed for restoration, so it must be recorded and managed.

B.3.3 Restore Method

B.3.3.1 File system

Use the restore function of AdvancedCopy Manager for restore processing.

B.3.3.2 Raw device

Use the restore function of AdvancedCopy Manager for restore processing.

B.3.4 Notes about using the sample shell scripts

Oracle resource location

Store Oracle data files, and files other than data files, on separate disks or in separate partitions.

Backup of Oracle resources

Note the following when performing this function:

- Initialization parameter file

  - Every time a change occurs, back up this file using the cp command of the OS is executed for backup.

- Control files

  - You are recommended to use mirroring using Oracle.

  - If a change is made in the DB environment, use the following command to create a trace file. If a control file is destroyed, execute the SQL statement in the trace file to recover the control file.

    ```
    alter database backup control file to trace;
    ```

- Data files

    - The recommended location of data files is on a raw device.

    - To store a data file on the file system, disable the automatic extended function of the data file (AUTOEXTEND OFF). Customization of the sample shell script is required to prevent unmounting during context processing.

    - Create backup copies in AdvancedCopy Manager.

    - The range of table spaces targeted for backup must be within the unit of synchronous backup processing (begin backup/end backup) in Oracle and the unit of backup in AdvancedCopy Manager. For example, if backup copies must be created in units of LUNs in AdvancedCopy Manager, place one table space in one LUN.

    - Archive all of the latest logs by executing the following command after backing up all table spaces in the system:

        ```
        alter system archive log current;
        ```

- REDO files

    - You are recommended to mirror the online REDO log files and archive log files using Oracle or the OS.

    - No online REDO log file needs to be backed up.

 Note

For online backup operations, the target Oracle database must use an archive log.

Recovery of Oracle databases

After an Oracle database has stopped, restore the database, and restart it. To recover the database to its latest state, use Oracle's recovery function.

Appendix C Pre-processing and Post-processing of Replication

The replication explains the replication pre-processing and post-processing that are performed in AdvancedCopy Manager.

C.1 Overview

The replication pre-processing and post-processing scripts are started before and after the replication (ie, the copy) processing.

AdvancedCopy Manager performs the processing required for a replication source volume and replication destination volume by executing these scripts.

This section describes the content and setup of pre-processing and post-processing.

 Note

. . . .

If replication is performed in group units, the replication pre-processing and post-processing scripts are not started. Therefore, replication pre-processing and post-processing must be performed manually when replication is performed in group units.

. . . .

C.2 Pre-processing and Post-processing of Replication

The pre-processing and post-processing need to be performed for both the replication source volume and the replication destination volume. The purpose of each is as follows.

- Pre-processing and post-processing for a replication source volume

  This is intended to freeze the replication source data in a state in which consistency is temporarily ensured. In pre-processing, source data is temporarily frozen, and after the replication is complete, the frozen data is released in post-processing.

- Pre-processing and post-processing for a replication destination volume

  This is intended to prevent applications from accessing the copy destination volume during replication. In pre-processing, replication data for the destination is given an access-disabled status. This status is released in post-processing, after the replication is complete.

The specific tasks executed during the pre-processing and post-processing differ according to the data contents of the volume. The next sub section explains the pre-processing and post-processing of the file system.

C.2.1 Pre-processing and post-processing for a file system

The pre-processing and post-processing for a file system is typically the unmounting and remounting of the file system. (For Windows, the processing consists of allocating and cancelling a drive letter.) Pre-processing and post-processing of the replication source volume is performed in order to guarantee the replication source data by disabling access from other processes during execution of synchronous processing and snapshot copy.

If no volume is mounted, the pre-processing and post-processing are not performed.

If the replication destination volume is a Logical Unit (disk), the post-processing for the replication destination disk notifies the OS that the management information (partition table, etc.) of the replication destination disk has been updated.

The implementation status of the pre-processing and post-processing for a file system is shown below.
This table shows whether pre-processing and post-processing are performed depending on the command type. It also shows the copy status at the time of the command execution.

Table C.1 Pre-processing and post-processing for a file system

| Command	Copy status	Source/Destination	Preprocessing	Post-processing
swsrpstartsync (Synchronous processing start command)	Copy not started, or replication establishment	Source	X	X
		Destination	O--*1	X
swsrpmake (Replication creation command)	Equivalency maintenance	Source	O	O
		Destination	X	O--*2
swsrpcancel (Replication cancellation command)	Equivalency maintenance	Source	O	O
		Destination	X	O--*2
	Replication establishment	Source	X	X
		Destination	X	X
	Copying	Source	X	X
		Destination	X	X
swsrpmake (Replication creation command)	Copy not started	Source	O	O
		Destination	O	O

O : Executed,

X : Not executed

The post-processing for the pre-processing executed in step (*1) is executed in step (*2). In the post-processing in (2), re-mount is performed for the mount point where unmount was performed in the pre-processing in (1).

The pre-processing and post-processing for a file system can be customized. The processing is started in accordance with the table above when a replication management command is executed. For replication between servers, these scripts are executed by remote execution using TCP/IP in the pre-processing and post-processing of a volume that is not connected.

The operation image before and after processing is shown in the following figure.

Figure C.1 Synchronized replication processing

If processing other than the standard pre-processing and post-processing is required in the pre-processing and post-processing, customize the scripts.

When customizing a script, strictly observe the following rules regarding error codes:

Table C.2 Error code rules

Error code	Usage
0-99	Unusable (reserved for AdvancedCopy Manager)
100-255	Usable

Note

- 417 -

AdvancedCopy Manager does not support pre-processing and post-processing of databases.Database pre-processing and post-processing must be executed uniquely before and after (respectively) the execution of the AdvancedCopy Manager command.

. . . .

C.2.2 Pre-processing when replication is executed

The replication source volume script (RepSrc.pre) and the replication destination volume script (RepDst.pre) are prepared for a pre-processing script, and it is stored in the following directory. Please customize these scripts according to the processing requirements.

In the case of non-cluster operation

```
/etc/opt/FJSVswsrp/sh
```

In the case of cluster operation

```
/etc/opt/FJSVswsrp/<logical node name>/sh
```

