



J2SZ-0180-04ENZO (A)

ETERNUS SF AdvancedCopy Manager V13.3

Operator's Guide (HP-UX)



FUJITSU



Preface

Purpose

This manual explains how to operate the Web-GUI with ETERNUS SF AdvancedCopy Manager.

Reader

This manual is intended for system managers who use the ETERNUS SF AdvancedCopy Manager for storage management.

Organization

This manual consists of ten chapters and three appendixes:

Chapter 1 Overview of AdvancedCopy Manager

This chapter provides an overview of AdvancedCopy Manager.

Chapter 2 Starting and Stopping Daemons

This chapter provides information on how to activate and stop the service daemons required by AdvancedCopy Manager.

Chapter 3 Security Operation Using Authentication Feature

This chapter provides information on access control for users who need to perform any of the backup management or replication management functions.

Chapter 4 Backup Operations

This chapter provides details of the backup operations performed by AdvancedCopy Manager.

Chapter 5 Backup Process That Uses the Suspend/Resume Function

This chapter provides details of the backup process that uses the Suspend/Resume function.

Chapter 6 Backup Operation by the QuickOPC Function

This chapter provides details of the backup operation using the QuickOPC function.

Chapter 7 Replication

This chapter provides details of AdvancedCopy Manager replication in an HP-UX system.

Chapter 8 Operational Maintenance

This chapter provides information on how to maintain a database, perform troubleshooting, and change the operation information.

Chapter 9 Commands

This chapter provides information on using commands.

Chapter 10 Notes on Operations

This chapter provides notes on the operations of AdvancedCopy Manager.

Appendix A Pre-processing and Post-processing of Backup and Restore

This appendix provides information on shell scripts used for backup and restore pre-processing and post-processing in HP-UX AdvancedCopy Manager.

Appendix B Replication Pre-processing and Post-processing

This appendix explains the shell scripts that perform the replication pre-processing and post-processing in AdvancedCopy Manager.

Appendix C Sample Shell Script for Oracle Integration 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 using AdvancedCopy Manager.

Related Manuals

This ETERNUS SF AdvancedCopy Manager manual is included in the following series of manuals:

- ETERNUS SF AdvancedCopy Manager Overview
Provides an overview of ETERNUS SF AdvancedCopy Manager.
- ETERNUS SF AdvancedCopy Manager Installation Guide
Describes the installation procedure for ETERNUS SF AdvancedCopy Manager.
- ETERNUS SF AdvancedCopy Manager Operator's Guide
Describes the operating procedures for ETERNUS SF AdvancedCopy Manager.
A different operator's guide is provided for each supported operating system.
- ETERNUS SF AdvancedCopy Manager GUI User's Guide
Describes the operating procedures for the GUI client of ETERNUS SF AdvancedCopy Manager.
- ETERNUS SF AdvancedCopy Manager Message Guide
Explains the messages output by ETERNUS SF AdvancedCopy Manager and associated troubleshooting information.
- ETERNUS SF AdvancedCopy Manager Operator's Guide for a Cluster Environment
Describes the installation and customization procedures for ETERNUS SF AdvancedCopy Manager in a cluster environment.
- ETERNUS SF AdvancedCopy Manager Operator's Guide for a Tape Backup Option
Describes the operating procedures for a tape backup with ETERNUS SF AdvancedCopy Manager.

Users are recommended to read the ETERNUS SF AdvancedCopy Manager Overview first to gain an understanding of the general concepts of this software before reading the other manuals as required.

Abbreviation

- Microsoft (R) Windows (R) 2000 Professional, Microsoft (R) Windows (R) 2000 Server, and Microsoft (R) Windows (R) 2000 Advanced Server are abbreviated to Windows 2000.
- Microsoft (R) Windows Server (R) 2003, Standard Edition, Microsoft (R) Windows Server (R) 2003, Enterprise Edition, Microsoft (R) Windows Server (R) 2003, Enterprise Edition for Itanium-based systems are abbreviated as Windows Server 2003.
- Microsoft (R) Windows Server (R) 2008 Standard, Microsoft (R) Windows Server (R) 2008 Enterprise, Microsoft (R) Windows Server (R) 2008 Datacenter, and Microsoft (R) Windows Server (R) 2008 for Itanium-Based Systems are abbreviated as Windows Server 2008.
- Microsoft (R) Windows (R) XP Professional, and Microsoft (R) Windows (R) XP Home Edition are abbreviated as Windows XP.
- Microsoft (R) Windows Vista (R) Home Basic, Microsoft (R) Windows Vista (R) Home Premium, Microsoft (R) Windows Vista (R) Business, Microsoft (R) Windows Vista (R) Enterprise, and Microsoft (R) Windows Vista (R) Ultimate are abbreviated as Windows Vista.
- Solaris (TM) Operating System is referred to as Solaris.
- "ETERNUS SF AdvancedCopy Manager" is abbreviated as "AdvancedCopy Manager".
- A Manager of ETERNUS SF AdvancedCopy Manager is abbreviated to Manager of AdvancedCopy Manager.
- A Agent of ETERNUS SF AdvancedCopy Manager is abbreviated to Agent of AdvancedCopy Manager.
- ETERNUS2000, ETERNUS3000, ETERNUS4000, ETERNUS6000, ETERNUS8000, and ETERNUS GR series are referred to as ETERNUS storage system.

Trademarks

Windows, Windows Vista, Windows Server, and the names of other Microsoft products are trademarks or registered trademarks of Microsoft Corporation in the United States and other countries.

Netscape, Netscape Navigator and Netscape Communicator are trademarks or registered trademarks of Netscape Communications Corporation in the United States and other countries. Sun, Sun Microsystems, the Sun Logo, Solaris and all Solaris based marks and logos are trademarks of Sun Microsystems, Inc. in the United States and other countries, and are used under license.

UNIX is a registered trademark exclusively owned by X/Open Company Limited in the United States and other countries.

ORACLE is a registered trademark of Oracle Corporation.

HP-UX is a registered trademark of Hewlett-Packard Company.

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 a registered trademark of Red Hat, Inc.

IBM, AIX and AIX 5L are a trademark or registered trademark of International Business Machines Corporation in the United States and other countries.

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

Table of Contents

Chapter 1	Overview of AdvancedCopy Manager	1
1.1	What is AdvancedCopy Manager?.....	2
1.1.1	High-speed backup (Online / Near Real time).....	2
1.1.2	High-speed replication.....	2
1.1.3	Non-stop backup operation [database].....	2
1.2	Product Configuration of AdvancedCopy Manager.....	3
1.2.1	Function system.....	3
1.2.2	System configuration.....	6
1.2.3	Combinations of products.....	6
1.3	Function Overview of AdvancedCopy Manager.....	8
1.3.1	What is the backup function.....	8
1.3.2	What is the replication function.....	8
1.3.3	Authentication feature function.....	9
1.3.4	GUI Client.....	9
1.4	Operating Environment of AdvancedCopy Manager.....	10
1.4.1	Hardware requirements.....	10
1.4.2	Software requirements.....	10
1.5	Managing a Device on AdvancedCopy Manager.....	11
1.5.1	Normal devices.....	11
1.5.2	Volumes under the control of Logical Volume Manager (LVM).....	11
1.5.3	Volumes under the control of VxVM (VERITAS Volume Manager).....	13
Chapter 2	Starting and Stopping Daemons	15
2.1	Daemons that Configure AdvancedCopy Manager.....	16
2.2	Starting Daemons.....	17
2.3	Stopping Daemons.....	18
Chapter 3	Security Operation Using Authentication Feature	19
Chapter 4	Backup Operations	21
4.1	Overview.....	22
4.1.1	Snapshot backup processing.....	23
4.1.2	Synchronized high-speed backup processing.....	25
4.1.3	Restoring data.....	27
4.2	Configuration.....	29
4.2.1	Selecting a server on which the backup operation should be performed.....	29
4.2.2	Selecting a backup target.....	30
4.2.3	Configuring a backup policy.....	31
4.2.4	Preparing a backup volume.....	31
4.2.5	Automating backup operations.....	33
4.3	Flowchart of Backup Operation.....	35
4.4	Preparing to Start the Backup.....	37
4.4.1	Activating daemons.....	37
4.4.2	Setting an access permission.....	37
4.4.3	Starting the GUI client.....	37
4.4.4	Registering a Storage Server.....	37
4.4.5	Fetching device information on a Storage Server.....	37
4.4.6	Displaying the backup management window.....	38
4.4.7	Configuring the environment for a backup operation server.....	38
4.4.8	Configuring the operation type for a device.....	38
4.4.9	Configuration options.....	40

4.4.10	Configuring backup policies.....	41
4.4.11	Customizing Pre-processing and Post-processing scripts.....	41
4.4.12	Preparing a device map file.....	42
4.5	Operation	44
4.5.1	Performing snapshot high-speed backup.....	44
4.5.2	Performing the synchronized high-speed backup.....	44
4.5.3	Restoring Data.....	45
4.5.4	Managing a backup history.....	46
4.6	Changing the configuration	46
4.6.1	Adding a device.....	46
4.6.2	Deleting a device.....	47
4.6.3	Adding a Storage Server.....	48
4.6.4	Deleting a Storage Server.....	49
4.7	Stopping the Operation	49
4.8	LVM Volume Operation	50
4.8.1	Backup operation in volume group units.....	50
4.8.2	Backup operation in units of physical disks.....	54
4.9	VxVM volume operation	62
4.9.1	Operational configuration.....	62
4.9.2	Preparations.....	63
4.9.3	Executing backup.....	64
4.9.4	Restoring.....	66
Chapter 5	Backup Process That Uses the Suspend/Resume Function.....	69
5.1	Overview	70
5.2	Backup configuration	72
5.3	Flow of Operations	73
5.4	Preparations	74
5.4.1	Configuring up the AdvancedCopy Manager environment.....	74
5.4.2	Configuring Storage Server information.....	74
5.4.3	Configuring device information.....	74
5.4.4	Configuring a backup policy.....	75
5.4.5	Preparing a device map file.....	75
5.5	Backup Operations	76
5.5.1	Starting synchronous processing.....	78
5.5.2	Backing up.....	79
5.5.3	Displaying the synchronous processing execution status.....	80
5.5.4	Canceling synchronous processing.....	81
5.5.5	History Management.....	84
5.5.6	Restoring.....	85
Chapter 6	Backup Operation by the QuickOPC Function.....	87
6.1	Overview	88
6.2	Configuring Backup Operation	90
6.3	Operation Flowchart	91
6.4	Preparation	92
6.4.1	Configuring the AdvancedCopy Manager environment.....	92
6.4.2	Configuring the environment for Storage Server.....	92
6.4.3	Configuring the operation type for a device.....	92
6.4.4	Configuring a backup policy.....	93
6.4.5	Preparing a device map file.....	93
6.5	Operation	94
6.5.1	Executing backup.....	94
6.5.2	History management.....	97
6.5.3	Restoration.....	98

6.5.4 Canceling restoration.....	99
6.5.5 Displaying the execution status of tracking processing.....	100
6.5.6 Stopping tracking processing.....	100
Chapter 7 Replication.....	101
7.1 Overview.....	102
7.1.1 Snapshot replication processing.....	102
7.1.2 Synchronized replication processing.....	111
7.2 Operational configuration.....	114
7.2.1 Determining the server that performs replication.....	114
7.2.2 Determining a replication target.....	115
7.2.3 Notes on the configuration of SnapOPC/SnapOPC+ replication.....	117
7.3 Flow of Operations.....	124
7.4 Preparations.....	125
7.4.1 Starting a daemon.....	125
7.4.2 Starting the GUI client.....	125
7.4.3 Registering a Storage Server.....	125
7.4.4 Fetching device information from a Storage Server.....	125
7.4.5 Configuring the source volume and replica volume.....	126
7.4.6 Customizing a script of Pre-processing and Post-processing.....	126
7.4.7 Creating groups.....	127
7.5 Operations.....	132
7.5.1 Executing snapshot replication.....	132
7.5.2 Executing synchronized replication.....	143
7.5.3 Backup and restoration using the replication function.....	152
7.5.4 Executing synchronous-type reversal.....	153
7.6 Changing the Configuration Information.....	154
7.6.1 Changing a source volume or replica volume.....	154
7.6.2 Deleting a source volume or replica volume.....	154
7.6.3 Changing the device information of a source or replica volume.....	154
7.7 Stopping the Operation.....	155
7.8 LVM Volume Operation.....	156
7.8.1 Replication in units of volume groups.....	156
7.8.2 Replication in units of physical disks.....	159
7.9 VxVM Volume Operation.....	164
7.9.1 Operational configuration.....	165
7.9.2 Preparations.....	165
7.9.3 Replication.....	166
Chapter 8 Operational Maintenance.....	169
8.1 Maintaining a Database.....	170
8.1.1 Making individual backups of databases.....	171
8.1.2 Making a package backup of databases.....	174
8.2 Troubleshooting.....	186
8.2.1 Insufficient free space in a repository.....	186
8.2.2 Collecting maintenance information.....	186
8.2.3 Failure in accessing a repository database.....	187
8.3 Troubleshooting a hardware error during backup.....	188
8.3.1 Overview.....	188
8.3.2 Troubleshooting.....	188
8.4 Troubleshooting a Hardware Error or Other Error Occurring during Replication.....	192
8.4.1 Overview.....	192
8.4.2 Troubleshooting.....	192
8.5 Changing the Operating Environment.....	197
8.5.1 Changing the IP address of a Storage Management Server.....	197