C.2.2.1 Replication source volume pre-processing sample script (RepSrc.pre)

```
 1: #!/bin/sh
 2:
 3: # AdvancedCopy Manager
 4: # Copyright FUJITSU LIMITED, 2002-2012
 5:
 6: #
 7: #   Preprocessing of Replication(Source) processing
 8: #
 9: #       Argument: $1 Device name of Source disk
10: #                 $2 Mount point of Source disk
11: #
12: #    Error number
13: #        2: Argument error
14: #       10: umount error
15: #       13: Illegal mount type (stack/bind mount)
16:
17: # Argument check
18: case $# in
19: 2)
20:     ;;
21: *)
22:     exit 2
23:     ;;
24: esac
25:
26: device=$1
27: mount_point=$2
28:
29: # Determination postprocessing file name
30: if [ "`echo $device | /bin/grep "/dev/sd"`" != "" ]
31: then
32:     # /dev/sd? -> sd?
33:     dev="`echo $device | /bin/sed "s/\/dev\///"`"
34: elif [ "`echo $device | /bin/grep "/dev/vd"`" != "" ]
35: then
36:     # /dev/vd? -> vd?
37:     dev="`echo $device | /bin/sed "s/\/dev\///"`"
38: elif [ "`echo $device | /bin/grep "/dev/FJSV"`" != "" ]
39: then
40:     # /dev/FJSVmphd/dsk/mplb?s? -> mplb?s?
41:     # /dev/FJSVmphd/dsk/mphd?s? -> mphd?s?
```

```
42:     dev="`echo $device | /bin/cut -d/ -f5`"
43: elif [ "`echo $device | /bin/grep "/dev/sfdsk/"`" != "" ]
44: then
45:     if [ "`echo $device | /bin/grep ":"`" != ""    ]
46:     then
47:         devnam="`echo $device | /bin/cut -d: -f2-`"
48:         # /dev/sfdsk/class/dsk/volume:sd? -> class_volume_sd?
49:         dev="`echo $device | /bin/sed "s/\/dev\/sfdsk\///"`"
50:         dev="`echo $dev | /bin/sed "s/\/dsk\//_/"`"
51:         dev="`echo $dev | /bin/sed "s/:/_/"`"
52:         device="`echo $device | /bin/cut -d: -f1`"
53:         if [ "`echo $devnam | /bin/grep "/dev/disk/by-id/"`" != "" ]
54:         then
55:            # /dev/sfdsk/class/dsk/volume:/dev/disk/by-id/<device> -> class_volume__by_id_<device>
56:             dev="`echo $dev | /bin/sed "s/\/dev\/disk\/by-id\//_by-id_/"`"
57:         elif [ "`echo $devnam | /bin/grep "/dev/disk/by-path/"`" != "" ]
58:         then
59:             # /dev/sfdsk/class/dsk/volume:/dev/disk/by-path/<device> ->
class_volume__by_path_<device>
60:             dev="`echo $dev | /bin/sed "s/\/dev\/disk\/by-path\//_by-path_/"`"
61:         fi
62:     else
63:         # /dev/sfdsk/class/dsk/volume -> _gds_class_volume
64:         dev="_gds_`echo $device | /bin/sed "s/\/dev\/sfdsk\///"`"
65:         dev="`echo $dev | /bin/sed "s/\/dsk\//_/"`"
66:     fi
67: elif [ "`echo $device | /bin/grep "/dev/vx/dmp/"`" != "" ]
68: then
69:     # /dev/vx/dmp/device -> _vx_pv_device
70:     dev="_vx_pv_`echo $device | /bin/sed "s/\/dev\/vx\/dmp\///"`"
71: elif [ "`echo $device | /bin/grep "/dev/mapper/"`" != "" ]
72: then
73:     # "/dev/mapper/<device>" -> "_mapper_<device>"
74:     dev="_mapper_`echo $device | /bin/sed "s/\/dev\/mapper\///"`"
75: elif [ "`echo $device | /bin/grep "/dev/disk/by-id/"`" != "" ]
76: then
77:     # "/dev/disk/by-id/<device>" -> "_by-id_<device>"
78:     dev="_by-id_`echo $device | /bin/sed "s/\/dev\/disk\/by-id\///"`"
79: elif [ "`echo $device | /bin/grep "/dev/disk/by-path/"`" != "" ]
80: then
81:     # "/dev/disk/by-path/<device>" -> "_by-path_<device>"
82:     dev="_by-path_`echo $device | /bin/sed "s/\/dev\/disk\/by-path\///"`"
83: else
84:     exit 0
85: fi
86: post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".spre"
87: fstype_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".sfstype"
88:
89: err_log_path="/var/opt/FJSVswsrp/"$SWSTGNODE"/log"
90:
91: if [ "$mount_point" != "" ]
92: # When device was mounted
93: #
94: then
95:
96:         if [ `/bin/cat /proc/mounts |/bin/cut -d' ' -f 2|/bin/grep "^$mount_point\$"|/usr/bin/wc
-w` != 1 ]; then
97:                 # stack mount (multi device on $mount_point)
98:                 /bin/mount  > $err_log_path/$dev.umount 2>&1
99:                 exit 13
100:        fi
101:        if [ `/bin/cat /proc/mounts |/bin/cut -d' ' -f 1|/bin/grep "^$device\$"|/usr/bin/wc -w` !
= 1 ]; then
```

```
102:              cdevice="/dev/`/usr/bin/readlink $device | /bin/sed "s/..\/..\///"`"
103:              if [ `/bin/cat /proc/mounts |/bin/cut -d' ' -f 1|/bin/grep "^$cdevice\
$"|/usr/bin/wc -w` != 1 ]; then
104:                        # bind mount (device has multi mount point)
105:                        /bin/mount  > $err_log_path/$dev.umount 2>&1
106:                        exit 13
107:                    fi
108:          fi
109:
110:      /bin/awk "\$2==\"$mount_point\" {print \$3}" /proc/mounts > $fstype_file
111:
112:      /bin/umount $mount_point 2>/dev/null
113:
114:      if [ $? != 0 ]
115:      then
116:              retry_count=3
117:              sleep_time=1
118:              result_flag=1
119:
120:              while [ $retry_count -gt 0 ]
121:              do
122:                  /bin/umount $mount_point > $err_log_path/$dev.umount 2>&1
123:                  if [ $? != 0 ]
124:                  then
125:                      retry_count=`expr $retry_count - 1`
126:                      /bin/sleep $sleep_time
127:                  else
128:                      /bin/rm -f $err_log_path/$dev.umount
129:                      result_flag=0
130:                      break
131:                  fi
132:              done
133:
134:              if [ $result_flag != 0 ]
135:              then
136:                  /sbin/fuser -vum $mount_point> $err_log_path/$dev.fuser 2>&1
137:                  /bin/ps -ef > $err_log_path/$dev.ps 2>&1
138:                  exit 10
139:              fi
140:      fi
141:
142:      echo "mount,$mount_point" > $post_file
143:
144: # When device was not mounted
145: #
146: else
147:      echo "none" > $post_file
148: fi
149:
150: exit 0
```

C.2.2.2 Replication destination volume pre-processing sample script (RepDst.pre)

```
1: #!/bin/sh
2:
3: # AdvancedCopy Manager
4: # Copyright FUJITSU LIMITED, 2002-2012
5:
6: #
7: #    Preprocessing of Replication(Destination) processing
8: #
```

```
 9: #        Argument: $1 Device name of Destination disk
10: #                 $2 Mount point of Destination disk
11: #
12: #    Error number
13: #         2: Argument error
14: #        10: umount error
15: #        13: Illegal mount type (bind/stack mount)
16:
17: # Argument check
18: case $# in
19: 2)
20:     ;;
21: *)
22:     exit 2
23:     ;;
24: esac
25:
26: device=$1
27: mount_point=$2
28:
29: # Determination postprocessing file name
30: if [ "`echo $device | /bin/grep "/dev/sd"`" != "" ]
31: then
32:     # /dev/sd? -> sd?
33:     dev="`echo $device | /bin/sed "s/\/dev\///"`"
34: elif [ "`echo $device | /bin/grep "/dev/vd"`" != "" ]
35: then
36:     # /dev/vd? -> vd?
37:     dev="`echo $device | /bin/sed "s/\/dev\///"`"
38: elif [ "`echo $device | /bin/grep "/dev/FJSV"`" != "" ]
39: then
40:     # /dev/FJSVmphd/dsk/mplb?s? -> mplb?s?
41:     # /dev/FJSVmphd/dsk/mphd?s? -> mphd?s?
42:     dev="`echo $device | /bin/cut -d/ -f5`"
43: elif [ "`echo $device | /bin/grep "/dev/sfdsk/"`" != "" ]
44: then
45:     if [ "`echo $device | /bin/grep ":"`" != ""   ]
46:     then
47:         devnam="`echo $device | /bin/cut -d: -f2-`"
48:        # /dev/sfdsk/class/dsk/volume:sd? -> class_volume_sd?
49:         dev="`echo $device | /bin/sed "s/\/dev\/sfdsk\///"`"
50:         dev="`echo $dev | /bin/sed "s/\/dsk\//_/"`"
51:         dev="`echo $dev | /bin/sed "s/:/_/"`"
52:         device="`echo $device | /bin/cut -d: -f1`"
53:         if [ "`echo $devnam | /bin/grep "/dev/disk/by-id/"`" != "" ]
54:           then
55:             # /dev/sfdsk/class/dsk/volume:/dev/disk/by-id/<device> -> class_volume__by_id_<device>
56:                dev="`echo $dev | /bin/sed "s/\/dev\/disk\/by-id\///_by-id_/"`"
57:         elif [ "`echo $devnam | /bin/grep "/dev/disk/by-path/"`" != "" ]
58:           then
59:                # /dev/sfdsk/class/dsk/volume:/dev/disk/by-path/<device> ->
class_volume__by_path_<device>
60:                dev="`echo $dev | /bin/sed "s/\/dev\/disk\/by-path\///_by-path_/"`"
61:         fi
62:     else
63:         # /dev/sfdsk/class/dsk/volume -> _gds_class_volume
64:         dev="_gds_`echo $device | /bin/sed "s/\/dev\/sfdsk\///"`"
65:         dev="`echo $dev | /bin/sed "s/\/dsk\//_/"`"
66:     fi
67: elif [ "`echo $device | /bin/grep "/dev/vx/dmp/"`" != "" ]
68: then
69:     # /dev/vx/dmp/device -> _vx_pv_device
70:     dev="_vx_pv_`echo $device | /bin/sed "s/\/dev\/vx\/dmp\///"`"
```

```
 71: elif [ "`echo $device | /bin/grep "/dev/mapper/"`" != "" ]
 72: then
 73:     # "/dev/mapper/<device>" -> "_mapper_<device>"
 74:     dev="_mapper_`echo $device | /bin/sed "s/\/dev\/mapper\///"`"
 75: elif [ "`echo $device | /bin/grep "/dev/disk/by-id/"`" != "" ]
 76: then
 77:     # "/dev/disk/by-id/<device>" -> "_by-id_<device>"
 78:     dev="_by-id_`echo $device | /bin/sed "s/\/dev\/disk\/by-id\///"`"
 79: elif [ "`echo $device | /bin/grep "/dev/disk/by-path/"`" != "" ]
 80: then
 81:     # "/dev/disk/by-path/<device>" -> "_by-path_<device>"
 82:     dev="_by-path_`echo $device | /bin/sed "s/\/dev\/disk\/by-path\///"`"
 83: else
 84:     exit 0
 85: fi
 86: post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".dpre"
 87: fstype_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".dfstype"
 88:
 89: err_log_path="/var/opt/FJSVswsrp/"$SWSTGNODE"/log"
 90:
 91: if [ "$mount_point" != "" ]
 92: # When device was mounted
 93: #
 94: then
 95:
 96:         if [ `/bin/cat /proc/mounts |/bin/cut -d' ' -f 2|/bin/grep "^$mount_point\$"|/usr/bin/wc
-w` != 1 ]; then
 97:                 # stack mount (multi device on $mount_point)
 98:                 /bin/mount  > $err_log_path/$dev.umount 2>&1
 99:                 exit 13
100:         fi
101:         if [ `/bin/cat /proc/mounts |/bin/cut -d' ' -f 1|/bin/grep "^$device\$"|/usr/bin/wc -w` !
= 1 ]; then
102:                 cdevice="/dev/`/usr/bin/readlink $device | /bin/sed "s/..\/..\///"`"
103:                 if [ `/bin/cat /proc/mounts |/bin/cut -d' ' -f 1|/bin/grep "^$cdevice\
$"|/usr/bin/wc -w` != 1 ]; then
104:                         # bind mount (device has multi mount point)
105:                         /bin/mount  > $err_log_path/$dev.umount 2>&1
106:                         exit 13
107:                 fi
108:         fi
109:
110:     /bin/awk "\$2==\"$mount_point\" {print \$3}" /proc/mounts > $fstype_file
111:
112:     /bin/umount $mount_point 2>/dev/null
113:
114:     if [ $? != 0 ]
115:     then
116:             retry_count=3
117:             sleep_time=1
118:             result_flag=1
119:
120:             while [ $retry_count -gt 0 ]
121:             do
122:                 /bin/umount $mount_point > $err_log_path/$dev.umount 2>&1
123:                 if [ $? != 0 ]
124:                 then
125:                     retry_count=`expr $retry_count - 1`
126:                     /bin/sleep $sleep_time
127:                 else
128:                     /bin/rm -f $err_log_path/$dev.umount
129:                     result_flag=0
130:                     break
```

```
131:                    fi
132:                done
133:
134:              if [ $result_flag != 0 ]
135:              then
136:                  /sbin/fuser -vum $mount_point> $err_log_path/$dev.fuser 2>&1
137:                  /bin/ps -ef > $err_log_path/$dev.ps 2>&1
138:                  exit 10
139:              fi
140:    fi
141:
142:    echo "mount,$mount_point" > $post_file
143:
144: # When device was not mounted
145: #
146: else
147:    echo "none" > $post_file
148: fi
149:
150: exit 0
```