8.5.2 Changing the IP address of a Storage Server	197
8.5.3 Changing the port number of a Storage Management Server or Storage Server.....	198
8.5.4 Changing the server name of a Storage Management Server or Storage Server.....	198
8.5.5 Changing the device configuration.....	199
8.5.6 Changing the IP address for GUI client connection of a Storage Management Server.....	200
8.5.7 Change of password or deletion of the repository access user.....	200
8.5.8 Changing the operational configuration in cluster operation.....	200
Chapter 9 Commands	207
9.1 Command list	208
9.2 Backup Management Commands	212
9.2.1 Environment definition commands.....	215
9.2.2 Operation commands.....	225
9.2.3 Maintenance commands.....	253
9.3 Configuration Management Commands	257
9.3.1 Server information change command (stgcmmodnode).....	258
9.3.2 Repository access user change command (stguserset).....	259
9.3.3 GUI client connection information Configuration command (stgguipset).....	259
9.3.4 Server information addition command (stgxfwcmaddsrv).....	259
9.3.5 Server information change command (stgxfwcmmodsrv).....	259
9.3.6 Server information deletion command (stgxfwcmdelsrv).....	259
9.3.7 Device information collection/reflection command (stgxfwcmsetdev).....	259
9.3.8 Device information deletion command (stgxfwcmdeldev).....	260
9.3.9 Server information display command (stgxfwcmdispsrv).....	260
9.3.10 Device information display command (stgxfwcmdispdev).....	260
9.3.11 Partition information display command (stgxfwcmdisppat).....	260
9.3.12 Environment information display command(stgenvdisp).....	260
9.3.13 Management information package backup server registration/deletion command (stgmgrinfoaset).....	264
9.3.14 Management information package backup command (stgmgrinfobkup).....	265
9.3.15 Management information package backup status display command (stgmgrinfodisp).....	268
9.4 Replication Management Commands	270
9.4.1 Environment definition commands.....	272
9.4.2 Operational commands.....	284
9.4.3 Maintenance commands.....	306
9.5 Start/Stop Daemon Commands	308
9.5.1 Starting and stopping the communication daemon.....	308
9.5.2 Starting and stopping AdvancedCopy Manager daemons.....	308
9.6 SnapOPC/SnapOPC+ Support Commands	309
9.6.1 Update volume measurement command (swstestupdate).....	309
9.6.2 Snap Data Volume operation/reference command (swstsdv).....	310
Chapter 10 Notes on Operations	315
10.1 Notes on Backup and Replication Operation	316
10.1.1 General notes.....	316
10.1.2 Notes on the snapshot high-speed backup.....	319
10.1.3 Notes on synchronous high-speed backup.....	319
10.1.4 Notes on restoration.....	319
10.1.5 Notes on synchronous replication.....	320
10.1.6 Notes on cluster operation.....	320
10.1.7 Notes on backups in cluster operation.....	321
10.1.8 Notes on replication in a cluster system.....	323
10.1.9 Notes on using ETERNUS VS900 (virtualization switch).....	326

Appendix A	Pre-processing and Post-processing of Backup and Restore	329
A.1	Overview	330
A.2	Pre-processing and Post-processing of Backup	331
A.2.1	Backup Pre-processing.....	331
A.2.2	Backup Post-processing.....	336
A.2.3	Customizing the shell script.....	342
A.3	Restore Pre-processing and Post-processing	345
A.3.1	Restore Pre-processing.....	345
A.3.2	Restore Post-processing.....	349
A.3.3	Customizing the shell script.....	354
Appendix B	Replication Pre-processing and Post-processing	357
B.1	Overview	358
B.2	Replication Pre-processing and Post-processing	359
B.2.1	Pre-processing and post-processing for a file system.....	359
B.2.2	Pre-processing when replication is executed.....	361
B.2.3	Post-processing when replication is executed.....	369
B.2.4	Customizing the shell script.....	378
Appendix C	Sample Shell Script for Oracle Integration Backup Operations	383
C.1	Overview	384
C.2	Backup and Restoration of a Transaction volume	385
C.2.1	Backup procedure.....	385
C.2.2	Restore procedure.....	385
C.3	Sample Shell Scripts	386
C.3.1	Overview of sample shell scripts.....	386
C.3.2	Applying the sample shell scripts.....	391
C.3.3	Restore Method.....	392
C.3.4	Notes about using the sample shell scripts.....	393

Chapter 1 Overview of AdvancedCopy Manager

This chapter provides an overview of AdvancedCopy Manager.

1.1 What is AdvancedCopy Manager?

AdvancedCopy Manager provides data storage management in an open systems environment and includes the following three features:

1.1.1 High-speed backup (Online / Near Real time)

The Advanced Copy Function (OPC, EC) is used to back up data quickly. The advanced copy function provides fast data copy regardless of data size, and can maintain a fixed backup time even if the amount of data to be backed up increases. Therefore, the backup time does not increase even if the data amount increases.

This function can dramatically reduce the duration for which applications must be stopped to perform backup.

With the remote copy function, it is also possible to backup data between disk array devices at remote locations.

Shell scripts are used in conjunction with DBMS environments to provide online or near real time, high-speed copying without interruption to live processing.

- Oracle

For applications using a database engine other than Oracle or that use a general file system, downtime or interruption to normal throughput is still significantly reduced.

1.1.2 High-speed replication

The Advanced Copy Function (OPC/EC) enables the same data to be made available to multiple, unrelated applications in addition to providing complete dataset data copies for redundancy and integrity management.

Using the remote copy functionality, it is possible to copy whole datasets to disk array devices at remote locations.

Shell scripts are used in conjunction with Oracle DBMS environments to provide online or near real time, high-speed copying without interruption to live processing.

1.1.3 Non-stop backup operation [database]

By using AdvancedCopy Manager, an Oracle database can be backed up without stopping the application. Refer to "Sample Shell Script for Oracle Integration Backup Operations" for information about the backup method for Oracle.



As an example of copying Oracle databases, a script created for replication is shown in "Sample Shell Script for Oracle Integration Backup Operations".

1.2 Product Configuration of AdvancedCopy Manager

1.2.1 Function system

The HP-UX version of AdvancedCopy Manager encompasses of the following functionality:

Table: Functions

Function name	Description
HP-UX version Agent of AdvancedCopy Manager	Operates on the Storage Management Server, and manages two or more Storage Servers. Backup management and replication management are also performed. The Agent component can be managed by an ACM client manager component running on Solaris, Windows, or Linux operating systems.
AdvancedCopy Manager GUI client	The GUI client can run on the user's local machine in the same way as other application and it can be used to monitor and manage the AdvancedCopy Manager environment. To do this, start the GUI client and then connect to the Storage Management Server.

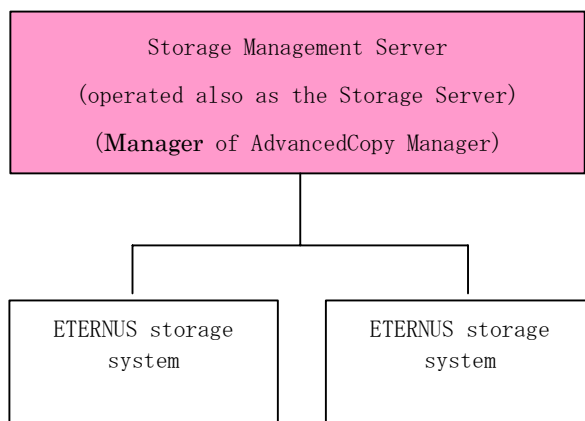
The AdvancedCopy Manager server component type will depend on the operating environment.

- Storage Server
The Storage Server executes the backup and replication operations. The Agent is installed on the Storage Server. If a Storage Server is also to be used as the Storage Management Server then both the Manager and the Agent will need to be installed. The Agent installation is included as part of the Manager installation.
- Storage Management Server
The Storage Management Server manages the Storage Servers, and only one Storage Management Server can be configured in the system.
The Manager is installed on the Storage Management Server.

Three AdvancedCopy Manager configuration scenarios are possible and their environments are described below:

- When only one Storage Server is used.
In this operating environment, only the Manager needs to be installed as the Manager incorporates the Agent functionality.


Figure: When only one Storage Server is installed



- When two or more Storage Servers are used
When two or more AdvancedCopy Manager servers are available, only one can be configured as the Storage Management Server or alternatively, the storage management component can be configured on a separate server unrelated to the core ETERNUS storage system. Each scenario is explained below.


- When the Storage Server is also operated as a Storage Management Server
In this scenario, one of the servers functions as a Storage Management Server. The Storage Management Server has a repository where the information on operating and managing all Storage Servers is stored.
When many Storage Servers are used, they all compete for access to the data repository. This will generally have an adverse effect on the performance of other applications running on the same server. It is therefore recommended that the Storage Management Server functionality should be installed on a dedicated server. The following figure illustrates the configuration of the Storage Server when it is also operated as the Storage Management Server.

Figure: When the Storage Server is also operated as a Storage Management Server



- Using a dedicated Storage Management Server
This scenario comprises two or more Storage Servers with a dedicated Storage Management Server management client unrelated to core ETERNUS storage system functionality.

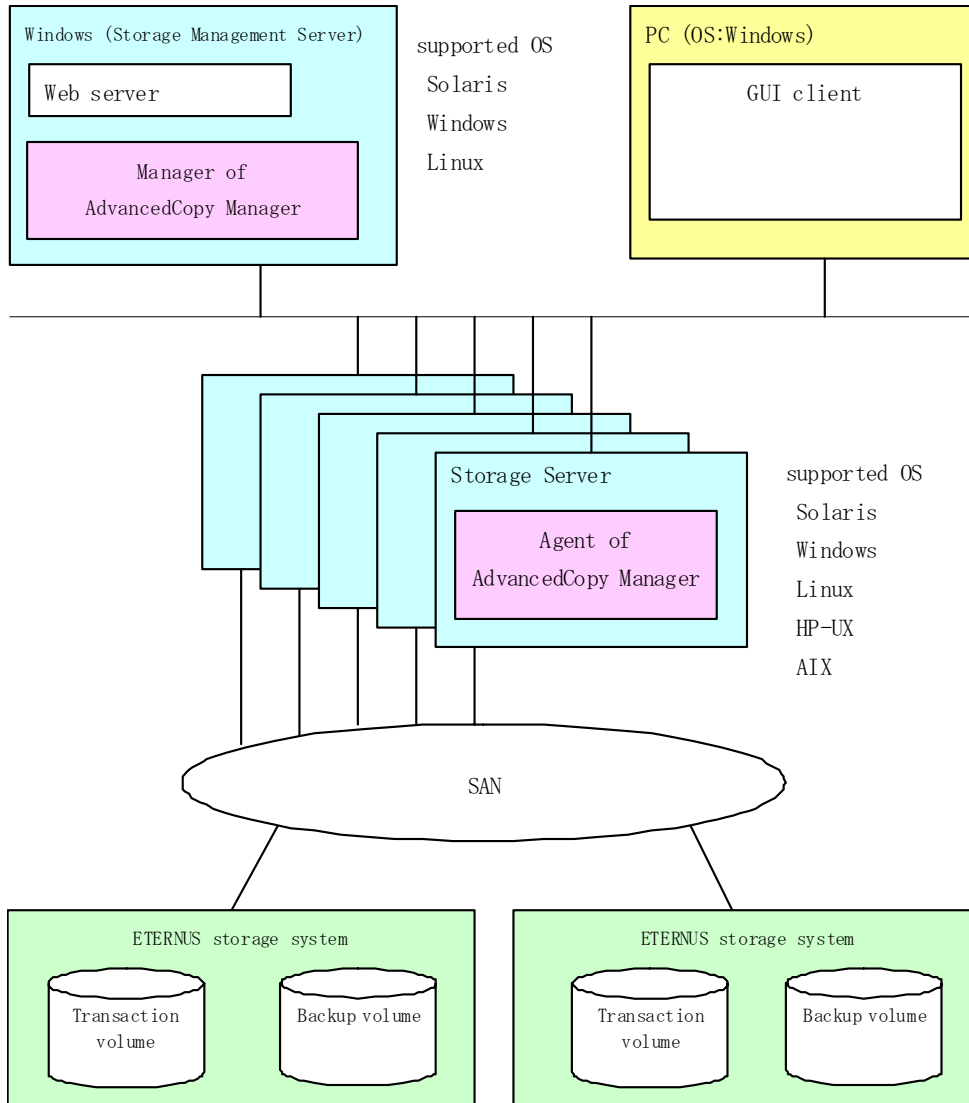
Figure: When a dedicated Storage Server is a Storage Management Server



1.2.2 System configuration

The system configuration of AdvancedCopy Manager is shown below.

Figure: System configuration of a Storage Management Server and two or more Storage Servers



1.2.3 Combinations of products

The following table details combinations of compatible product revision levels of the AdvancedCopy Manager Agent and Manager components that can be used on different OS platforms.

Table: Combinations of products

		Agent platform (Storage Server)															
		Solaris			Windows			Linux			HP			AIX			
		10. x	11. x	13. x	V10.0 Lxx	V11.0 Lxx	13. x	V10.0 Lxx	V11.0 Lxx	13. x	10. x	11. x	13. x	10. x	11. x	13. x	
Manager platform (Storage Management Server)	Solaris	10. x	A	C	C	A	C	C	A	C	C	A	C	C	A	C	C
		11. x	B	A	C	B	A	C	B	A	C	B	A	C	B	A	C
		13. x	B	B	A	B	B	A	B	B	A	B	B	A	B	B	A
	Windows	V10.0 Lxx	A	C	C	A	C	C	A	C	C	A	C	C	A	C	C
		V11.0 Lxx	B	A	C	B	A	C	B	A	C	B	A	C	B	A	C
		13. x	B	B	A	B	B	A	B	B	A	B	B	A	B	B	A
	Linux	V10.0 Lxx	A	C	C	A	C	C	A	C	C	A	C	C	A	C	C
		V11.0 Lxx	B	A	C	B	A	C	B	A	C	B	A	C	B	A	C
		13. x	B	B	A	B	B	A	B	B	A	B	B	A	B	B	A

A : Compatible

B : Compatible (However, any functionality supported by later versions of the Storage Server agent is unusable.)

C : Incompatible



Note

- There is a difference in functionality between some product revision levels. Installation of a Manager and Agent which operate at the same revision level is recommended.
- The following restrictions apply if different levels of the same version are combined:
 - If the Manager is a newer level than the Agent
Compatibility is possible (but functionality only supported by levels later than that of the Storage Server cannot be used)
 - If the Agent is a newer level than the Manager
Compatibility not possible
- The versions of ETERNUS SF AdvancedCopy Manager listed as 10. x, 11. x, V10. 0Lxx, V11. 0Lxx in "Table: Combinations of products" are Softek versions.

1.3 Function Overview of AdvancedCopy Manager

1.3.1 What is the backup function

The backup function provides support for high-speed backup by using the One Point Copy (OPC), Remote One Point Copy (ROPC), Equivalent Copy (EC), and Remote Equivalent Copy (REC) functions of a Fujitsu ETERNUS Storage System.

The following sub-sections provide more information on these backup functions.

High-speed backup operation

The advanced copy functions of the disk array unit (for example the OPC and EC functions) are 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 updates of differential copies using the Advanced Copy function.

Policy management function

The policy management function manages the number of preserved backup generations and the number of days between backups(interval). If this interval is exceeded, resulting in a delay in performing the scheduled backup, a warning is provided to the operator when displaying information relating to affected volumes.

Storage configuration management function

In the AdvancedCopy Manager environment, a volume to be backed up is called a transaction volume and a volume on which a backup should be saved is called a backup volume.

The storage configuration management function manages the configurations of both source transaction volumes and destination backup volumes.

Backup and restore functions

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



Backup using ROPC/REC is limited to disk array devices within the same Storage Server that support ROPC/REC.

History information management function

The history information management function manages the backup history information in preserved backup generations using previously backed up data.

Consistency recovery function

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

1.3.2 What is the replication function

The replication function manages data replication by using the One Point Copy (OPC), Remote One Point Copy (ROPC), Equivalent Copy (EC), and Remote Equivalent (REC) functions of the Fujitsu ETERNUS Storage Systems.

Features of the replication function include the following:

High-speed replication operation

The advanced copy functions of the disk array unit (ie, the OPC and EC functions) are used to perform high-speed replication at the hardware level between volumes regardless of volume capacities. This function greatly reduces the operational impact of a replication operation.

Faster replication can be achieved by performing differential copies where only data which has changed or has been updated is copied using the Advanced Copy function.

Consistency recovery function

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

Main differences with backup function

While the backup function offers policy management and history management, these features are not available as part of the replication function.

The replication function merely replicates data between two associated volumes.

1.3.3 Authentication feature function

The GUI client provided by AdvancedCopy Manager can be used to manage the users who need to have access to the system. It allows the configuration of secure operational information specifying the range of functionality each user is entitled to access. Refer to "Security Operation Using Authentication Feature".

1.3.4 GUI Client

The following environmental configuration and functionality is available using the GUI client:

- Security management
Operational restrictions on users can be set up using the GUI client.
- Backup management
Backup management can be performed via the GUI client (eg, to define a server or a device to be used for backup or restoration).
- Replication management
AdvancedCopy Manager replication management tasks (such as defining volumes and associations and executing replications) can be performed using the GUI client.

A GUI client can be used for accessing a management server.

The AdvancedCopy Manager GUI client must be installed as the GUI client. For details, refer to the "ETERNUS SF AdvancedCopy Manager Installation Guide".

The hardware on which the Web browser is running does not have to have AdvancedCopy Manager installed.

Backup and replication management tasks may also be performed by executing the relevant commands at a command prompt as well as the GUI client.

For more information on the GUI client, refer to "GUI Client" in "ETERNUS SF AdvancedCopy Manager GUI User's Guide".

1.4 Operating Environment of AdvancedCopy Manager

1.4.1 Hardware requirements

One of the following hardware devices is required to perform backup or replication management. In addition, an advanced copy license and a suitable cache table size need to be set up. The advanced copy functions are different for each hardware device, so it will be necessary to check the required settings with a Fujitsu systems engineer (SE) before use.

- ETERNUS2000 disk array unit (model 100 or higher)
- ETERNUS3000 disk array unit (model 100 or higher)
- ETERNUS4000 disk array unit (model 100 or higher)
- ETERNUS6000 disk array unit
- ETERNUS8000 disk array unit
- ETERNUS GR720 disk array unit
- ETERNUS GR730 disk array unit
- ETERNUS GR740 disk array unit
- ETERNUS GR820 disk array unit
- ETERNUS GR840 disk array unit

1.4.2 Software requirements

1.4.2.1 Agent of AdvancedCopy Manager

Table: Software requirements for the Agent of AdvancedCopy Manager

Product name	Requirement	Remarks
HP-UX 11.00 HP-UX 11i v1 HP-UX 11i v2 (PA-RISC) HP-UX 11i V2.0 (Itanium)	Required	Support OS
MC/Service Guard 11.13 VERITAS Cluster Server 5.0	Conditionally required	For cluster operation with the Agent of AdvancedCopy Manager
VERITAS Volume Manager 3.5, 5.0	Conditionally required	When a logical volume of VERITAS Volume Manager is the target of the operation performed.

1.5 Managing a Device on AdvancedCopy Manager

1.5.1 Normal devices

The following devices are supported for backup and replication management:

- sd driver resources

Note that the following device is not supported for management by HP-UX 11i v2.0 (Itanium):

- GPT disk

1.5.2 Volumes under the control of Logical Volume Manager (LVM)

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

- Volume group
- Physical disk that constitutes a volume group

For information on the backup and replication of LVM volumes, refer to the "LVM Volume Operation" sections in the chapters on Backup and Replication.



Note

Do not set LVM physical volume groups. If set, LVM device information cannot be fetched correctly.

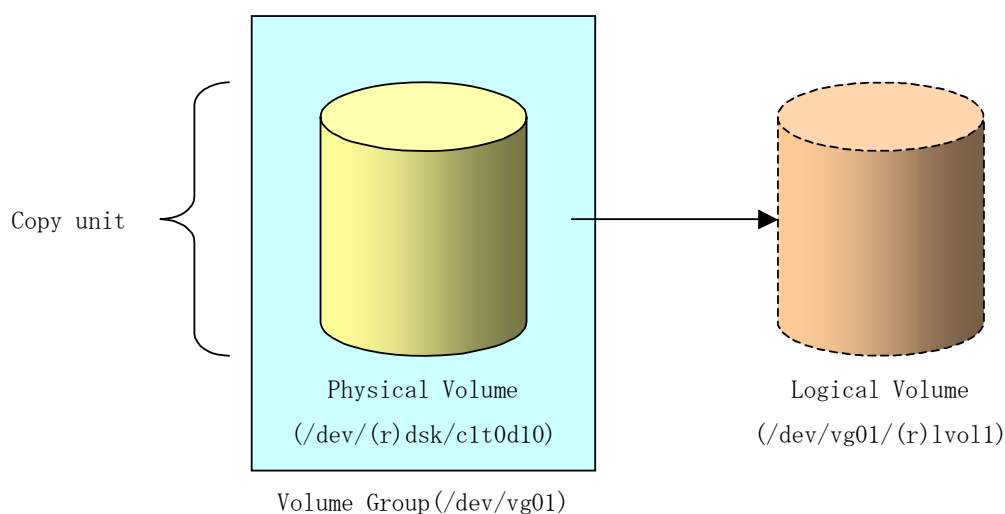
1.5.2.1 Backup/replication in units of volume groups

LVM volumes are volumes which are managed in units of volume groups containing logical volumes. AdvancedCopy Manager copies each physical disk that constitutes a volume group.

An LVM volume that can be operated in units of logical volumes must satisfy the following conditions:

- One volume group contains only one physical disk, and logical volumes are configured so that one physical disk includes <n> logical volumes.

Figure: Support configuration



LVM allows multiple logical volumes to be defined on a single physical disk (a logical volume is referred to as LU: Logical Unit). The boundaries of logical volumes can be defined independently of the boundaries of the host physical disk.

A configuration in which one volume group contains multiple logical volumes is also supported.

For operation, please note the following:

- A logical volume for which a file system was constructed must be unmounted before backup or replication can be performed, and remounted after the backup or replication has completed.

1.5.2.2 Physical disks that constitute a volume group


In this case, the individual physical disks that constitute a volume group in which a logical volume exists are used as management units.

Use this type of operation when a volume group has a configuration that is not supported for operations in units of volume groups.



For operations in a cluster configuration, the device name (`/dev/(r) dsk/c#t#d#`) of the physical disks that comprise a volume group must be the same at all servers comprising the cluster, and all the ETERNUS disks indicated by the device name must be the same.

Figure: Configuration that is not supported for operation in units of volume groups



Copying during backup or replication must be performed while maintaining the integrity of an entire volume group.

1.5.3 Volumes under the control of VxVM (VERITAS Volume Manager)

AdvancedCopy Manager can replicate VxVM volumes in the following units:

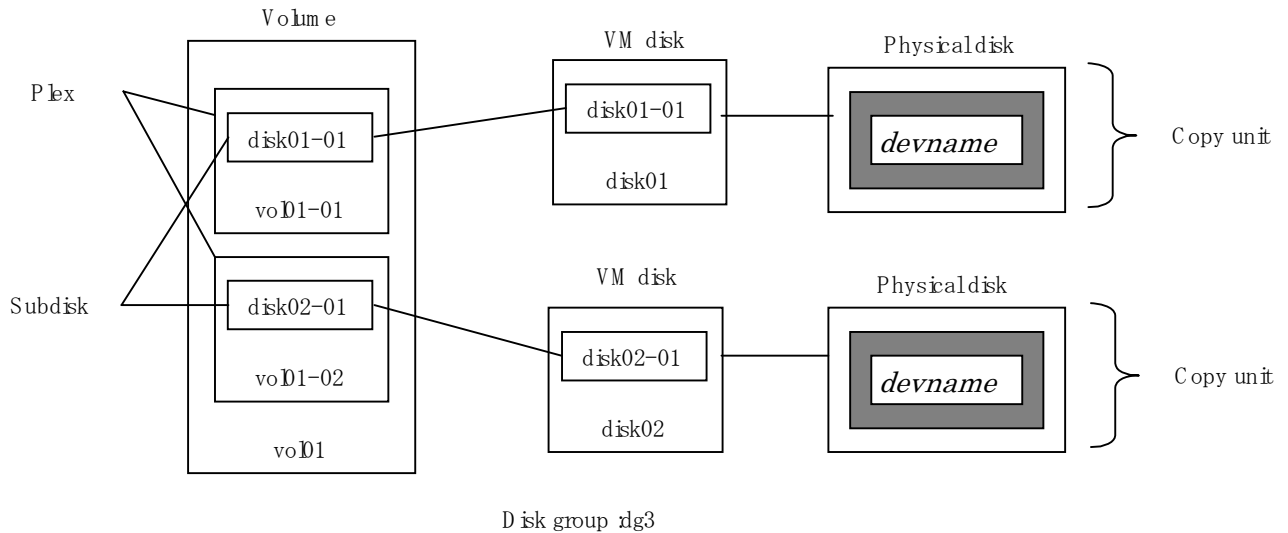
- Physical disks that constitute a disk group

For information on the replication operation of VxVM volumes, see "VxVM Volume Operation".

1.5.3.1 Physical disks that constitute a disk group

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

Figure: Example of VxVM configuration



Since AdvancedCopy Manager copies an entire VM disk, it copies not only the shared area that constitutes a subdisk but also the occupied area containing VxVM-internal configuration information.

Therefore, the integrity of VxVM configuration information must be maintained during copying when replication is performed.

Specify a device name as follows:

`/dev/vx/dmp/c#t#d#`



Operation under the naming convention based on the enclosure is not supported.



For operations in a cluster configuration, the device name (`/dev/(r) dsk/c#t#d#`) of the physical disks that comprise a disk group must be the same at all servers comprising the cluster, and all the ETERNUS disks indicated by the device name must be the same.

Chapter 2 Starting and Stopping Daemons

This chapter provides information on how to activate and stop the service daemons required by AdvancedCopy Manager.

For more information on how to activate and stop the daemons described, see "Start/Stop Daemon commands".

2.1 Daemons that Configure AdvancedCopy Manager

The Agent of AdvancedCopy Manager is configured with the following daemons:

Table: Types of daemons that are required by the Agent of AdvancedCopy Manager

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

In cluster operation, the following daemon runs in the Storage Server:

Table: Type of daemon running Storage Server transactions

Daemon type	Description
Communication daemon	Daemon handling communication with Storage Management Servers and other Storage Servers



Note

The following daemon cannot run in a Storage Server participating in a cluster operation:

- Communication daemon (local daemon)

2.2 Starting Daemons

Normally, AdvancedCopy Manager daemons are automatically activated as part of the server startup routine. However, if activation fails or a daemon is stopped, activate the daemon on each server by executing the following command at a command prompt:

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



Note

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

In cluster operation, the cluster system activates the daemon automatically.



Note

To start a daemon independently of some AdvancedCopy Manager transactions (instance) for any reason, follow the steps described below.

You must start the daemon for each AdvancedCopy Manager transaction (including storage local transactions). Take the following steps for the respective transactions.

- 1) Execute the following command on the command line:

```
# ps -ef | grep stgxfws
```

Confirm that command execution does not display:

```
/opt/FJSVswstf/bin/stgxfws logical-node-name
```

If the above message is displayed, the daemon has already started.

- 2) For environment variable SWSTGNODE, specify the name of the logical node on which the transaction is to run.
- 3) Start the communication daemon.
For the startup command, refer to "Start/Stop Daemon Commands".

2.3 Stopping Daemons

.Normally, daemons are automatically stopped when the system stops.

To stop a daemon on a Storage Management Server or a Storage Server, execute the following command on the server:

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



Note

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

In cluster operation, stop the daemon from the cluster system. For the stop procedure, see the relevant cluster software manual.



Note

To stop a daemon independently of certain AdvancedCopy Manager transactions for any reason, follow the steps described below.

You must stop the daemon in each AdvancedCopy Manager transaction (including storage local transactions). Take the following steps for the respective transactions.

- 1) Execute the following command on the command line:

```
# ps -ef | grep stgxfws
```

Confirm that command execution displays:

```
/opt/FJSVswstf/bin/stgxfws logical-node-name
```

If the message above is not displayed, the daemon has already stopped.

- 2) For the environment variable SWSTGNODE, specify the name of the logical node on which the transaction is to run.
- 3) Stop the communication daemon.
For the stop command, see "Start/Stop Daemon Commands".

Chapter 3 Security Operation Using Authentication Feature

AdvancedCopy Manager provides access control for users who need to perform any of the backup management or replication management functions.

To achieve this, a security feature is enabled which specifies the range of functionality for each users

For details of this, see "Security Operation Using Authentication Feature" in "ETERNUS SF AdvancedCopy Manager Operator's Guide the following operating environments:

- Windows edition
- Solaris edition
- Linux edition.

Chapter 4 Backup Operations

This chapter provides details of the backup operations performed by AdvancedCopy Manager.

The backup operation using the HP-UX version of AdvancedCopy Manager's Agent as described in this chapter can be used with AdvancedCopy Manager's Manager component running under other OS environments such as Windows, Solaris or Linux.

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

4.1 Overview

This chapter describes the backup and restore functions of AdvancedCopy Manager.

AdvancedCopy Manager performs high-speed backup between volumes, regardless of the volume capacities, using the advanced copy function (i. e., the OPC and EC functions) of a disk array unit.

A Backup that uses either the One Point Copy (OPC) or Remote One Point Copy (ROPC) function of a Fujitsu ETERNUS storage system is referred to as a "Snapshot Fast Backup".

A Backup that uses the Equivalent Copy (EC) or Remote Equivalent Copy (REC) function is referred to as a "Synchronized High speed Backup".