C.2.3 Post-processing when replication is executed

The replication source volume script (RepSrc.post) and the replication destination volume script (RepDst.post) are prepared for a post-processing script, which is stored in the following directory. This script must be customized according to the processing requirements.

In the case of non-cluster operation

```
/etc/opt/FJSVswsrp/sh
```

In the case of cluster operation

```
/etc/opt/FJSVswsrp/<logical node name>/sh
```

C.2.3.1 Replication source volume post-processing sample script (RepSrc.post)

```
 1: #!/bin/sh
 2:
 3: # AdvancedCopy Manager
 4: # Copyright FUJITSU LIMITED, 2002-2012
 5:
 6: #
 7: #   Postprocessing of Replication(Source) processing
 8: #
 9: #         Argument: $1 Device name of Source disk
10: #
11: #  Error number
12: #          2: Argument error
13: #         11: mount error
14:
15: # Argument check
16: case $# in
17: 1)
18:    ;;
19: *)
20:    exit 2
21:    ;;
22: esac
23:
24: device=$1
```

```
 25:
 26: # Determination of postprocessing file name
 27: if [ "`echo $device | /bin/grep "/dev/sd"`" != "" ]
 28: then
 29:     # /dev/sd? -> sd?
 30:     dev="`echo $device | /bin/sed "s/\/dev\///"`"
 31: elif [ "`echo $device | /bin/grep "/dev/vd"`" != "" ]
 32: then
 33:     # /dev/vd? -> vd?
 34:     dev="`echo $device | /bin/sed "s/\/dev\///"`"
 35: elif [ "`echo $device | /bin/grep "/dev/FJSV"`" != "" ]
 36: then
 37:     # /dev/FJSVmphd/dsk/mplb?s? -> mplb?s?
 38:     # /dev/FJSVmphd/dsk/mphd?s? -> mphd?s?
 39:     dev="`echo $device | /bin/cut -d/ -f5`"
 40: elif [ "`echo $device | /bin/grep "/dev/sfdsk/"`" != "" ]
 41: then
 42:     if [ "`echo $device | /bin/grep ":"`" != ""   ]
 43:     then
 44:             devnam="`echo $device | /bin/cut -d: -f2-`"
 45:             # /dev/sfdsk/class/dsk/volume:sd? -> class_volume_sd?
 46:             dev="`echo $device | /bin/sed "s/\/dev\/sfdsk\///"`"
 47:             dev="`echo $dev | /bin/sed "s/\/dsk\//_/"`"
 48:             dev="`echo $dev | /bin/sed "s/:/_/"`"
 49:             device="`echo $device | /bin/cut -d: -f1`"
 50:             if [ "`echo $devnam | /bin/grep "/dev/disk/by-id/"`" != "" ]
 51:             then
 52:                     # /dev/sfdsk/class/dsk/volume:/dev/disk/by-id/<device> ->
 class_volume__by_id_<device>
 53:                     dev="`echo $dev | /bin/sed "s/\/dev\/disk\/by-id\//_by-id_/"`"
 54:             elif [ "`echo $devnam | /bin/grep "/dev/disk/by-path/"`" != "" ]
 55:             then
 56:                     # /dev/sfdsk/class/dsk/volume:/dev/disk/by-path/<device> ->
 class_volume__by_path_<device>
 57:                     dev="`echo $dev | /bin/sed "s/\/dev\/disk\/by-path\//_by-path_/"`"
 58:             fi
 59:     else
 60:             # /dev/sfdsk/class/dsk/volume -> _gds_class_volume
 61:             dev="_gds_`echo $device | /bin/sed "s/\/dev\/sfdsk\///"`"
 62:             dev="`echo $dev | /bin/sed "s/\/dsk\//_/"`"
 63:     fi
 64: elif [ "`echo $device | /bin/grep "/dev/vx/dmp/"`" != "" ]
 65: then
 66:     # /dev/vx/dmp/device -> _vx_pv_device
 67:     dev="_vx_pv_`echo $device | /bin/sed "s/\/dev\/vx\/dmp\///"`"
 68: elif [ "`echo $device | /bin/grep "/dev/mapper/"`" != "" ]
 69: then
 70:     # "/dev/mapper/<device>" -> "_mapper_<device>"
 71:     dev="_mapper_`echo $device | /bin/sed "s/\/dev\/mapper\///"`"
 72: elif [ "`echo $device | /bin/grep "/dev/disk/by-id/"`" != "" ]
 73: then
 74:     # "/dev/disk/by-id/<device>" -> "_by-id_<device>"
 75:     dev="_by-id_`echo $device | /bin/sed "s/\/dev\/disk\/by-id\///"`"
 76: elif [ "`echo $device | /bin/grep "/dev/disk/by-path/"`" != "" ]
 77: then
 78:     # "/dev/disk/by-path/<device>" -> "_by-path_<device>"
 79:     dev="_by-path_`echo $device | /bin/sed "s/\/dev\/disk\/by-path\///"`"
 80: else
 81:     exit 0
 82: fi
 83: post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".spre"
 84: fstype_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".sfstype"
 85:
```

```
 86: err_log_path="/var/opt/FJSVswsrp/"$SWSTGNODE"/log"
 87:
 88: # Confirmation of postprocessing
 89: if [ ! -r $post_file ]
 90: then
 91:     exit 0
 92: fi
 93: post="`/bin/cat $post_file | /bin/cut -d',' -f1`"
 94: mount_point="`/bin/cat $post_file | /bin/cut -d',' -f2`"
 95:
 96: # Confirmation of FStype
 97: if [ ! -r $fstype_file ]
 98: then
 99:     fs=""
100: else
101:     fs="`/bin/cat $fstype_file`"
102: fi
103:
104: # No processing
105: if [ "$post" = "none" ]
106: then
107:     /bin/rm -rf $post_file 2> /dev/null
108:     /bin/rm -rf $fstype_file 2> /dev/null
109:     exit 0
110: fi
111:
112: # mount processing
113: if [ "$post" = "mount" ]
114: then
115:     if [ "`echo $device | /bin/grep "/dev/disk/by-id/"`" != "" \
116:            -o "`echo $device | /bin/grep "/dev/disk/by-path/"`" != "" ]
117:     then
118:            cdevice="/dev/`/usr/bin/readlink $device | /bin/sed "s/..\/..\///"`"
119:            Result="`/bin/df -l | /bin/grep "$cdevice " | /bin/awk 'END {print NR}'`"
120:     else
121:            Result="`/bin/df -l | /bin/grep "$device " | /bin/awk 'END {print NR}'`"
122:     fi
123:         if [ "$Result" != "1" ]
124:     then
125:            if [ ! -r $fstype_file ]
126:            then
127:                    /bin/mount $device $mount_point 2> /dev/null
128:            else
129:                    Result1="`echo $fs | /bin/awk 'END {print NR}'`"
130:                    if [ "$Result1" != "1" ]
131:                    then
132:                            /bin/mount $device $mount_point 2> /dev/null
133:                    else
134:                            /bin/mount -t $fs $device $mount_point 2> /dev/null
135:                    fi
136:            fi
137:            if [ $? != 0 ]
138:            then
139:                    retry_count=3
140:                    sleep_time=1
141:                    result_flag=1
142:
143:                    while [ $retry_count -gt 0 ]
144:                    do
145:                            if [ ! -r $fstype_file ]
146:                            then
147:                                    /bin/mount $device $mount_point > $err_log_path/$dev.mount 2>&1
148:                            else
```

```
149:                                           Result1="`echo $fs | /bin/awk 'END {print NR}'`"
150:                                           if [ "$Result1" != "1" ]
151:                                           then
152:                                                   /bin/mount $device $mount_point > $err_log_path/
$dev.mount 2>&1
153:                                           else
154:                                                   /bin/mount -t $fs $device $mount_point > $err_log_path/
$dev.mount 2>&1
155:                                           fi
156:                                   fi
157:                                   if [ $? != 0 ]
158:                                   then
159:                                           retry_count=`expr $retry_count - 1`
160:                                           /bin/sleep $sleep_time
161:                                   else
162:                                           /bin/rm -f $err_log_path/$dev.mount
163:                                           result_flag=0
164:                                           break
165:                                   fi
166:                           done
167:
168:                           if [ $result_flag != 0 ]
169:                           then
170:                                   exit 11
171:                           fi
172:               fi
173:       fi
174:       /bin/rm -rf $post_file 2> /dev/null
175:       /bin/rm -rf $fstype_file 2> /dev/null
176:       exit 0
177: fi
178:
179: exit 0
```