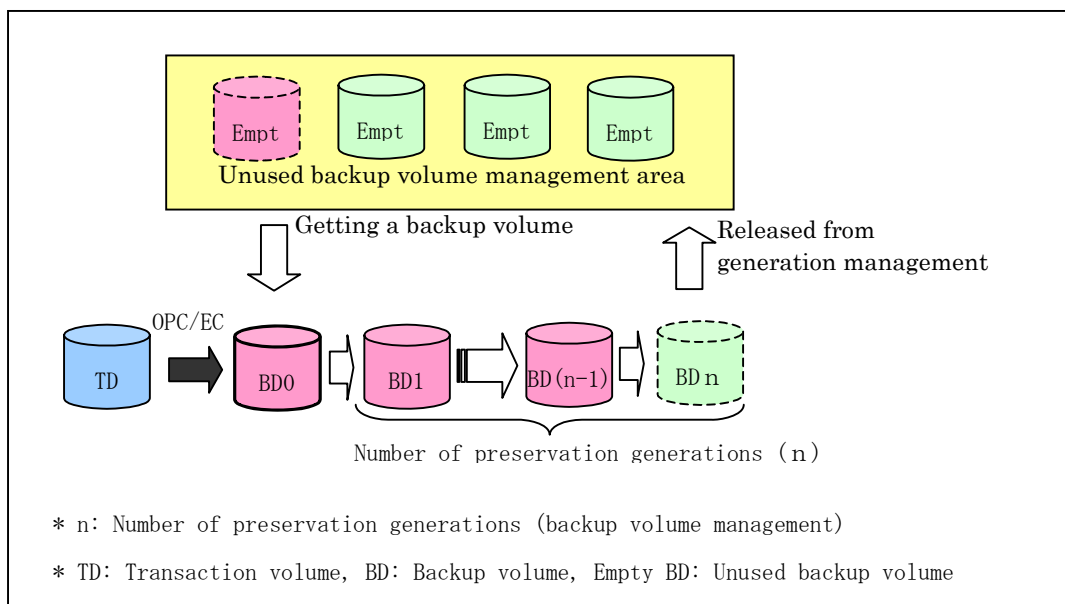
Automatic snapshot backup can be scheduled without first estimating the time required because it does not use synchronous backup processing. However, when the generation number is "1", a state may occur in which there is no backup data, because a snapshot backup deletes the history prior to copy processing.

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 may be "1". However, synchronized high-speed backup requires one more backup volume than is required by snapshot fast backup, also requires synchronized backup processing before the backup execution command is entered. Therefore, execution of automatic synchronized high-speed backup must be scheduled so that the time required for synchronized processing is estimated in advance. Synchronized backup processing is then started based on the estimated time, and the backup execution command is entered only after equivalency maintenance status has been achieved.

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

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

Figure: Managing a backup volume



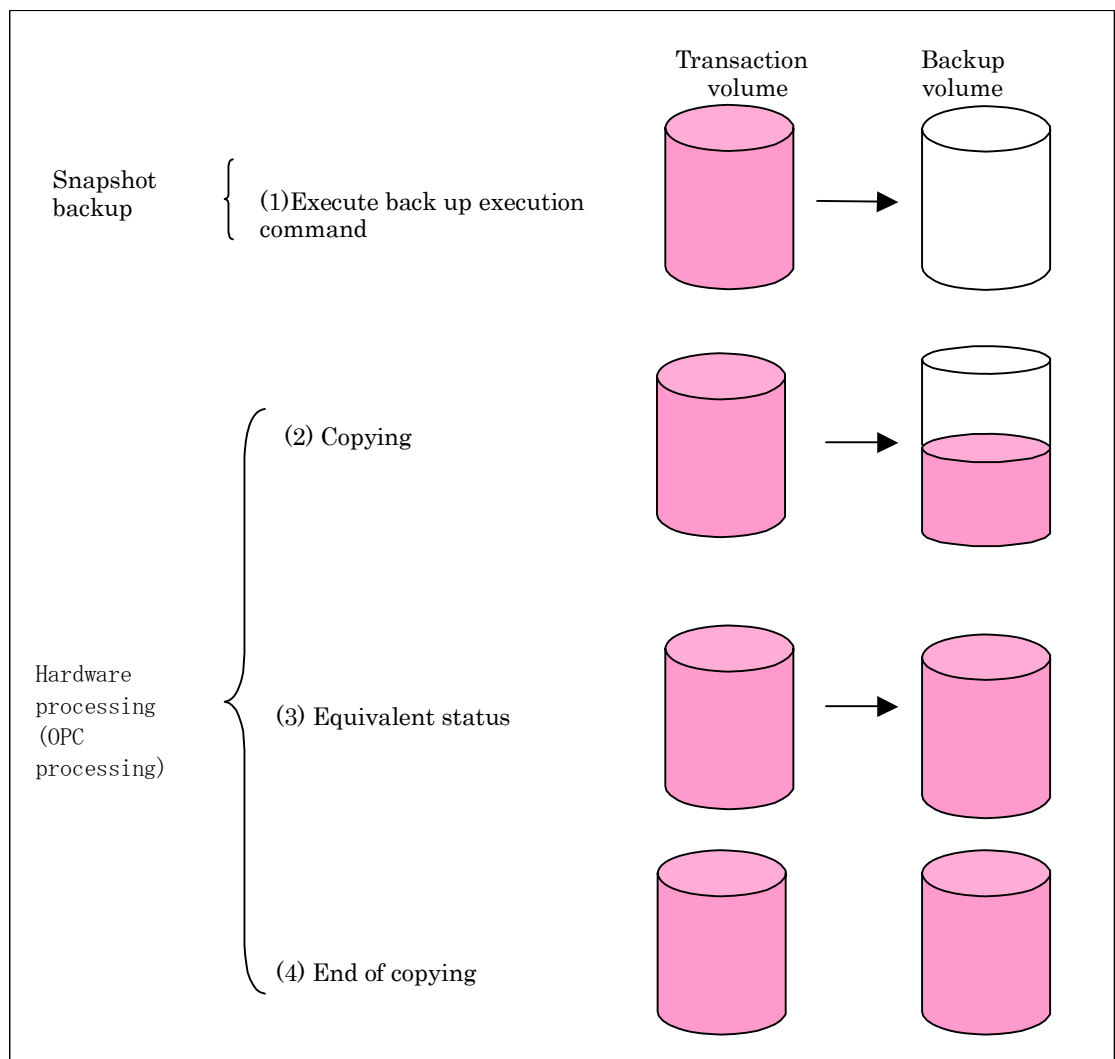
4.1.1 Snapshot backup processing

The One Point Copy (OPC) and Remote One Point Copy (ROPC) functions of the Fujitsu ETERNUS storage system are used to copy data from a transaction volume to an unused backup volume.

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

1. Execute the backup execution command (swstbackup) ((1) in the following figure). The following processing will be performed:
 - If backup volumes under generation management exist for all the generations, the oldest backup volume will be assumed to be unused.
 - An unused backup volume will be selected from the unused backup volume management area.
2. Data on the transaction volume will be copied to the selected unused backup volume ((2) to (4) in the following figure).
3. Backup carrier information will be set.

Figure: 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 the Fujitsu ETERNUS storage system.



Note

Using the QuickOPC function in the snapshot high speed backup, physically copies only the differences in the data from the start of the previous backup to create a complete snapshot image, which greatly reduces the time taken for physical copying.

For information on the QuickOPC function, see "Backup Operation by the QuickOPC Function."



Note

- If all the generations of backup volumes are managed, and no backup volume can be obtained from the unused region in the backup volume management area, backup copies cannot be created.
- Backup using ROPC is limited to disk array devices within the same Storage Server that support ROPC.

4.1.1.1 Snapshot backup processing of a Volume Group

When the transaction volume is a volume group, copy processing by the OPC and the ROPC functions of disk array equipment is performed for the entire physical disk comprising the volume group. Therefore, when two or more logical volumes exist in a volume group, the data contained in all logical volumes is backed up (refer to the following figure).

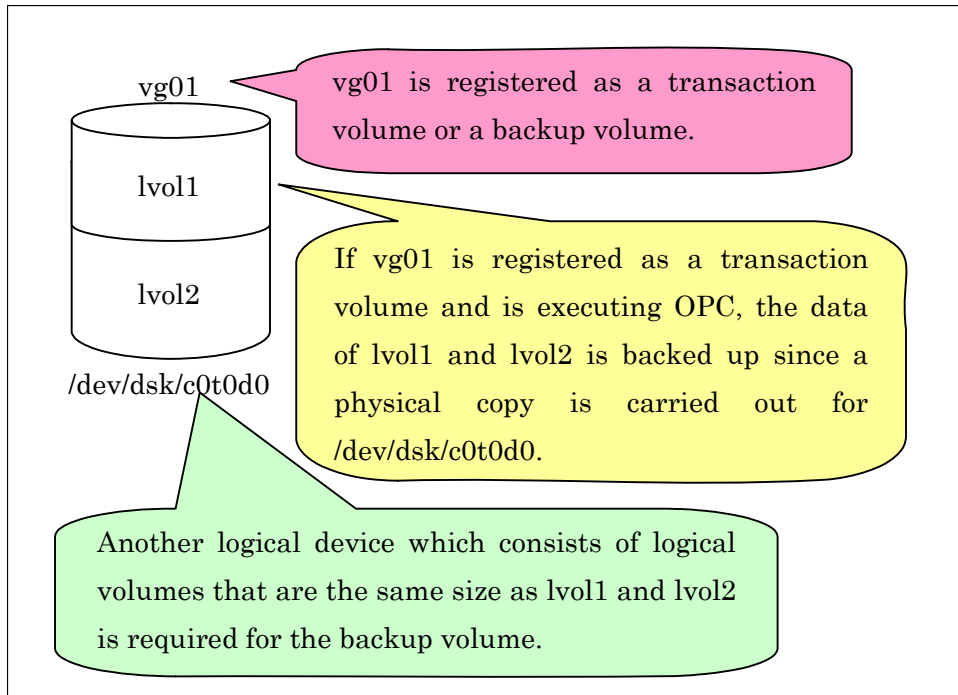


Note

When using a volume group for backup operation, care needs to be taken with the following points:

- When the transaction source volume is a volume group, construct a logical disk whose size matches that of the transaction volume, in a volume group different from the transaction group, and register this volume group whose physical disk size is the same as the transaction volume as the backup volume.
- A "device map file" is created to describe the pairing of matching transaction and backup volumes. A "Device map" is mandatory if you are backing up using a volume group unit.
- In backup processing, since the LVM configuration information on the backup volume will be rewritten for the LVM information on the transaction volume, LVM configuration information is restored in the processing after backup. For more information, see "Pre-processing and Post-processing of Backup and Restoration".
- When backing up a volume group using an account with Administrator authority for backup, modify the scripts used for backup pre-processing and post-processing to unmount/mount all logical volumes in the volume group. For more information, see "Pre-processing and Post-processing of Backup and Restoration".
- When two or more physical disks exist in the volume group which is registered with a transaction volume, the backup operation will result in an error.

Figure: Transaction volume and backup volume for two or more logical disks



4.1.2 Synchronized high-speed backup processing

The Equivalent Copy (EC) or Remote Equivalent Copy (REC) function of the Fujitsu ETERNUS storage system of disk array units is used to copy data from a transaction volume to an unused backup volume.

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

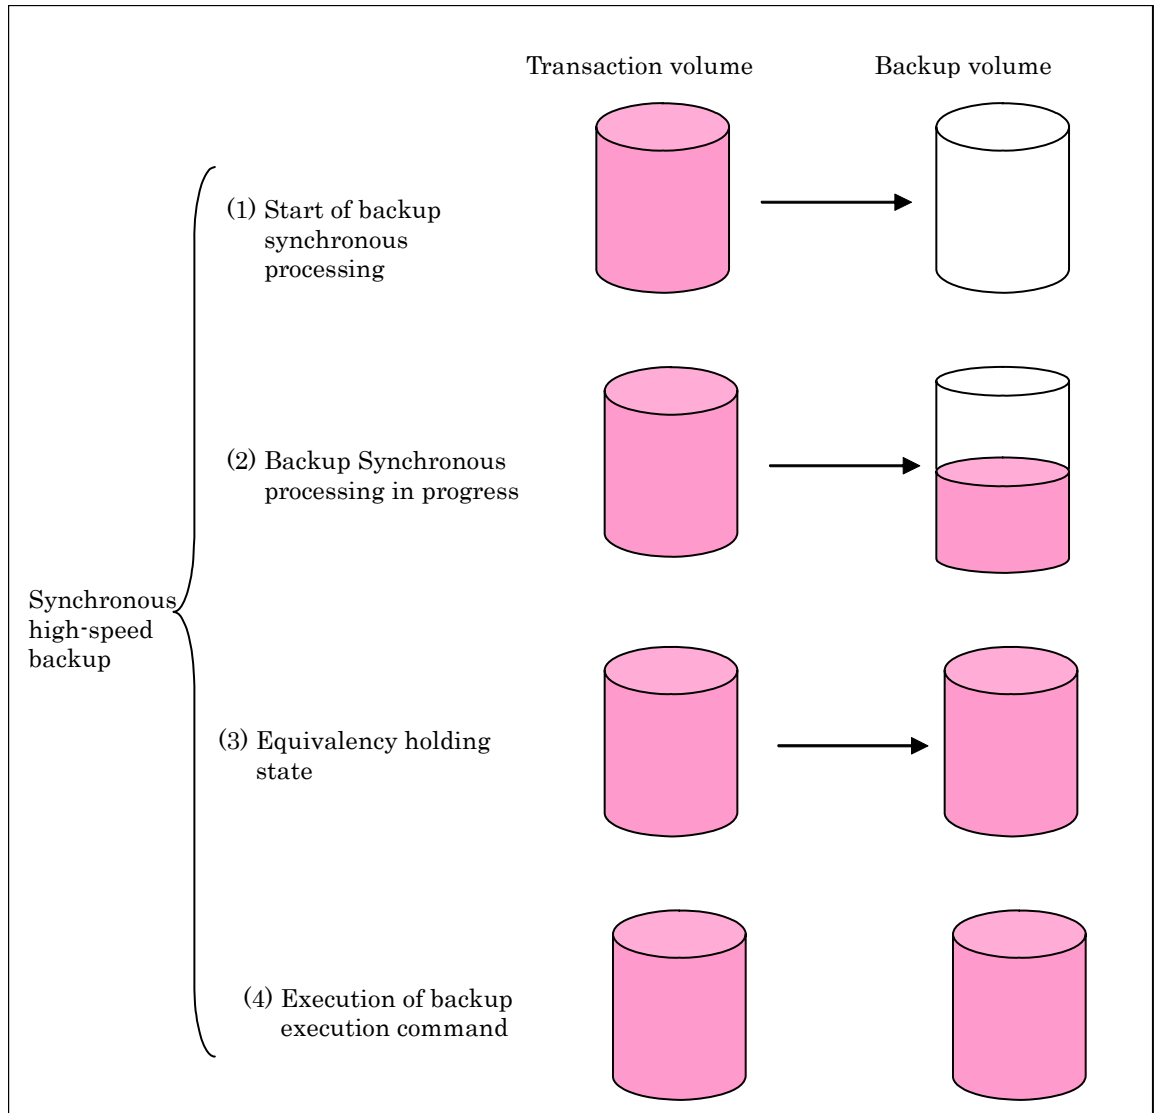
1. Execute the Backup synchronous processing start command (step (1) in the following figure). An unused backup volume with the same capacity as the transaction volume is selected from the unused backup management area, and the synchronized backup processing starts.
2. The transaction and backup volumes are both placed in an equivalency state. From this point on, equivalence between the transaction volume and backup volume is maintained. This status is referred to as equivalency maintenance status (step (3) in the following figure).
3. Execute the backup execution command (step (4) in the following figure). The backup synchronization processing stops and the backup history information is set.
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, and is overwritten with the new data.



Note

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

Figure: 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, see "Backup Process That Uses the Suspend/Resume Function"



Note

- If all the generations of backup volumes are being managed, and no backup volume can be obtained from the unused region in the backup volume management area, backup copies cannot be created.
- The backup function using REC is limited to disk array devices within the same Storage Server that support REC/ROPC.

4.1.2.1 Synchronized high-speed backup processing of Volume Group

When the transaction volume is a volume group, copy processing by the EC and REC functions is performed to the physical entire disk comprising the volume group. Therefore, when two or more logical volumes exist in a volume group, the data contained in all logical volumes is backed up.



Note

When using a volume group for backup operation, care needs to be taken with the following points:

- When the transaction source volume is a volume group, construct a logical disk whose size matches that of the transaction volume, in a volume group different from the transaction group, and register this volume group whose physical disk size is the same as the transaction volume as the backup volume.
- A "device map file" is created to describe the pairing of matching transaction and backup volumes. A "Device map" is mandatory if you are backing up using a volume group unit.
- In backup processing, since the LVM configuration information on the backup volume will be rewritten for the LVM information on the transaction volume, LVM configuration information is restored in the processing after backup. For more information, see "Pre-processing and Post-processing of Backup and Restoration".
- When backing up a volume group using an account with Administrator authority for backup, modify the scripts used for backup pre-processing and post-processing to unmount/mount all logical volumes in the volume group. For more information, see "Pre-processing and Post-processing of Backup and Restoration".
- When two or more physical disks exist in the volume group which is registered with a transaction volume, the backup operation will result in an error.
- In the case where the backup volume is a volume group, do not execute the `vgchange (1M)` command while executing synchronous processing (EC/REC). If the backup volume becomes deactivated during the execution of synchronous processing, it will not be possible to reactivate it. In addition, it will not be possible to stop synchronous processing, the backup execution command, the backup synchronous processing cancel command, etc. If accidentally deactivated, use `ETERNUSmgr`, `GRmgr`, or similar, to forcibly stop synchronous processing, and then activate the backup volume.

4.1.3 Restoring data

This section describes how data is restored in AdvancedCopy Manager.

The restore 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 the disk array unit 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.

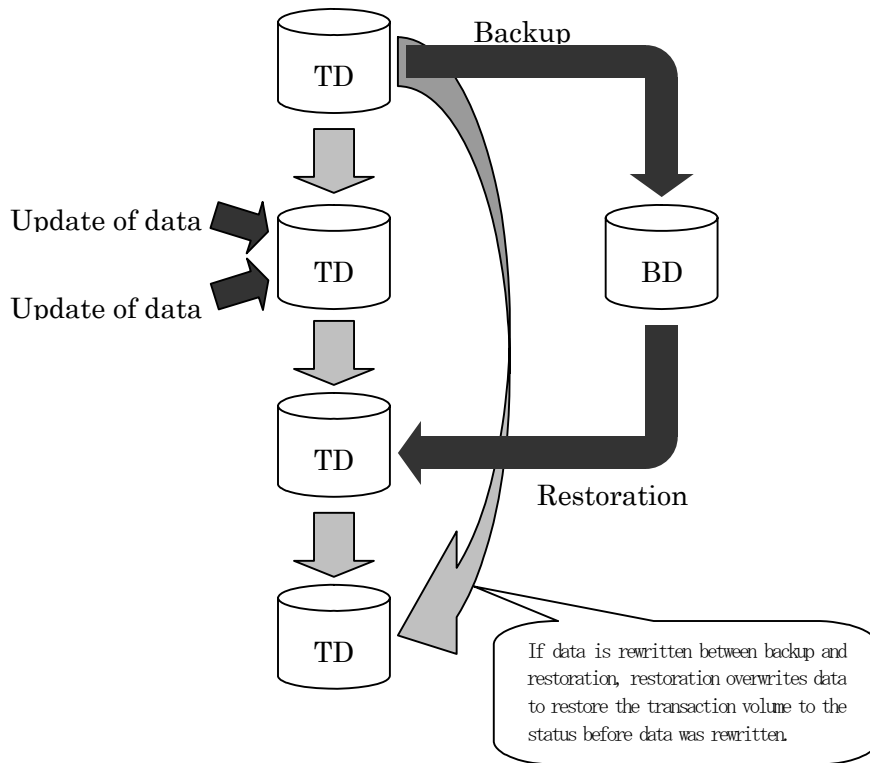
Restoration processing is performed to the transaction volume that was the backup source. However, it is also possible to change the restore destination.



Note

If the contents of a transaction volume are rewritten or modified in the period between the latest backup of that volume and a subsequent restore operation, the accuracy of any data rewritten or modified cannot be guaranteed.

Figure: Notes on restore operations in AdvancedCopy Manager



TD: Transaction volume
BD: Backup volume

4.1.3.1 Restoration processing of Volume Group

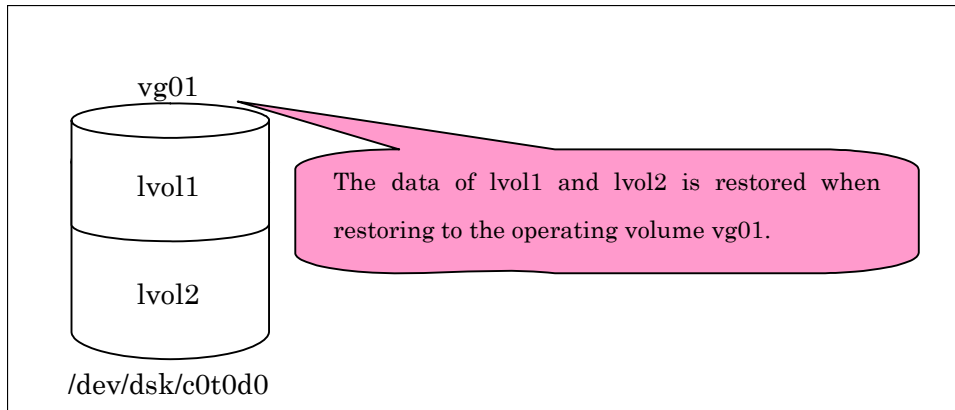
When a transaction volume is a volume group, copy processing by the OPC and ROPC functions of a disk array is performed to the entire physical disk comprising the volume group. Therefore, when two or more logical volumes exist in a volume group, the data contained in all logical volumes is backed up.



When using a volume group for backup operation, care should be taken with the following point:

When restoring a volume group, using an account with Administrator authority for backup, correct the script used for backup pre-processing and post-processing to unmount/mount all logical volumes. For more information, see "Pre-processing and Post-processing of Backup and Restoration".

Figure: Restoration for two or more logical disks



4.2 Configuration

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:

- Selecting a server on which the backup operation should be performed
- Selecting a backup source target
- Configuring a backup policy
- Preparing a backup destination volume
- Automating the backup operation

4.2.1 Selecting a server on which the backup operation should be performed

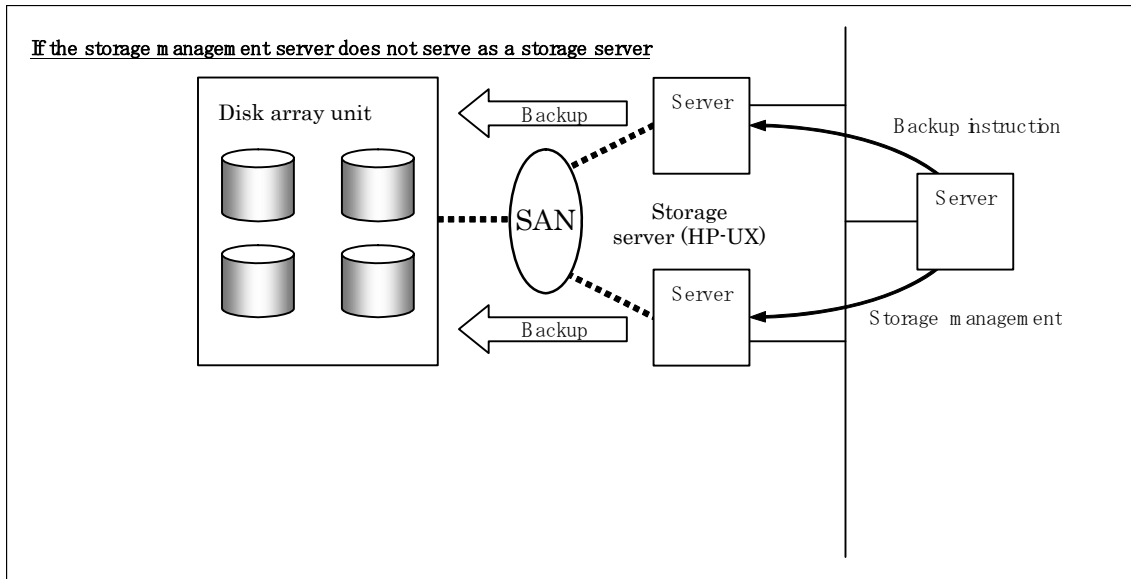
The backup operation can be performed on one of the following servers:

- Storage Management Server
Multiple Storage Servers are always managed and operated from a single participating server hosting the management functionality.
A Storage Management Server can also serve as a Storage Server.
Note that AdvancedCopy Manager running on a Storage Management Server must run in either a Windows, Solaris or Linux OS environment.
- Storage Server
The HP-UX version of AdvancedCopy Manager can also be used to operate a Storage Server.



The locale of the internal code system (i.e., the code system specified when installing AdvancedCopy Manager on the Storage Management Server) used by AdvancedCopy Manager must be installed on all Storage Servers.

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



4.2.2 Selecting a backup target

The backup or restoration unit of the HP-UX AdvancedCopy Manager is a physical disk (LU: Logical Unit). However, in the case of a physical disk which constitutes the LVM (Logical Volume Manager), it is managed and operated as a volume group. This disk is called a transaction volume (source). A transaction volume to be backed up must be placed on the Fujitsu ETERNUS storage system.

For information on setting up a transaction volume, see "Setting the operation type for a device".



Note

The following should be noted when using a volume group as operating volume:

- Do not register two or more physical disks into one volume group.



Note

Do not back up the systems disk, or the disk on which AdvancedCopy Manager is installed.



Note

Refer to the "General notes" for points to consider when deciding on a candidate volume for backup. An example of a point to consider is the following:

- Exclusions from Backup Targets or Replication Targets

4.2.3 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 the generation management. If a system failure or any other failure occurs while the backup is running, there may not be as many generations of backup data as required. If this happens, you should immediately perform backup again.

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 complete. Therefore, the required number of backup volumes is the number of backup generations plus one (1).

Interval days

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

A notification that backup is overdue is displayed if a period greater than the specified interval has passed since the last backup date.

Note that backup is not performed automatically in this instance, even if you have specified the interval days between backups.

For information on setting a backup policy, see "Configuring backup policies".

4.2.4 Preparing a backup volume

This section describes how to prepare the backup volume.

In AdvancedCopy Manager, the disk in which the live transaction data is stored is called a transaction volume. A backup volume for this transaction volume must also be placed on the Fujitsu ETERNUS storage system.

The backup volume size must be exactly the same number of bytes as the transaction volume. The number of backup volumes used for both snapshot type high speed backup and synchronous high speed backup operations vary as shown below:

Table: Number of backup volumes required

Backup	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 in advance:

Table: Number of backup volumes required for a newly registered transaction volume

Backup	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

Refer to "Setting the device operation type" for the backup volume setting method.



Note

The following should be noted when using a volume group as a backup volume:

- Do not register two or more physical disks into one volume group.
- Create a logical volume so that it does not span two or more physical disks.
- The backup volume must be the same physical size as transaction volume. In addition, if the backup and transaction volumes are volume groups, they must be in different volume groups. However, these volume groups must have the same configuration.



Note

When selecting backup volumes, do not select systems disks where the operating system resides or where AdvancedCopy Manager is installed or invoked.



Note

Note the following points when performing backup to a backup volume in a cabinet that is not the same as the cabinet containing the transaction volume.

- Disk array devices must support the REC or ROPC function.
- Both cabinets must be connected to an FCRA (FC Remote Adapter).
- Since data flows only from the Initiator side to the Target side through an FCRA connection, FCRA connections must be established in minimum sets of two. In a backup operation involving a disk array device that does not support the ROPC function (and only the REC function can be used), backing up data to a backup volume in a cabinet containing no transaction volume is impossible since the backup operation uses the ROPC function during restoration.



Point

If the number of backup volumes that can be prepared in synchronized high-speed backup is not greater than the number of preservation generations, the backup operation can still be carried out by saving the backup volumes to secondary media, such as tapes.

To achieve 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.

If the history information is deleted, 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. In this event, the data must be restored directly from the secondary media.

4.2.5 Automating backup operations

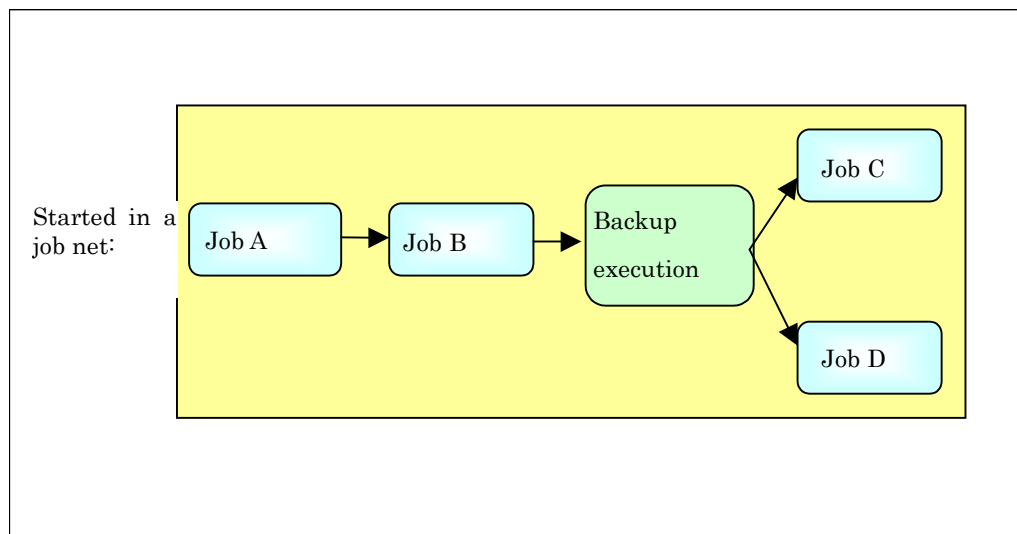
The backup operation of AdvancedCopy Manager can be automated by using SystemWalker/OperationMGR, etc.