C.2.3.2 Replication destination volume post-processing sample script (RepDst.post)

```
 1: #!/bin/sh
 2:
 3: # AdvancedCopy Manager
 4: # Copyright FUJITSU LIMITED, 2002-2012
 5:
 6: #
 7: #    Postprocessing of Replication(Destination) processing
 8: #
 9: #          Argument: $1 Device name of Destination disk
10: #
11: #  Error number
12: #           2: Argument error
13: #          11 mount error
14:
15: # Argument check
16: case $# in
17: 1)
18:    ;;
19: *)
20:    exit 2
21:    ;;
22: esac
23:
24: device=$1
25:
```

```
26: # Determination of postprocessing file name
27: if [ "`echo $device | /bin/grep "/dev/sd"`" != "" ]
28: then
29:     # /dev/sd? -> sd?
30:     dev="`echo $device | /bin/sed "s/\/dev\///"`"
31: elif [ "`echo $device | /bin/grep "/dev/vd"`" != "" ]
32: then
33:     # /dev/vd? -> vd?
34:     dev="`echo $device | /bin/sed "s/\/dev\///"`"
35: elif [ "`echo $device | /bin/grep "/dev/FJSV"`" != "" ]
36: then
37:     # /dev/FJSVmphd/dsk/mplb?s? -> mplb?s?
38:     # /dev/FJSVmphd/dsk/mphd?s? -> mphd?s?
39:     dev="`echo $device | /bin/cut -d/ -f5`"
40: elif [ "`echo $device | /bin/grep "/dev/sfdsk/"`" != "" ]
41: then
42:     if [ "`echo $device | /bin/grep ":"`" != ""   ]
43:     then
44:             devnam="`echo $device | /bin/cut -d: -f2-`"
45:             # /dev/sfdsk/class/dsk/volume:sd? -> class_volume_sd?
46:             dev="`echo $device | /bin/sed "s/\/dev\/sfdsk\///"`"
47:             dev="`echo $dev | /bin/sed "s/\/dsk\//_/"`"
48:             dev="`echo $dev | /bin/sed "s/:/_/"`"
49:             device="`echo $device | /bin/cut -d: -f1`"
50:             if [ "`echo $devnam | /bin/grep "/dev/disk/by-id/"`" != "" ]
51:             then
52:                     # /dev/sfdsk/class/dsk/volume:/dev/disk/by-id/<device> ->
class_volume__by_id_<device>
53:                     dev="`echo $dev | /bin/sed "s/\/dev\/disk\/by-id\//_by-id_/"`"
54:             elif [ "`echo $devnam | /bin/grep "/dev/disk/by-path/"`" != "" ]
55:             then
56:                      # /dev/sfdsk/class/dsk/volume:/dev/disk/by-path/<device> ->
class_volume__by_path_<device>
57:                     dev="`echo $dev | /bin/sed "s/\/dev\/disk\/by-path\//_by-path_/"`"
58:             fi
59:     else
60:             # /dev/sfdsk/class/dsk/volume -> _gds_class_volume
61:             dev="_gds_`echo $device | /bin/sed "s/\/dev\/sfdsk\///"`"
62:             dev="`echo $dev | /bin/sed "s/\/dsk\//_/"`"
63:     fi
64: elif [ "`echo $device | /bin/grep "/dev/vx/dmp/"`" != "" ]
65: then
66:     # /dev/vx/dmp/device -> _vx_pv_device
67:     dev="_vx_pv_`echo $device | /bin/sed "s/\/dev\/vx\/dmp\///"`"
68: elif [ "`echo $device | /bin/grep "/dev/mapper/"`" != "" ]
69: then
70:     # "/dev/mapper/<device>" -> "_mapper_<device>"
71:     dev="_mapper_`echo $device | /bin/sed "s/\/dev\/mapper\///"`"
72: elif [ "`echo $device | /bin/grep "/dev/disk/by-id/"`" != "" ]
73: then
74:     # "/dev/disk/by-id/<device>" -> "_by-id_<device>"
75:     dev="_by-id_`echo $device | /bin/sed "s/\/dev\/disk\/by-id\///"`"
76: elif [ "`echo $device | /bin/grep "/dev/disk/by-path/"`" != "" ]
77: then
78:     # "/dev/disk/by-path/<device>" -> "_by-path_<device>"
79:     dev="_by-path_`echo $device | /bin/sed "s/\/dev\/disk\/by-path\///"`"
80: else
81:     exit 0
82: fi
83: post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".dpre"
84: fstype_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".dfstype"
85:
86: err_log_path="/var/opt/FJSVswsrp/"$SWSTGNODE"/log"
```

```
 87:
 88: # Confirmation of postprocessing
 89: if [ ! -r $post_file ]
 90: then
 91:    exit 0
 92: fi
 93: post="`/bin/cat $post_file | /bin/cut -d',' -f1`"
 94: mount_point="`/bin/cat $post_file | /bin/cut -d',' -f2`"
 95:
 96: # Confirmation of FStype
 97: if [ ! -r $fstype_file ]
 98: then
 99:    fs=""
100: else
101:    fs="`/bin/cat $fstype_file`"
102: fi
103:
104: # No processing
105: if [ "$post" = "none" ]
106: then
107:    /bin/rm -rf $post_file 2> /dev/null
108:    /bin/rm -rf $fstype_file 2> /dev/null
109:    exit 0
110: fi
111:
112: # mount processing
113: if [ "$post" = "mount" ]
114: then
115: #  df -l $device > /dev/null 2>&1
116: #  if [ $? != 0 ]
117:    if [ "`echo $device | /bin/grep "/dev/disk/by-id/"`" != "" \
118:            -o "`echo $device | /bin/grep "/dev/disk/by-path/"`" != "" ]
119:    then
120:            cdevice="/dev/`/usr/bin/readlink $device | /bin/sed "s/..\/..\///"`"
121:            Result="`/bin/df -l | /bin/grep "$cdevice " | /bin/awk 'END {print NR}'`"
122:    else
123:            Result="`/bin/df -l | /bin/grep "$device " | /bin/awk 'END {print NR}'`"
124:    fi
125:      if [ "$Result" != "1" ]
126:    then
127:            if [ ! -r $fstype_file ]
128:            then
129:                    /bin/mount $device $mount_point 2> /dev/null
130:            else
131:                    Result1="`echo $fs | /bin/awk 'END {print NR}'`"
132:                    if [ "$Result1" != "1" ]
133:                    then
134:                            /bin/mount $device $mount_point 2> /dev/null
135:                    else
136:                            /bin/mount -t $fs $device $mount_point 2> /dev/null
137:                    fi
138:            fi
139:            if [ $? != 0 ]
140:            then
141:                    retry_count=3
142:                    sleep_time=1
143:                    result_flag=1
144:
145:                    while [ $retry_count -gt 0 ]
146:                    do
147:                            if [ ! -r $fstype_file ]
148:                            then
149:                                    /bin/mount $device $mount_point > $err_log_path/$dev.mount 2>&1
```

```
150:                             else
151:                                     Result1="`echo $fs | /bin/awk 'END {print NR}'`"
152:                                     if [ "$Result1" != "1" ]
153:                                     then
154:                                         /bin/mount $device $mount_point > $err_log_path/$dev.mount
2>&1
155:                                     else
156:                                             /bin/mount -t $fs $device $mount_point > $err_log_path/
$dev.mount 2>&1
157:                                     fi
158:                             fi
159:                             if [ $? != 0 ]
160:                             then
161:                                     retry_count=`expr $retry_count - 1`
162:                                     /bin/sleep $sleep_time
163:                             else
164:                                     /bin/rm -f $err_log_path/$dev.mount
165:                                     result_flag=0
166:                                     break
167:                             fi
168:                     done
169:
170:                     if [ $result_flag != 0 ]
171:                     then
172:                             exit 11
173:                     fi
174:             fi
175:     fi
176:     /bin/rm -rf $post_file 2> /dev/null
177:     /bin/rm -rf $fstype_file 2> /dev/null
178:     exit 0
179: fi
180:
181: exit 0
```

C.2.4 Customizing the shell script

The method to customize the shell script for replication pre/post-processing is explained below.

When you do not want to unmount the replication target volume

If you do not want to unmount the mounted copy source volume, then customize the script while referring to "A.2 Pre-processing and Post-processing of Backup".

This is possible only in the following file system:

- ext2

- ext3

Appendix D System Volume Backup Operation

This appendix provides how AdvancedCopy Manager can be used to back up and restore system volumes.

D.1 Overview

In a SAN boot environment where the system volume is allocated to a machine in the ETERNUS Disk storage system, the system volume can be backed up and restored using the replication function.

Figure D.1 Overview of backup and restoration of a system volume

> ## Note
> - A hardware device that supports SAN boot is required.
>
> - The backup target server must be in the stopped state to back up or restore the system volume.
>   This is to ensure the consistency between the data in the backup process and the backed up data.

D.2 Preparation

The server must be stopped in order to back up or restore its system volume, so a separate server is required for backup operations.

That is, at least two servers are required for this operation.

A configuration example for backup and restoration of a system volume is shown below.

Figure D.2 Configuration example for backup and restoration of a system volume

Note

. . . .

The device names (i.e.: /dev/sda1) should be used to specify the root partitions of /etc/grub.conf and /etc/fstab of the server used for backup operations.

Normally, volume labels such as "/" or "/boot" are set for the system volume. The system volume for which the volume label is set is copied to the server used for backup operations by performing backup. Therefore, multiple partitions with the same volume label exist in the server used for backup operations. If the root partition of /etc/grub.conf and /etc/fstab is specified using the volume label name (i.e. LABEL=/ or /boot), the system may not operate correctly because the volume cannot be identified.

. . . .

Point

. . . .

If the backup source server and the server used for backup operations have different operating systems, refer to "10.1.1.10 Notes on executing copy between different operating systems" in "10.1.1 General notes".

. . . .

D.2.1 Registration the Managed Server

The server being backed up and the server used for backup operations must be registered as Managed Servers.

Refer to "7.4.3 Registering a Managed Server" to register Managed Servers.

D.2.2 Registration of device information

The system volume on the server being backed up and the device information of the backup volume on the server used for backup operations must be registered.

Refer to "7.4.4 Fetching device information on a Managed Server" to register the device information.

Note

To back up a system volume that is a GDS logical volume, the SDX object operation unit on the server must be slices.

D.2.3 Settings of replication volume information

In replication volume information, the system volume to be backed up must be set as the source volume and the backup destination volume as the replica volume.

Refer to "7.4.5 Configuring the source volume and replica volume" to set the replication volume information.

Point

Register the backup volume as a volume without a mount point.

Specify the server used for backup operations as the replication destination.

```
# /opt/FJSVswsrp/bin/swsrpsetvol -o REP SYSVOL@TRG-SV SBKVOL@SBK-SV
```

D.3 Backing up the system volume

The replication function is used to back up the system volume.

Stop the server that is to be backed up before performing backup.

Note

If backup is performed while the backup target server is running, the backup may take place while a data update is in progress and without maintaining data consistency. Even though backup is completed, the correct data may not have been backed up.

If the incorrect backed up data is restored, the system may not operate correctly.

Figure D.3 Backing up the system volume

The backup procedure is as follows:

1. Stop the backup target server.

   After the server stops, check that the server power is off.

2. Execute swsrpmake (Replication creation command) on the server used for backup.

   ```
   # /opt/FJSVswsrp/bin/swsrpmake -m SYSVOL@TRG-SV SBKVOL@SBK-SV
   ```

   (*) The command operand was specified in the format "volume name@Managed Server name", as shown in the examples below.

   - name of backup/restore target server: TRG-SV

   - name of system volume: SYSVOL

   - name of server that performs the backup/restore operation: SBK-SV

   - name of backup volume: SBKVOL

3. Activate the backup target server.

D.4 Restoring a system volume

The replication function is used to restore a system volume.

Figure D.4 Restoring a system volume

The restoration procedure is as follows:

1. Stop the restoration target server.

   After the server stops, check that the server power is off.

2. If the system volume on the restoration target server is unusable because of a hardware failure, recover the system volume so that the data can be restored.

3. On the server used for restore operations, execute swsrpmake (Replication creation command) with replication in the reverse direction.

   ```
   # /opt/FJSVswsrp/bin/swsrpmake -m SBKVOL@SBK-SV SYSVOL@TRG-SV
   ```

   (*) The command operand was specified in the format "volume name@Managed Server name", as shown in the examples below.

   - name of backup/restore target server: TRG-SV

   - name of system volume: SYSVOL

   - name of server that performs the backup/restore operation: SBK-SV

   - name of backup volume: SBKVOL

4. Activate the restoration target server.

Appendix E Backup Operations Using Power-Saving

E.1 Overview

Under the power-saving operation, the ETERNUS Disk storage system Eco-mode functionality is used. The processing time of a specific disk drive (RAID group unit) is managed according to the schedule, and the disk drive's motor is stopped if there is no access to the disk for a certain period of time (the period can be set at 10 minutes, 20 minutes, 30 minutes, 40 minutes, 50 minutes and 60 minutes). In this way, power-saving can be achieved.

If there is an access to the disk after the motor has stopped, the disk drive is started again and the access is enabled in one to three minutes' time.

Moreover, the number of accesses is judged to be high if the number of operations of the disk drive in one day exceeds the set range (one to five times can be set). In this situation, the echo mode is not applied, and the operation will continue.

It is possible to save power consumption under the following operations.

- Backup management

  - Snapshot type high-speed backup (OPC) with no access to the backup volume for a long period of time.

  - Suspend/Resume functionality of synchronous type high-speed backup (EC, REC) with a long suspend time.

- Replication management

  - Snapshot type high-speed replication (OPC) with no access to the copied volume as the backup destination for a long period of time.