Automatic snapshot fast backup

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

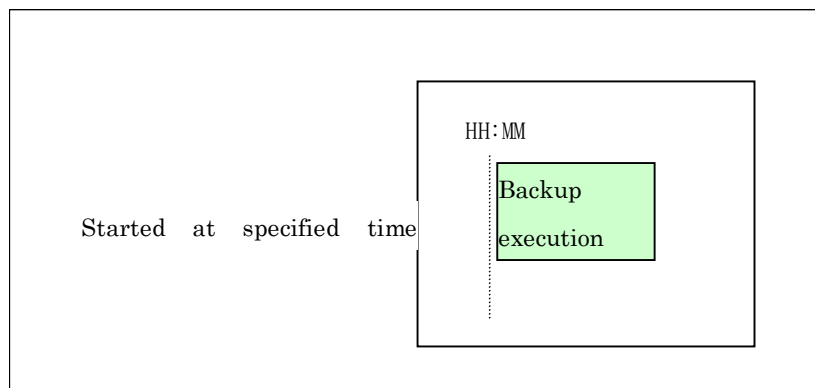
- Start the backup command in a job net of a batch job every time that a backup should be collected.

Figure: Start in a job net



- Start the backup command at the specified time.

Figure: Start at a specified time



Automatic synchronized high-speed backup

For synchronized high speed backup, the transaction and backup volumes must have attained equivalency maintenance status as at the scheduled backup time.

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

1. 1. Estimate the expected time required for the synchronized backup processing.
2. 2. Execute the synchronized backup processing start command. To do this, specify a start 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 attained equivalency maintenance status when backup is executed.

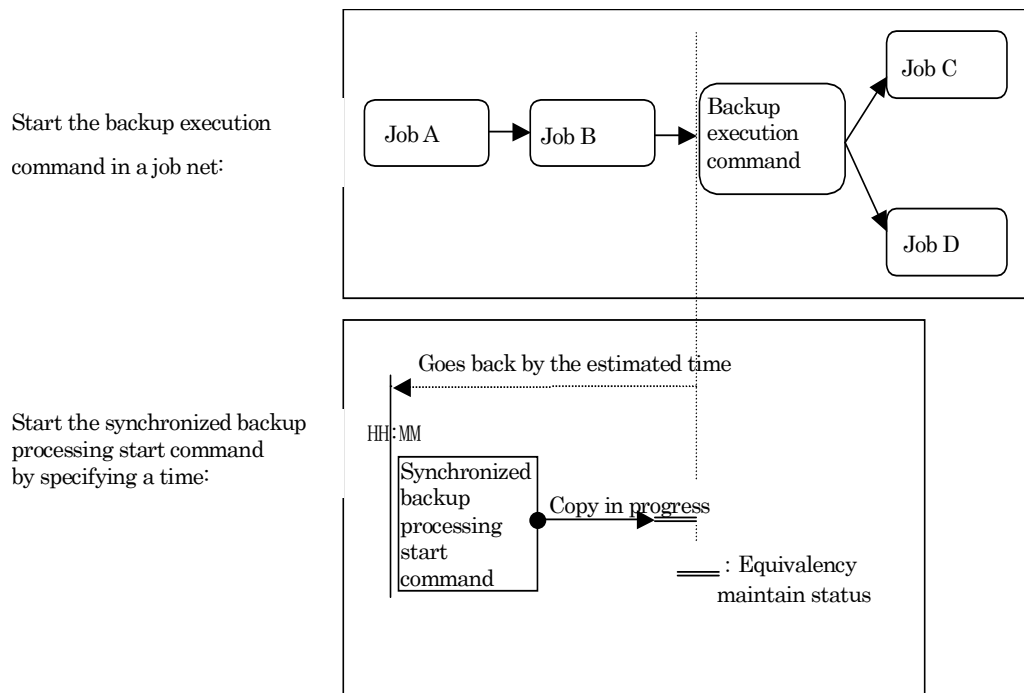
The time between the start of synchronized backup processing and the transition to equivalency maintenance status depends on the capacity of the volume to be copied.

Starting the backup execution command

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

- Start the backup execution command in a job net of a batch job at the same time that the backup data is to be collected.

Figure: Start in a job net

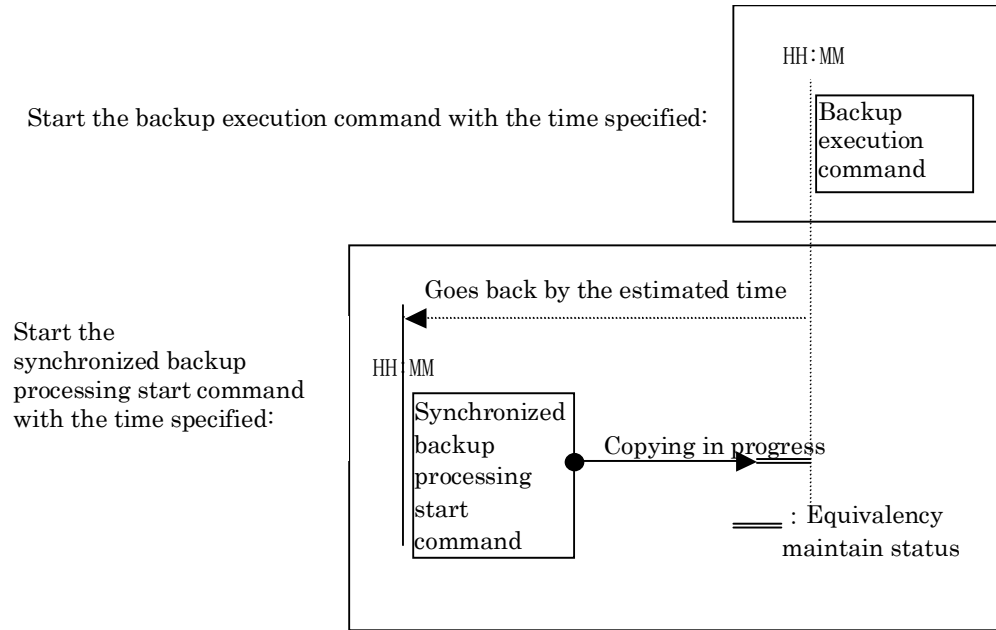


Note

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.

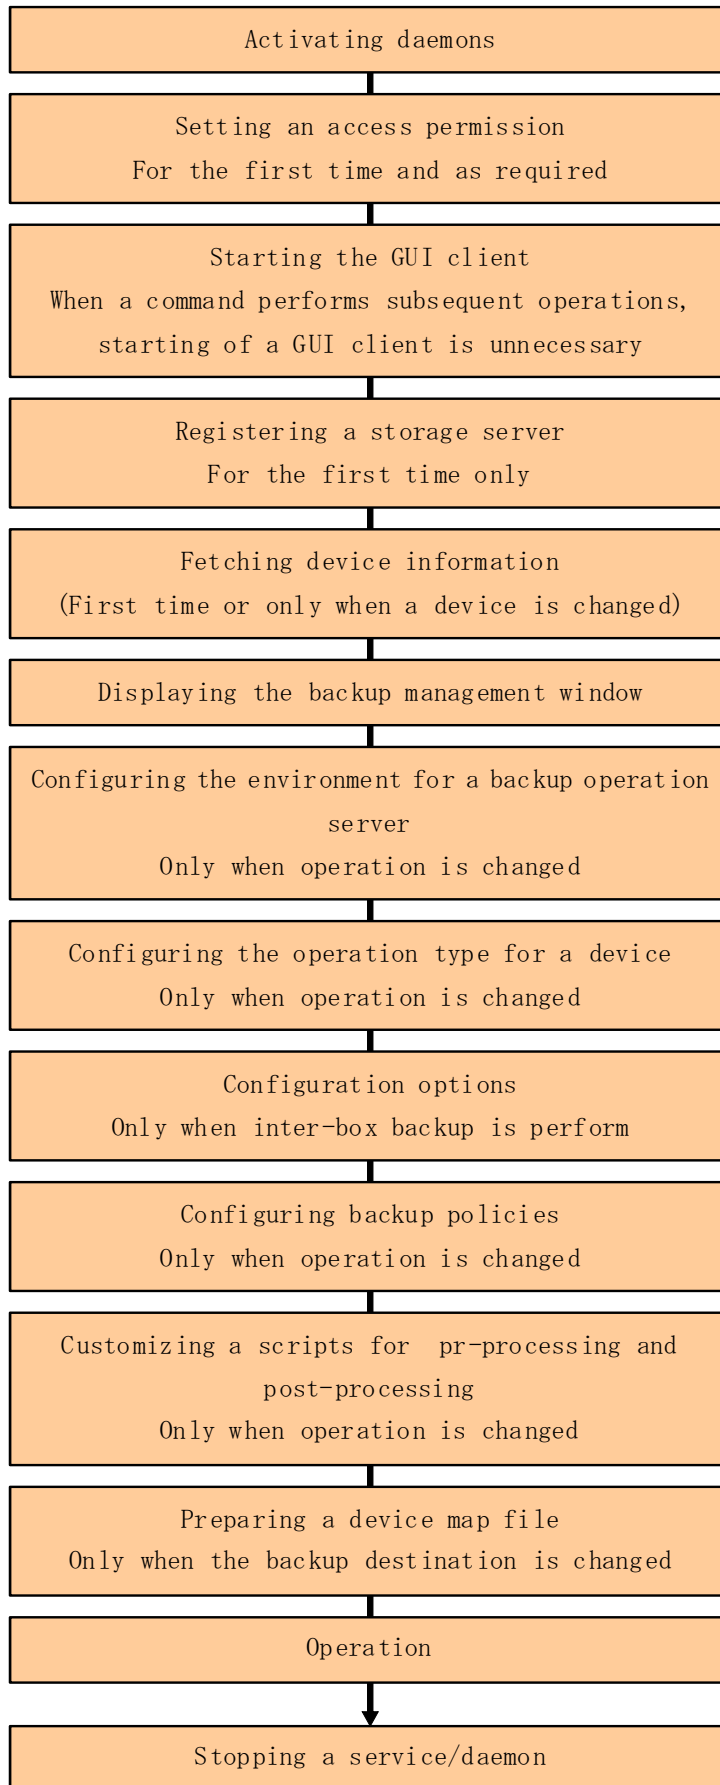
Figure: Start at a specified time



4.3 Flowchart of Backup Operation

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

Figure: Flow of backup operation



4.4 Preparing to Start the Backup

The following preparations are required before the backup operation is started:

4.4.1 Activating daemons

Before starting a backup operation, daemons on the Storage Management Server and Storage Servers must be activated. These daemons are normally automatically activated on system startup.

If startup fails for any reason, or if a daemon is stopped once it has been started, it is necessary to restart a daemon for each server.

Please refer to "Starting and Stopping Daemons" for additional information.

4.4.2 Setting an access permission

When performing backup from a backup management screen, access permission must be set up for each function of the backup management screen. For details on how to set permissions, see "Security Operation Using Authentication Feature" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide".

When performing backup using a command at a command prompt, the setup of access permission is not required.

4.4.3 Starting the GUI client

Start the GUI client. For details of this, refer to "Starting the AdvancedCopy Manager Workbench" in the "ETERNUS SF AdvancedCopy Manager User's Guide". If using commands only, this operation is not required.

4.4.4 Registering a Storage Server

The Storage Server to be managed is registered from the Storage Management Server. When the Storage Server is operated as a Storage Management Server the server does not need to be registered.

The process of registering the Storage Server can be carried out using GUI operations or commands. For information on the operation method using GUI, refer to "Register a New Storage Server" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For information on the operation method using commands, refer to Server information addition command (stgxfwcmaddsrv).

4.4.5 Fetching device information on a Storage Server

This process can be carried out using GUI operations or commands. For information on the operation method using GUI, refer to "Manage Devices" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For information on the operation method using commands, refer to the Device information collection/reflection command (stgxfwcmsetdev) in the "ETERNUS SF AdvancedCopy Manager User's Guide" for the OS of the Storage Management server.



Note

As part of replication management, device information on a Storage Server must first be stored in the repository.



Note

There may not be sufficient area available for updating the repository if information is to be fetched for devices under two or more Storage Servers. Check if the area required when updating the repository is available before fetching the device information. If there is insufficient area, allocate the required area, then fetch the device information.

For the procedure, refer to the version of the following manual that applies to the operating system of the Storage Management Server:

- "ETERNUS SF AdvancedCopy Manager Operator's Guide", step 2 of "Action required when access to repository fails"



Note

The time required to add a new device is in proportion to the number of devices defined on the Storage Server. For this reason, if a large number of devices have been defined, perform this process only while the CPU load on the system is low. As a guide, each operation takes about 0.5 seconds per device (i. e., partition) in a state without load.

4.4.6 Displaying the backup management window

Configuring the environment for a server can be carried out using either GUI operations or commands. For information on the operation method using GUI, refer to "Setting up Storage server configuration information" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For information on the operation method using commands, refer to Storage server configuration information display command (swstsvrset). When performing backup using a command only, the Storage Server configuration information setting command must be executed at the time that operations start.

If the Storage Management Server is combined with the Storage Server, these environment settings must also be carried out for the Storage Management Server.

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



Point

If the Storage Management Server is combined with the Storage Server, these environment settings are also required for the Storage Management Server.

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

4.4.7 Configuring the environment for a backup operation server

Set up the volume type for a device that is connected to the Storage Server.

The setting of the operation type is explained below.

4.4.8 Configuring the operation type for a device

Set up the volume type for a device that is connected to the Storage Server.

The setting of the operation type is explained below.