  - Suspend/Resume functionality of synchronous type high-speed replication with a long suspend time.

E.2 Structure

This section explains the overview of the structure.

Figure E.1 Overview of structure

E.3 Operating Environment

This section explains the operating environment of the backup operation under the power-saving mode.

Device	Description
Management Server	Required operating systems Operating systems that support both the AdvancedCopy Manager's manager and Storage Cruiser Version 14.0 or later (the Management Server function). Required software Storage Cruiser Version 14.0 or later(the Management Server function) **Note**: For installation and setup of Storage Cruiser, refer to Storage Cruiser manual.
Managed Server	Required operating system Operating systems supported by AdvancedCopy Manager Required software None
ETERNUS Disk storage system	For information on the supported ETERNUS Disk storage system, refer to "Supported devices" in the "ETERNUS SF Installation and Setup Guide". For ETERNUS8000, the Eco-mode functionality is available only on ETERNUS8000 model 700 or higher. **Note**: Firmware that can control Eco-mode software is required.

E.4 Environment Settings

E.4.1 Procedure for Configuring the Environment

This section describes how to configure the environment.

1. E.4.2 Settings for ETERNUS Disk storage system

2. E.4.3 Settings for AdvancedCopy Manager

3. E.4.4 Settings for Storage Cruiser

E.4.2 Settings for ETERNUS Disk storage system

To use the Eco-mode, it is required to set the device Eco-mode, create the Eco-mode schedule, and then set the Eco-mode scheduling for the RAID group.
For details on how to configure an ETERNUS Disk storage system, refer to the ETERNUS Web GUI manual.

Creating a RAID group and Logical Unit

The Eco-mode is applied and the RAID group is created.
A RAID group that is to store the data to be backed up (hereafter called "backup source RAID group") and a RAID group that is to store the backup data (hereafter called "backup destination RAID group") are created. A logical Unit is allocated to each RAID group. To effectively reduce power consumption, the backup destination RAID group sets a Logical Unit for each backup that is performed simultaneously. Other Logical Units should not be included in the backup destination RAID group.
The example below illustrates the procedure to simultaneously back up three Logical units (Backup operation A). There are no Logical Units other than the ones for backup operation A in the backup destination RAID group. Therefore, a disk drive can be stopped, except when a backup is running.

The example below illustrates the case where Logical Units for both backup operation B and for business operations are included in the backup destination RAID group, in addition to the Logical Units for backup operation A. As a result, although backup operation A has ended, the Logical Units for backup operation B and for business applications are required to be running.

Point

A backup destination RAID group must be created at each generation if a backup over multiple generations is performed.

Note

A Snap Data Volume and a Snap Data Pool should not be included in a backup destination RAID group.
The RAID group disk drive which includes a Snap Data Volume or a Snap Data Pool does not reach a stop status.

Settings in the Eco-mode

This describes the setting details of the Eco-mode that supports RAID group.
For more details about the settings, refer to the ETERNUS Web GUI manual.

Table E.1 Eco-mode common settings

Setting items	Details of settings
Eco-mode settings	Starting the Eco-mode
Host I/O monitoring time (10, 20, 30, 40, 50, 60 minutes)	Any
Limit of the number of starts/stops in one day (1 to 5 times)	Any

Table E.2 Backup source RAID group - Eco-mode settings

Setting items	Details of settings
Eco-mode schedule ON/OFF	OFF

Table E.3 Backup destination RAID group - Eco-mode settings

Setting items	Details of settings
Eco-mode schedule ON/OFF	ON
Eco-mode schedule	Software control

E.4.3   Settings for AdvancedCopy Manager

AdvancedCopy Manager needs to be configured so that the RAID group for the Eco-mode is used.

Using the backup operation

> When assigning a backup volume, use a logical unit or partition (slice) that belongs to a RAID group for the Eco-mode.

Using the replication operation

> For the destination volumes, use a logical unit or partition (slice) that belongs to a RAID group.

E.4.4   Settings for Storage Cruiser

A Storage Cruiser command is used for the power-saving operation. Therefore, Storage Cruiser needs to be set. For details, refer to the Storage Cruiser manual.

E.5   Operation

This section describes the procedure for establishing power-saving operations. This procedure is to be followed for the Management Server.

1. E.5.1 Start the Disk Drive

2. E.5.2 Operate AdvancedCopy Manager

3. E.5.3 Stop the Disk Drive

E.5.1   Start the Disk Drive

A RAID group disk drive is started using Storage Cruiser disk control command. It is possible to wait for the start time of the disk drive by specifying the -sync option.

Example when the Management Server is Windows

```
C:\ETERNUS_SF\ESC\Manager\opt\FJSVssmgr\sbin\storageadm.exe spindle start -ipaddr <EternusIP> -
raid <RaidGroup> -sync
```

> *<EternusIP>*: IP address of ETERNUS Disk storage system
> *<RaidGroup>*: RAID group number to be used for the backup destination

Example when the Management Server is Solaris or Linux

```
/opt/FJSVssmgr/sbin/storageadm spindle start -ipaddr <EternusIP> -raid <RaidGroup> -sync
```

> *<EternusIP>*: IP address of ETERNUS Disk storage system
> *<RaidGroup>*: RAID group number to be used for the backup destination

For details on this command, refer to the Storage Cruiser manual.

E.5.2   Operate AdvancedCopy Manager

Backup, restoration and replication are performed using AdvancedCopy Manager commands. For details, refer to the description of each operation in this manual.

E.5.3 Stop the Disk Drive

Stopping of a RAID group disk drive is scheduled using the disk control command of the Storage Cruiser. This command returns control after the end of the stopping schedule of the disk drive.

Example when the Management Server is Windows

```
C:\ETERNUS_SF\ESC\Manager\opt\FJSVssmgr\sbin\storageadm.exe spindle stop -ipaddr <EternusIP> -raid
<RaidGroup> -s
```

<*EternusIP*>: IP address of ETERNUS Disk storage system
<*RaidGroup*>: RAID group number to be used for the backup destination

Example when the Management Server is Solaris or Linux

```
/opt/FJSVssmgr/sbin/storageadm spindle stop -ipaddr <EternusIP> -raid <RaidGroup> -s
```

<*EternusIP*>: IP address of ETERNUS Disk storage system
<*RaidGroup*>: RAID group number to be used for the backup destination

For details on this command, refer to the Storage Cruiser manual.

Note

. . . .

Depending on the status of the back up destination RAID group, a disk drive may not stop immediately.

- While OPC physical copying:

  If the OPC physical copy is in progress on the disk drive, the disk drive stops after a certain period once the physical copy is completed.

- While the disk is being accessed:

  If the RAID group for backup destination is being accessed, for example when backing up to a tape, the drive stops after a certain period once the access is stopped.

Other than the above, there are cases that the disk does not stop. For details, refer to the Storage Cruiser manual.

. . . .

Appendix F Operation on KVM Guests

This appendix explains operation using the KVM Support Function.

F.1 Overview

The KVM Support Function serves to execute Advanced Copy functions from a Linux KVM guest OS via the AdvancedCopy Manager's agent on a KVM host.

Figure F.1 Overview of KVM Support Function

F.2 Configuration

The diagram below shows the configuration for using the KVM Support Function.

Figure F.2 Configuration for Using KVM Support Function

- Install ETERNUS SF Manager or AdvancedCopy Manager's agents on both of the following servers:

    - KVM host server

    - Guest server that is used for backup or replication operations

F.3 Software requirements

The table below shows the operating systems requirements for using the KVM Support Function.

Server type	Operating system
Management Server	No requirements.
Managed Server (guest)	Red Hat Enterprise Linux 6 Update 2 (for Intel64) Red Hat Enterprise Linux 6 Update 2 (for x86)
Managed Server (host)	Red Hat Enterprise Linux 6 Update 2 (for Intel64)

F.4 Available functions

This section explains the functions that are available with the KVM Support Function.

F.4.1 AdvancedCopy Manager functions

With the KVM Support Function, the following AdvancedCopy Manager functions can be utilized:

- Advanced Copy function

- Backup function

- Replication function

Out of the Advanced Copy functions, the following types of copy can be utilized:

- OPC

- QuickOPC

- EC

- REC

See
. . . .

For detailed information on the AdvancedCopy Manager functions, refer to "AdvancedCopy Manager functions" in the "ETERNUS SF AdvancedCopy Manager Overview".
. . . .

F.4.2 Online backup

- Symfoware

F.4.3 Cluster operations

On KVM guest OSes, cluster operation utilizing PRIMECLUSTER is supported.

Cluster operation with software other than PRIMECLUSTER is not supported.

F.5 Supported devices

This section explains devices that can utilize the KVM Support Function.

F.5.1 Configuration of Virtual disks

The following notes show the precautions to be observed for adding a virtual disk.

- Specify "virtio" as the bus type and create a virtual disk as a VirtIO device.

    If you use Virtual Machine Manager to create the virtual disk, select "Virtio Disk" as the device type.

    Virtual IDE disks (bus type "IDE") are not supported.

- Set the size of the host disk (logical unit of ETERNUS Disk storage system) to less than 2 TB.
    As a consequence of this setting, the disk size of the guest will also be less than 2 TB.

- As the disk area on the host side, you can specify the device names below:

    - Device names generated by the Linux udev facility (udev device names)

        Example: /dev/disk/by-id/xxxxxxxx (for a disk)
        Example: /dev/disk/by-id/xxxxxxxx-part1 (for a partition)

- SDX object names of PRIMECLUSTER GDS

    Example: /dev/sfdsk/class0001/dsk/volume0001

- If you are using an SDX object of PRIMECLUSTER GDS, you can only use a volume in a single configuration.

    Non-single configuration volumes (mirrored volumes, striped volumes, concatenation volumes, or volumes consisting of tiered disk groups) cannot be used.

- Allocate a disk or partition of the host to the guest.

Figure F.3 Allocating a disk of the host to the guest
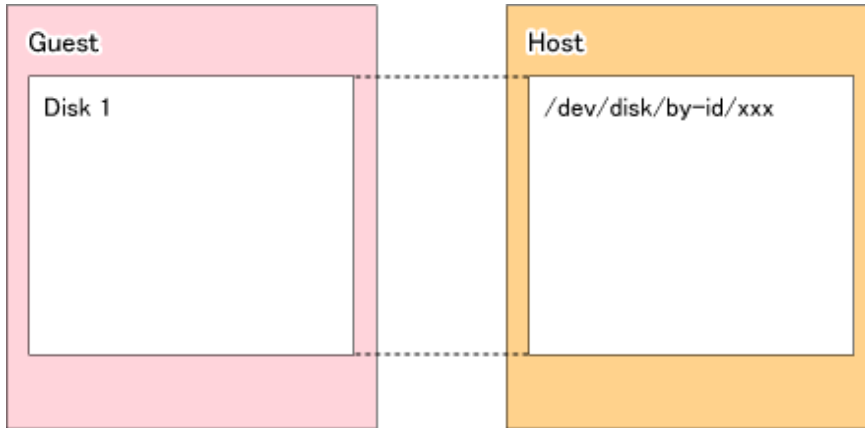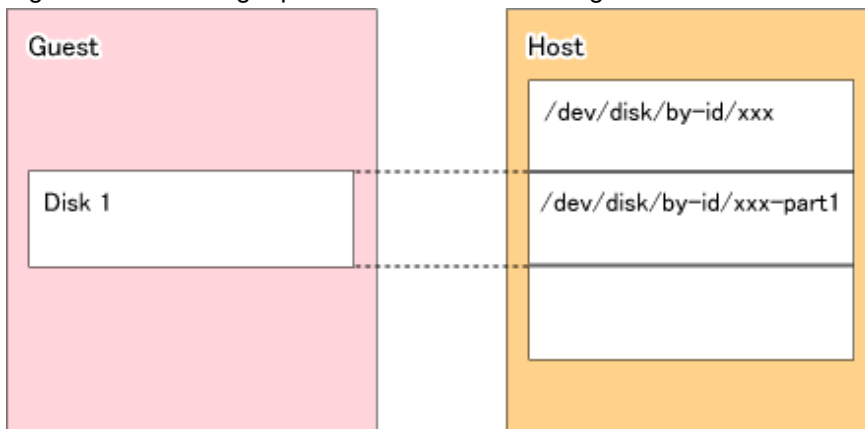

Figure F.4 Allocating a partition of the host to the guest

- If you allocate partitions of the host to the guest, do not connect multiple partitions on the same disk to the same guest.

Figure F.5 Connecting multiple partitions on the same disk to different guests (OK)


Figure F.6 Connecting multiple partitions on the same disk to the same guest (Bad)

F.5.2 Devices that can be specified in AdvancedCopy Manager

Device names that can be specified by AdvancedCopy Manager on KVM guests are as listed in below table.

Device type	Specifiable device name
Normal devices	virtio-blk device names can be specified. When specifiable, the format is as follows: /dev/vd?? Example: /dev/vda1 (for disk partition units) Example: /dev/vda (for Logical Unit units) For details on device names, refer to "1.5.1 Normal devices". Note: Device names generated by Linux udev facility (udev device names) are not supported.
SDX objects of PRIMECLUSTER GDS	Names of devices that operate slice units can be specified. When specifiable, the format is as follows: /dev/sfdsk/<Class name>/dsk/<Volume name>:<DEVNAM value of sdxinfo> Example: /dev/sfdsk/class0001/dsk/volume0001:vda

Device type	Speciﬁable device name
	For details on device names, refer to "1.5.2 SDX objects of PRIMECLUSTER GDS".
	Note: Names of devices that operate per logical volume are not supported.

F.6 Flow of operation

This section explains the flow of operations utilizing the KVM Support Function.

Figure F.7 Flow of operations utilizing the KVM Support Function
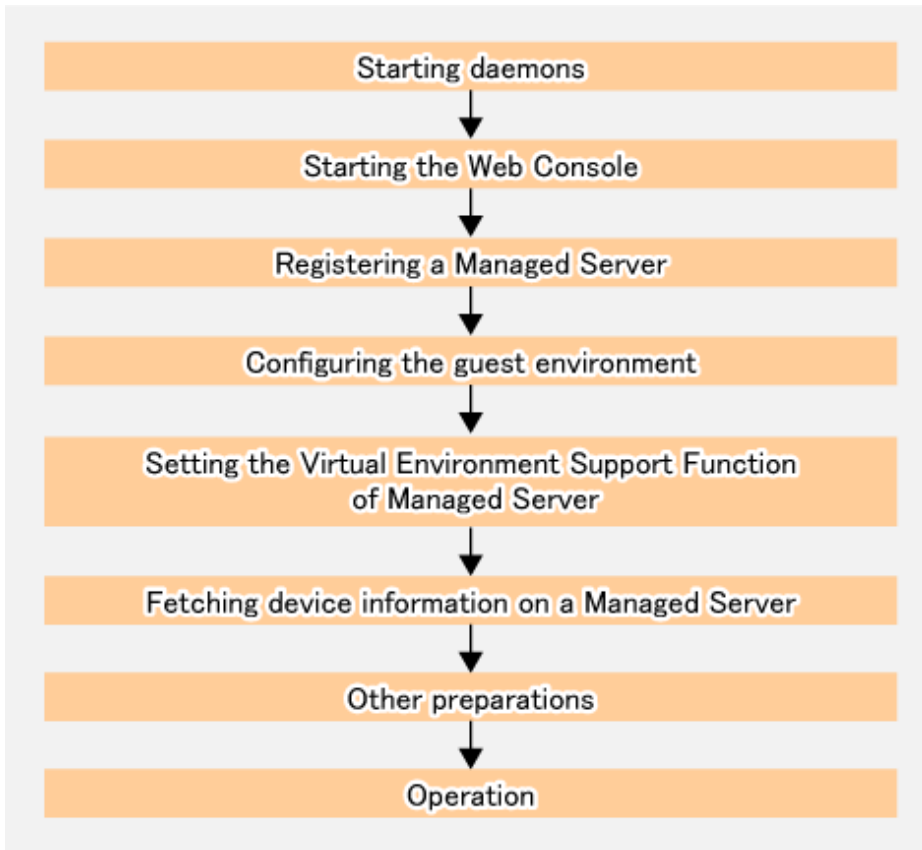
Above work is carried out for the following servers to be controlled and processed.

No.	Contents of work	Object of control	Server to be processed
1	"F.7.1 Starting daemons"	- Management Server - All guests and hosts	- Management Server - All guests and hosts
2	"F.7.2 Starting the Web Console"	PC	PC
3	"F.7.3 Registering a Managed Server"	- Web Console (when using the GUI) - Management Server (when using the CLI)	All guests and hosts
4	"F.7.4 Configuring the guest environment"	All guests	All guests
5	"F.7.5 Setting the Virtual Environment Support Function of Managed Server"	All guests	All guests
6	"F.7.6 Fetching device information on a Managed Server"	- Web Console (when using the GUI)	All guests

No.	Contents of work	Object of control	Server to be processed
		- Management Server (when using the CLI)	
7	"F.7.7 Other preparations"	- Web Console (when using the GUI) - Management Server (when using the CLI)	All guests

F.7 Preparations

This section explains the preparations required before using the KVM Support Function.