- Defining the volume to be backed up as a transaction volume.
- Preparing backup volumes. Prepare as many volumes, each with the same size as the transaction volume, as required for the backup type, then define them as the backup volumes:

Table: Number of backup volumes required

Backup	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 GUI operations or commands. For information on the operation method using GUI, refer to "Configure a Device for Use with Backup" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For information on the operation method using commands, refer to Device information setting command (swstdevinforeset)



Note

A logical volume of LVM cannot be registered as transaction volume or backup volume. Register it in a volume group.



Note



Note

In order to use a volume group as a transaction volume, it is necessary to create a logical disk as follows:

- Do not register two or more physical disks into one volume group.
- Create a logical disk so that one logical disk does not span two or more physical disks.



Note

Do not set the disk where the system is stored or the disk where AdvancedCopy Manager is installed as a backup target or as a backup volume.



Note

To change the configuration of a disk registered as the transaction volume or backup volume, take the following steps:

1. Delete the registered configuration, and then change the configuration.
2. Attempt to fetch the device information on the Storage Server.
3. Register the information by executing the device information configuration command (swstdevinforeset).



Note

A mounted device cannot be configured as a backup volume. (It is only possible to configure it as a transaction volume.)



You can check the volumes that have been registered as follows:

- Checking in a GUI client
You can view information on a transaction volume in the transaction volume list view, and information on a backup volume can be referred to from the GUI client. For more information, see the "ETERNUS SF AdvancedCopy Manager User's Guide".
- Checking by executing a command
You can view information on a registered transaction volume or backup volume using device usage status display command.



When the registered contents are checked by the GUI client or by the command, the value displayed on the "Size" column is the physical disk size. When a volume group is registered, this column also displays the physical disk size of the volume group.

4.4.9 Configuration options

It will be necessary to set a number of options when performing backup to a backup volume in a cabinet that is different from the cabinet containing the transaction volume. Backup operations may not be able to continue if the options setup file is changed after operations start. Therefore, do not change the options setup file after operations start. If the options setup file is changed, the backup policy must also be reset.



- Disk array devices must support the REC and ROPC functions.
- Both cabinets must be connected to an FCRA (FC Remote Adapter).
Since data flows only from the Initiator side to the Target side through an FCRA connection, FCRA connections must be established in sets of two or more.
- In a backup operation involving a disk array device that does not support the ROPC function, only the REC function can be used. In this case, you cannot backup data to a backup volume in a cabinet that does not contain a transaction volume, since the backup operation uses the ROPC function during restore.

Create the following files to set the options:

- For normal operation (non-cluster)
`/etc/opt/FJSVswsts/data/DEFAULT/check.ini`
 - For cluster operation
`/etc/opt/FJSVswsts/logical node name(*1)/data/DEFAULT/check.ini`
- (*1) Logical node name of AdvancedCopy Manager specified during setting up of a cluster.

Create code in the files as follows:

```
[check]
RemoteCopy=Yes
```

4.4.10 Configuring backup policies

In the case of all transaction volume registered on a Storage Server, set the following backup policy:

- The number of preservation generations
This refers to how many generations of backup data should be retained.
- The interval days
This refers to 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 GUI client.

This process can be carried out using GUI operations or commands. For information on the operation method using GUI, refer to "Set Backup Policy" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For information on the operation method using commands, refer to Backup policy setting command (swstbkpolset)



Note

- Even if the number of interval days has been specified, AdvancedCopy Manager does not automatically invoke 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, see "Preparing a backup volume".
- Even though the number of backup volumes required for the backup may not have been registered beforehand when backup policies for synchronized high-speed backup were set, the number can be set by default provided the number of backup volumes that are required for snapshot fast backup have already been registered. In this case however, it may not be possible to execute synchronized high-speed backup.



Point

You can check the volumes that have been defined in the following ways:

- Checking in a GUI client
To check backup policy, refer to "Set Backup Policy" in the "ETERNUS SF AdvancedCopy Manager User's Guide".
- Checking using a command
You can view information on a defined backup policy using the backup policy display command (swstbkpoldisp).

4.4.11 Customizing Pre-processing and Post-processing scripts

In AdvancedCopy Manager, backup or restore processing must be performed while transaction volumes are unmounted. Therefore, when backup or restore is executed, mount and unmount processing of the transaction volume is done by using the pre-processing and post-processing scripts.

Backup or restore processing is not executed in cases when the transaction volume cannot be unmounted.

For details of the pre-processing and post-processing scripts, refer to "Pre-processing and Post-processing of Backup and Restoration."

When the operation corresponds to any of the following cases, it is necessary to customize the scripts that are used for pre-processing and post-processing of backup.

- When the backup target is a volume group.
- For operational reasons, avoid unmount/mount processing of a transaction volume into which the file system has been built.
- To describe special processing in a pre-processing and post-processing script.

For information about the customization method, see "Pre-processing and Post-processing of Backup and Restoration".



When AdvancedCopy Manager is upgraded

The scripts may be updated in an upgraded version of AdvancedCopy Manager. Therefore, you should customize the scripts again after upgrading rather than continuing to use the script from the previous version.

4.4.12 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 registered backup volumes. If a specified backup volume is to be used to back up a given transaction volume, a "Device map file" (which is a file defining the correspondence between transaction and backup volumes) must be created.

Create a device map file anywhere on the Storage Server that performs backup. If you specify this file while performing backup, an operation that recognizes the backup destination is enabled.



To manage multiple generations of backups, 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, or
- The backup volume used in the history that was deleted by this backup must be specified.

You should ensure that you configure and use the device map file appropriately, depending on how you are using the backup system with two or more backup volumes.

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

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

4.4.12.1 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/dsk/c1t1d1             /dev/dsk/c1t1d20
/dev/dsk/c1t1d1             /dev/dsk/c1t1d20
/dev/dsk/c1t1d1             /dev/dsk/c1t1d20

# Volume Group
/dev/vg01                   /dev/vg02 # A comment can also be entered in this section.
/dev/vg03                   /dev/vg04
:
:
```

These are the rules for creating a device map file.

- Describe a transaction volume name and a corresponding target backup volume name in one line. Separate the transaction volume name and the output destination backup volume name with one or more space or tab characters.
- One or more space or tab characters may be included between the beginning of a line and a transaction volume name or between the end of an output destination 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.
- If a line starts with the symbol "#" then that entire line is a comment.
- You cannot specify more than one target backup volume for one working disk in one device map file. If you do, only the information in the line first found will be valid. While the device map file will still be read, the lines after the first line will be ignored.
- A device map file may include descriptions of transaction volumes other than the processing target. (See the description example provided in the beginning of this chapter.)

4.5 Operation

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

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

AdvancedCopy Manager has two kinds of backup operation:

- Snapshot fast backup
- Synchronous high-speed backup

4.5.1 Performing snapshot high-speed backup

Snapshot high-speed backup can be carried out using GUI operations or commands.

For information on the operation method using GUI, refer to "Perform a Snapshot Backup" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For information on the operation method using commands, refer to Backup execution command (swstbackup)

It is also possible to use the Backup execution status display command (swstbackstat) to check the backup status.



Note

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

- Execute the tracking cancel command (swstcanceltrk).

4.5.2 Performing the synchronized high-speed backup

Synchronous high-speed backup can be carried out using either GUI operations or commands.

For information on the operation method using GUI, refer to "Execution of synchronous high-speed backup" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

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: "Backup synchronous processing start command (swststartsync)".
- Checking status of backup synchronous process: "Backup synchronous processing progress status command (swstsyncstat)".
- Executing of backup: "Backup execution command (swstbackup)".
- Checking backup status: "Backup execution status display command (swstbackstat)".



Note

- The 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.



Note

When the Suspend/Resume function is used for synchronous high-speed backup, the operations listed below cannot be performed from the GUI client. Enter the appropriate command for operation.

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

4.5.3 Restoring Data

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

4.5.3.1 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 GUI operations or commands. For information on the operation method using GUI, refer to "Restore" in "Perform a Snapshot Backup" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For information on the operation method using commands, refer to the Restore execution command (swstrestore)

To check restore status, use the Restore execution status display command (swstreststat)



Note

If restoring for a general file system, stop operations.



Note

Notes about the restoration of a volume group are available in this chapter at "Restoration processing of Volume Group".

4.5.3.2 Restoring an individual file

Restore an individual file as shown below. Note that you cannot restore an individual file from the GUI client.

1. Mount the backup volume. Note that you can check the backup volume in the history list view, or by executing the history information display command.
2. Copy the files to be restored with the cp command.
3. Unmount the backup volume.

4.5.4 Managing a backup history

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



- 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.

4.5.4.1 Displaying a backup history

Backup history can be deleted using either GUI operations or commands. For information on the operation method using GUI, refer to "Delete Backup history" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For information on the operation method using commands, refer to theHistory information display command (swsthistdisp).

4.5.4.2 Deleting a backup history

Backup history can be deleted using either GUI operations or commands. For information on the operation method using GUI, refer to "Delete Backup history" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For information on the operation method using commands, refer to the History information deletion command (swsthistdel).

4.6 Changing the configuration

It will be necessary to change the configuration of AdvancedCopy Manager if any changes are made to a Storage Server or device used in the backup operation.

This section describes how to perform these changes.

4.6.1 Adding a device

After a device is added to any Storage Server, register the new device into the backup operation as shown below.

This process can be carried out using either GUI operations or commands. For information on the operation method using GUI, refer to "Adding a new device" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For information on adding a device using commands, refer to the following steps:

3. Use the Device information collection/reflection command (stgxfwcmssetdev) in the "ETERNUS SF AdvancedCopy Manager User's Guide" for the OS of the Storage Management server to acquire information for the newly added device, and accept management information.
4. Use the Device information configuration command (swstdevinfofet) to set up the operation type of the added device.
5. If the added device was defined as a transaction volume, use the Backup policy configuration command (swstbkpolset) 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" in the "ETERNUS SF AdvancedCopy Manager Installation 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 fetching the device information. If there is insufficient area, allocate the required area, then fetch the device information.
- For the procedure, refer to the version of the following manual that applies to the operating system of the Storage Management Server:
 - "ETERNUS SF AdvancedCopy Manager Operator's Guide", step 2 of "Action required when access to repository fails"
- The repository (directory for the DB space that stores repository data) of the Storage Management Server may be too small when a device is added. Re-estimate the repository size before adding a device. Refer to "Estimating database area" in the "ETERNUS SF AdvancedCopy Manager Installation 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 "Action required when the repository is too small" for details.

4.6.2 Deleting a device

This section describes how to delete different types of devices.

4.6.2.1 Deleting a transaction volume

The process of deleting a transaction volume can be carried out using either GUI operations or commands. For information on the operation method using GUI, refer to "Deleting a transaction volume" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For information on the operation method using commands, refer to the following steps.

6. Use History information deletion command (swsthistdel) to delete the backup history.
7. Use Backup policy deletion command (swstbkpoldel) to delete the backup policy.
8. Use Device information configuration command (swstdevinfofet) to delete the device information for the deleted transaction volume.



Point

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

4.6.2.2 Deleting a backup volume

The process of deleting a backup volume can be carried out using either GUI operations or commands. For information on the operation method using GUI, refer to "Remove a Backup Volume" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For information on the operation method using commands, refer to the following steps:

1. If deleting the backup volume, ensure that you first use the Backup policy display command (swstbkpoldisp) 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, that 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.

9. Use the History information deletion command (swsthistdel) to delete the backup history, removing the relevant volume from the backup history. If there is no backup history, proceed to the next step.
10. Use the Device information setting command (swstdevinforeset) to delete the device information for the deleted backup volume.



Point

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

4.6.3 Adding a Storage Server

The process of adding a Storage Server to the backup operation can be carried out using GUI operations or commands. For information on the operation method using GUI, refer to "Add a New Device" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For information on the operation method using commands, refer to Server information addition command (stgxfwcmaddsrv)

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

1. Register a Storage Server. For more information, refer to "Registering a Storage Server"
2. Fetch all the information on the devices connected to the newly added Storage Server. For more information on this, refer to "Fetching device information on a Storage Server".
3. Set the environment information for the Storage Server. For more information on this, refer to "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 "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 "Setting backup policies".



Note

- System resources such as disk and memory resources may be insufficient when a Storage Server is added. Re-estimate the system resource requirements before adding a Storage Server. Refer to "Required resources" in the "ETERNUS SF AdvancedCopy Manager Installation Guide" for details.
- There may not be sufficient area available for updating the repository when a Storage Server is added. Check if the area required when updating the repository is available before adding the Storage Server. If there is insufficient area, allocate the required area, then add the Storage Server.

For the procedure, refer to the version of the following manual that applies to the Storage Management Server:

- "ETERNUS SF AdvancedCopy Manager Operator's Guide", step 2 of "Action required when access to repository fails"

- The repository (directory for the DB space that stores repository data) of the Storage Management Server may be too small when a Storage Server is added. Re-estimate the repository size before adding a Storage Server. Refer to "Estimating database area" in the "ETERNUS SF AdvancedCopy Manager Installation Guide" for the estimation method. 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 Storage Server. Refer to "Action required when the repository is too small" for the extension method.

4.6.4 Deleting a Storage Server

The process of deleting a Storage Server can be carried out either using GUI operations or commands. For information on the operation method using GUI, refer to "Delete a Device" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For information on the operation method using commands, refer to Server information deletion command (stgxfwcmdelsrv) in the "ETERNUS SF AdvancedCopy Manager User's Guide" for the OS of the Storage Management server.



Point

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

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

1. Delete the backup history for all the transaction volumes on the Storage Server that is being deleted. For more information, refer to "Deleting a backup history".
2. Delete the backup policy for all the transaction volumes on the Storage Server to be deleted.

The process of deleting a backup policy can be carried out using either GUI operations or commands. For information on the operation method using GUI, refer to "Delete Backup Policy" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For information on the operation method using commands, refer to the Backup policy deletion command (swstbkpodel)

3. Delete the volume type that was registered using Setting the operation type for a device for the device connected to the Storage Server
4. Delete the Storage Server.

4.7 Stopping the Operation

To stop a backup operation, stop the daemons on the Storage Server. Normally, they are automatically stopped when the system is shutdown.

It is also possible to individually stop a demon if required. For more details of this, refer to "Starting and Stopping a Deamons".



Note

When the daemons stop, all functions of AdvancedCopy Manager running on the Storage Server also stop.

You should stop daemons on a Storage Management Server only after checking that all the Storage Servers that it manages have shutdown.

4.8 LVM Volume Operation

The backup operation of LVM volumes can be classified into the following two modes, depending on the volume group configuration:

- Backup operation units of volume groups
- Backup operation in units of physical disks (LU: Logical Unit)



Before starting this operation, you need to understand the basic operation of ordinary volumes.

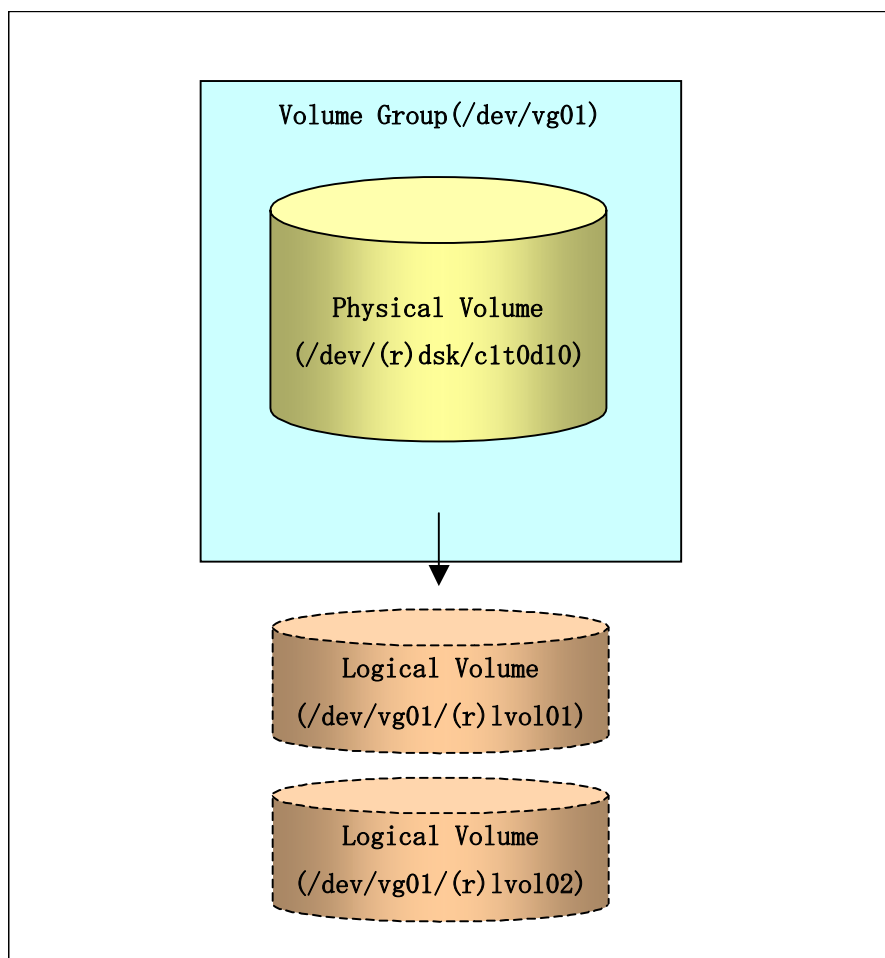
4.8.1 Backup operation in volume group units

Provided that all volume group configurations satisfy the following conditions, backup operation can be performed in units of volume groups:

- One volume group has only one physical disk, and logical volumes are configured so that one physical disk includes n logical volumes.

If the above conditions are not satisfied, backup operation must be performed in units of physical disks.

Figure: Example of a configuration in which backup operation can be performed in volume group units

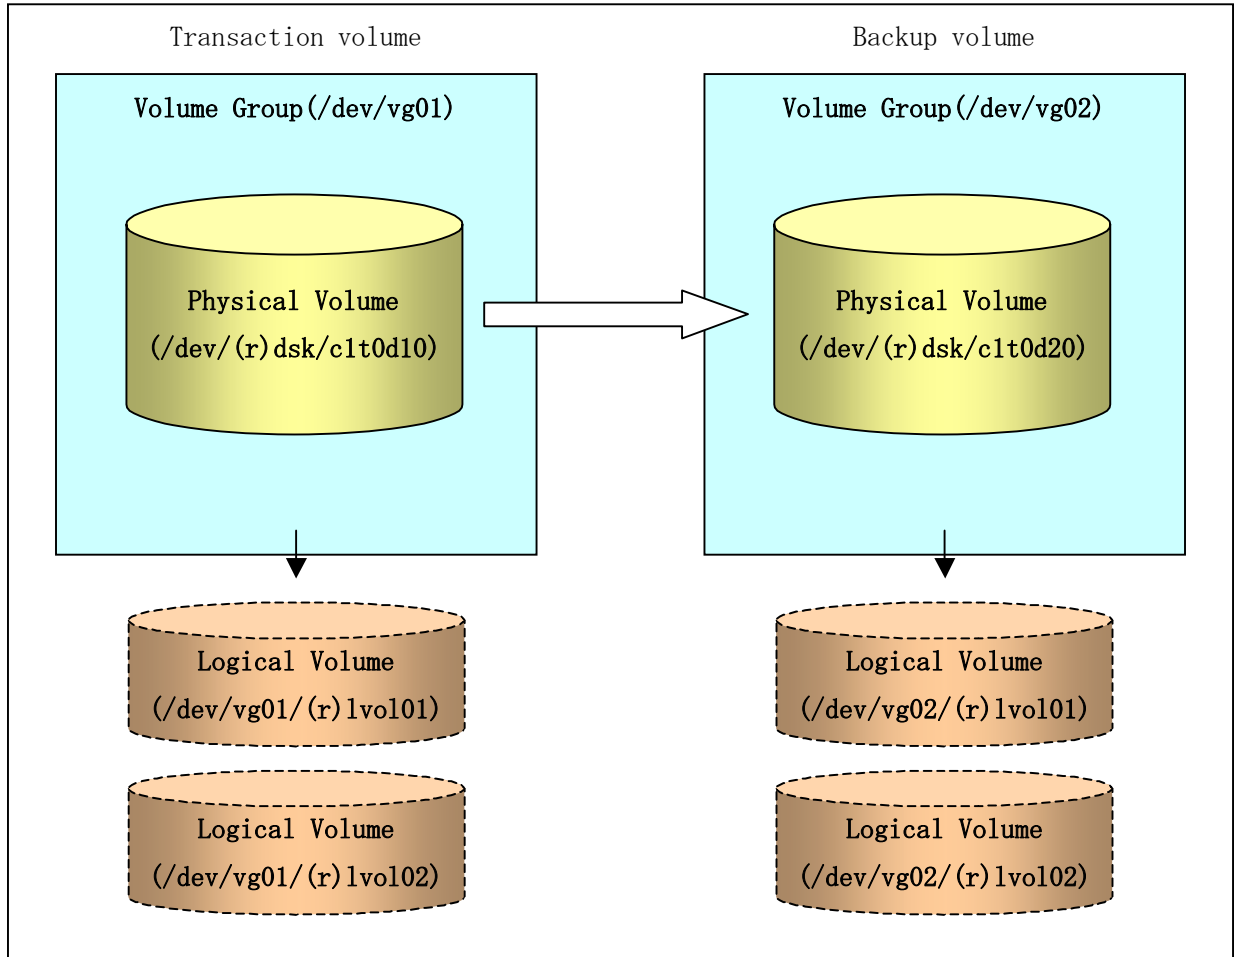


4.8.1.1 Operation design

Take note of the following conditions regarding volume groups when configuring volume groups for use as transaction volumes or backup volumes:

- All physical disks must be the same size.
- The configurations of all logical volumes must be the same.

Figure: Transaction volume and backup volume



4.8.1.2 Preparations

4.8.1.2.1 Saving the volume group configuration information file

The standard post-processing script restores a volume group configuration information file from the following location. Confirm that a volume group configuration information file has been saved.