F.7.1 Starting daemons

The daemons of AdvancedCopy Manager must be started on both the Management Server and the Managed Servers. Normally, these daemons are automatically started up when the system is started. If the startup fails the system for some reason or if a daemon has been stopped, you will need to start the daemon on each server.

For information on starting a daemon, refer to "Chapter 2 Starting and Stopping Daemons".

F.7.2 Starting the Web Console

Start the Web Console. For details of this, refer to "How to start Web Console" in the "ETERNUS SF Web Console Guide".

F.7.3 Registering a Managed Server

Register all guests and hosts used for KVM Support Function as the Managed Server.

This process is carried out using Web Console.

F.7.4 Configuring the guest environment

Create the following 2 setting files on all guests.

Table F.1 Setting files for KVM Support Function

Setting file	Description	File path
Host information setting file	This file sets the IP addresses and port numbers of the communication daemon for Managed Servers (hosts) to which copying instructions are sent from guests.	/etc/opt/swstorage/vstxapi_hostacm.ini
Guest information setting file	This file sets the domain names of virtual machines on which local guests are started.	/etc/opt/swstorage/vstxapi_domname.ini

[Example of Host information setting file]

```
ipaddr=10.10.10.22
port=1226
```

[Example of Guest information setting file]

```
domname=domain_1
```

For details about the setting files, refer to "F.11 Setting file specifications".

. . . .

For confirming domain names of virtual machines, execute the virsh list --all command on a KVM host.

. . . .

F.7.5   Setting the Virtual Environment Support Function of Managed Server

Use stgxfwcmsetmode (Information collection mode configuration command) to enable the KVM Support Function on all guests.

In a clustered environment, set the logical node name in the environment variable SWSTGNODE before executing the command.

[Example]

```
# export SWSTGNODE=nodeagt
# /opt/FJSVswstf/bin/stgxfwcmsetmode -virt kvm
stgxfwcmsetmode completed (SWSTGNODE=nodeagt)
#
```

F.7.6   Fetching device information on a Managed Server

Fetch the device information that is stored on Managed Servers (guests).
It is not required to carry out this work for Managed Servers (hosts).

Refer to the following sections for details on various control procedures, depending on your particular method of operation.

  - For backup operation on normal transaction volumes: "3.4.4 Fetching device information on a Managed Server"

  - For replication operation: "7.4.4 Fetching device information on a Managed Server"

F.7.7   Other preparations

Carry out the preparations that are required for your particular method of operation.

Refer to the following sections for details on the control procedures:

  - For backup operation on normal transaction volumes: "3.4.5 Configuring the environment for a backup operation server" or later

  - For replication operation: "7.4.5 Configuring the source volume and replica volume" or later

F.8   Operation

This section explains operation of AdvancedCopy Manager using the KVM Support Function.

Refer to the following sections for details on your particular method of operation. However, that there is no need to implement any of the preparations that are described in these reference sections:

  - For backup operation on normal transaction volumes: "3.5 Operation"

  - For backup operation in conjunction with Symfoware: "4.1.5 Operation"

  - For replication operation: "7.5 Operations"

. . . .

When you use the KVM Support Function, data communication with the ETERNUS Disk storage system is carried out via the host. Therefore, compared to a physical server environment, processing takes more time, in proportion with the number of volumes to be processed.

. . . .

F.9　Changing the operating environment

The following section shows how to modify the operating environment for using the KVM Support Function.

For other procedures for modifying the operating environment, refer to "8.5 Changing the Operating Environment".

- "F.9.1 Changing the IP address of Managed Server"

- "F.9.2 Changing the port number of Managed Server"

- "F.9.3 Implementing Live Migration"

F.9.1　Changing the IP address of Managed Server

Change the IP address of a Managed Server (host) as follows:

1. If replication management function was used at the target Managed Server to define it as a copy source or copy destination volume, delete the copy source or copy destination volume settings by referring to "7.6.2 Deleting a source volume or replica volume".

2. On the target Managed Server, stop the AdvancedCopy Manager daemons. For more information on stopping daemons, refer to "Chapter 2 Starting and Stopping Daemons".

3. On the target Managed Server, change the system IP address.

4. On the target Managed Server, restart the AdvancedCopy Manager daemons. For more information on starting daemons, refer to "Chapter 2 Starting and Stopping Daemons".

5. On the all guests, change the IP address set to Host information setting file.
   For details, refer to "F.7.4 Configuring the guest environment".

6. On the Management Server, create a server information change instruction file, and then execute stgxfwcmmodsrv (Server information change command) with the -f option specified to change the IP address.

7. If the replication management function is used, reconstruct the replication management environment (perform the "7.4.5 Configuring the source volume and replica volume" task and subsequent tasks).

F.9.2　Changing the port number of Managed Server

Change the port number of a Managed Server (host) as follows:

1. If replication management function was used at the target Managed Server to define it as a copy source or copy destination volume, delete the copy source or copy destination volume settings by referring to "7.6.2 Deleting a source volume or replica volume".

2. On the target Managed Server, change the port number of the communication daemon (stgxfws) of AdvancedCopy Manager indicated by /etc/services.

3. On the target Managed Server, restart the AdvancedCopy Manager daemons. For information on restarting daemons, refer to "Chapter 2 Starting and Stopping Daemons".

4. On the all guests, change the port number set to Host information setting file.
   For details, refer to "F.7.4 Configuring the guest environment".

5. On the Management Server, create a server information change instruction file, and then execute stgxfwcmmodsrv (Server information change command) with the -f option specified to change the port number.

6. If the replication management function is used, reconstruct the replication management environment (perform the "7.4.5 Configuring the source volume and replica volume" task and subsequent tasks).

See
. . . .
For information about performing the above change in cluster operation, refer to "8.5.7.3 Changing the port number of Management Server transaction or Managed Server transaction".
. . . .

F.9.3   Implementing Live Migration

If you implement live migration while AdvancedCopy Manager is operating, afterwards, you will have to update all Host information setting files.

Point
. . . .
If you know the target hosts before live migration, you can update Host information setting files without editing them during implementation of live migration by preparing templates for Host information setting files in advance. The following is a sample procedure for this update:

1. Before live migration, create a Host information setting file for each target host.

   Store the created Host information setting file in a different folder from the path described in "F.11.1 Host information setting file (vstxapi_hostacm.ini)".

2. Implement live migration.

3. After live migration, update the Host information setting file for the target host.

   Run the OS commands to copy the Host information setting file that you created in step 1 and overwrite the file on the path described in "F.11.1 Host information setting file (vstxapi_hostacm.ini)".
. . . .

F.10   Available commands

The following are the commands that are available with the KVM Support Function.

- "9.2 Backup Management Commands"

- "9.3 Configuration Management Commands"

- "9.4 Replication Management Commands"

- "9.5 Start/Stop Daemon Commands"

- "9.7 Troubleshooting data collection command"

F.11   Setting file specifications

This section explains the setting files that are used by the KVM Support Function.

F.11.1   Host information setting file (vstxapi_hostacm.ini)

The Host information setting file sets the IP addresses and port numbers of the communication daemon for Managed Servers (hosts) to which Advanced Copy instructions are sent. This file has to be created and edited by the users.

Create the Host information setting file so as to be on the following path:

[Path for Host information setting file]

```
/etc/opt/swstorage/vstxapi_hostacm.ini
```

[Contents of Host information setting file]

```
ipaddr=xxx.xxx.xxx.xxx
port=xxxxx
```

Parameter	Value to be set
ipaddr=xxx.xxx.xxx.xxx	IP address of the communication daemon for Managed Server (host) to which Advanced Copy instructions are sent.
port=xxxxx	Port number (1 - 65535) of the communication daemon for Managed Server (host) to which Advanced Copy instructions are sent.

F.11.2 Guest information setting file (vstxapi_domname.ini)

The Guest information setting file sets the domain names of virtual machines on which local guests are started. This file has to be created and edited by the users.

Create the Guest information setting file so as to be on the following path:

[Path for Guest information setting file]

```
/etc/opt/swstorage/vstxapi_domname.ini
```

[Contents of Guest information setting file]

```
domname=xxxxxx
```

Parameter	Value to be set
domname=xxxxxx	Domain name of virtual machine on which local guest is started (max. 50 characters).

Index