```
/etc/lvmconf/<volume-group-name>.conf
```



Note

Backup cannot be performed if a volume group configuration information file has not been saved. In this case, post-processing will result in an error if backup is executed. Customize the post-processing script if volume group configuration information has been saved to another file.

4.8.1.2.2 Configuring the transaction volume and backup volume

When configuring the transaction volume and backup volume, specify their volume groups.

Example:

```
# /opt/FJSVswsts/bin/swstdevinfoaset -t /dev/vg01
swstdevinfoaset completed
# /opt/FJSVswsts/bin/swstdevinfoaset -b /dev/vg02
swstdevinfoaset completed
#
```

4.8.1.2.3 Customizing the pre-processing and post-processing scripts

If a volume group is to be backed up, the pre-processing and post-processing scripts must be customized accordingly.

See "Pre-processing and Post-processing for Backup and Restoration" for information on the customization procedure.



Note

If backup is attempted without customization of the scripts, pre-processing results in an error and backup cannot be performed.

4.8.1.3 Backup in units of volume groups

When a transaction volume is part of a volume group, copy processing covers the entire physical disk corresponding to the volume group. Therefore, if the volume group contains multiple logical volumes, data on all the logical volumes is backed up.

Example of a snapshot backup

```
# /opt/FJSVswsts/bin/swstbackup /dev/vg01
/dev/vg01 swstbackup completed
#
```

Example of a synchronous backup

```
# /opt/FJSVswsts/bin/swststartsync /dev/vg01
/dev/vg01 swststartsync completed
(After state of equivalency upkeep)
# /opt/FJSVswsts/bin/swstbackup /dev/vg01
```



```
/dev/vg01 swstbackup completed
```

```
#
```



Note

When using a volume group for backup operation, you should take note of the following:

- Configure the same logical disk as the transaction volume in a volume group separate from the transaction volume, and configure a volume group with the same physical disk size as the backup volume.
To perform the backup by associating transaction volumes with backup volumes, create an association file called "device map file" that defines the relationship between the transaction volume and the backup volume. The "device map" specification is required when performing backup in units of volume groups.

- During backup processing, the LVM configuration information for the backup volume is overwritten by the LVM information from the transaction volume. Therefore, it is critical that the original LVM configuration information for the backup volume must be restored by backup post-processing. For details, see Pre-processing and Post-processing of Backup and Restoration.

- A file system may be configured on a logical volume. In this case, modify both the backup pre-processing and post-processing scripts, and execute the backup using an account with Administrator authority, so that all the logical volumes in the volume group are unmounted or mounted respectively. For details, see Pre-processing and Post-processing of Backup and Restoration.

- When the backup volume is a volume group, do not deactivate the backup volume with the "vgchange" (IM) command during synchronous processing (that is, while running the EC or REC functions).

If the backup volume is deactivated during synchronous processing, not only can the backup volume not be activated again, but the functions for stopping synchronous processing (eg, the backup command and the synchronous processing cancel command) are also disabled. If the backup volume is deactivated in error, forcibly stop synchronous processing by using functions like ETERNUSmgr/GRmgr, and then reactivate the backup volume.

4.8.1.4 Restoration in units of volume groups

When a transaction volume is part of a volume group, copy processing covers the entire physical disk corresponding to the volume group. Therefore, if the volume group contains multiple logical volumes, data on all the logical volumes is restored.

Example of restoration

```
# /opt/FJSVswsts/bin/swstrestore /dev/vg01
```

```
/dev/vg01 swstrestore completed
```

```
#
```



Note

When using a volume group for restoration, note the following:

- A file system may be configured on a logical volume. In this case, modify the backup pre-processing and post-processing scripts, and execute the backup using an account with Administrator authority, so that all logical volumes are unmounted and mounted, respectively.

For details, see "Pre-processing and Post-processing of Backup and Restore".

4.8.2 Backup operation in units of physical disks

When the volume group configuration does not satisfy the conditions of operation in units of volume groups, backup can be performed by operation in units of physical disks.

When backup is performed in units of physical disks, volume group integrity must be maintained. Therefore, all the physical disks in the volume group must be operated synchronously.



Point

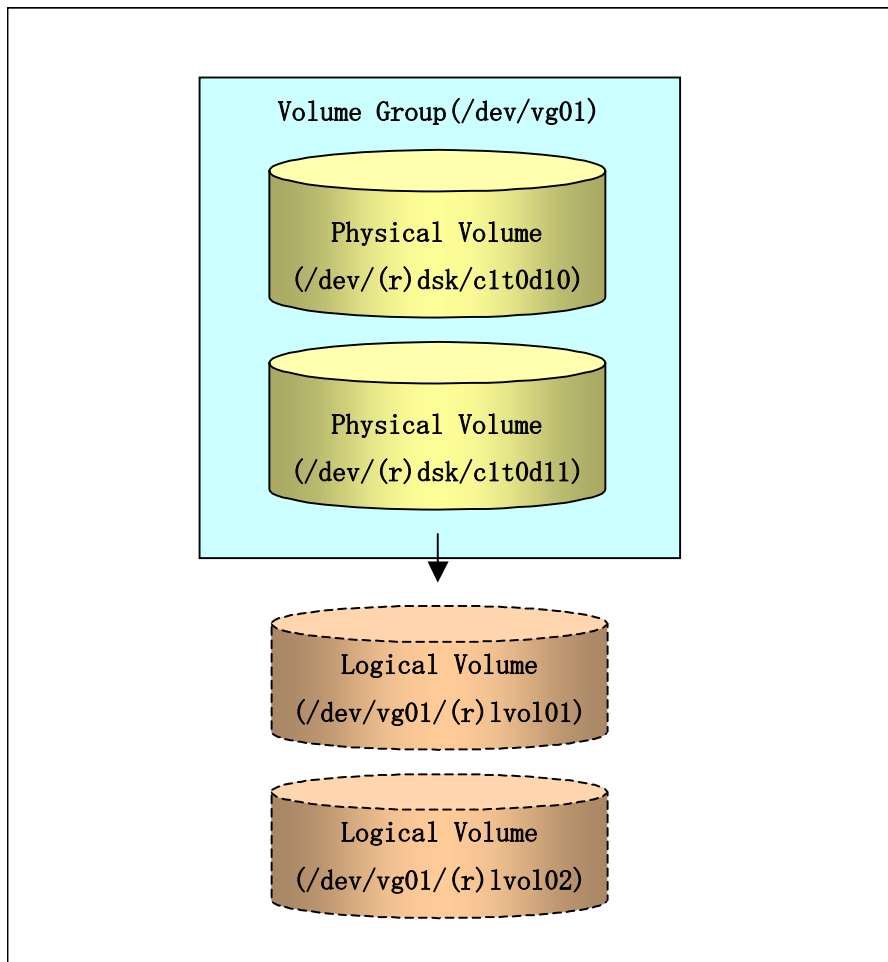
Refer to "Device information display command (stgxfwcmdispdev)", or the "Confirmation of devices in the save logical group" in an "ETERNUS SF AdvancedCopy Manager User's Guide" for details of how to determine which physical volumes should be operated synchronously.



Note

For operation in a cluster configuration, the device name (/dev/(r)dsk/c#t#d#) of the physical disks that comprise the volume group must be the same at all servers that comprise the cluster, and the ETERNUS disk indicated by the device name must also be the same.

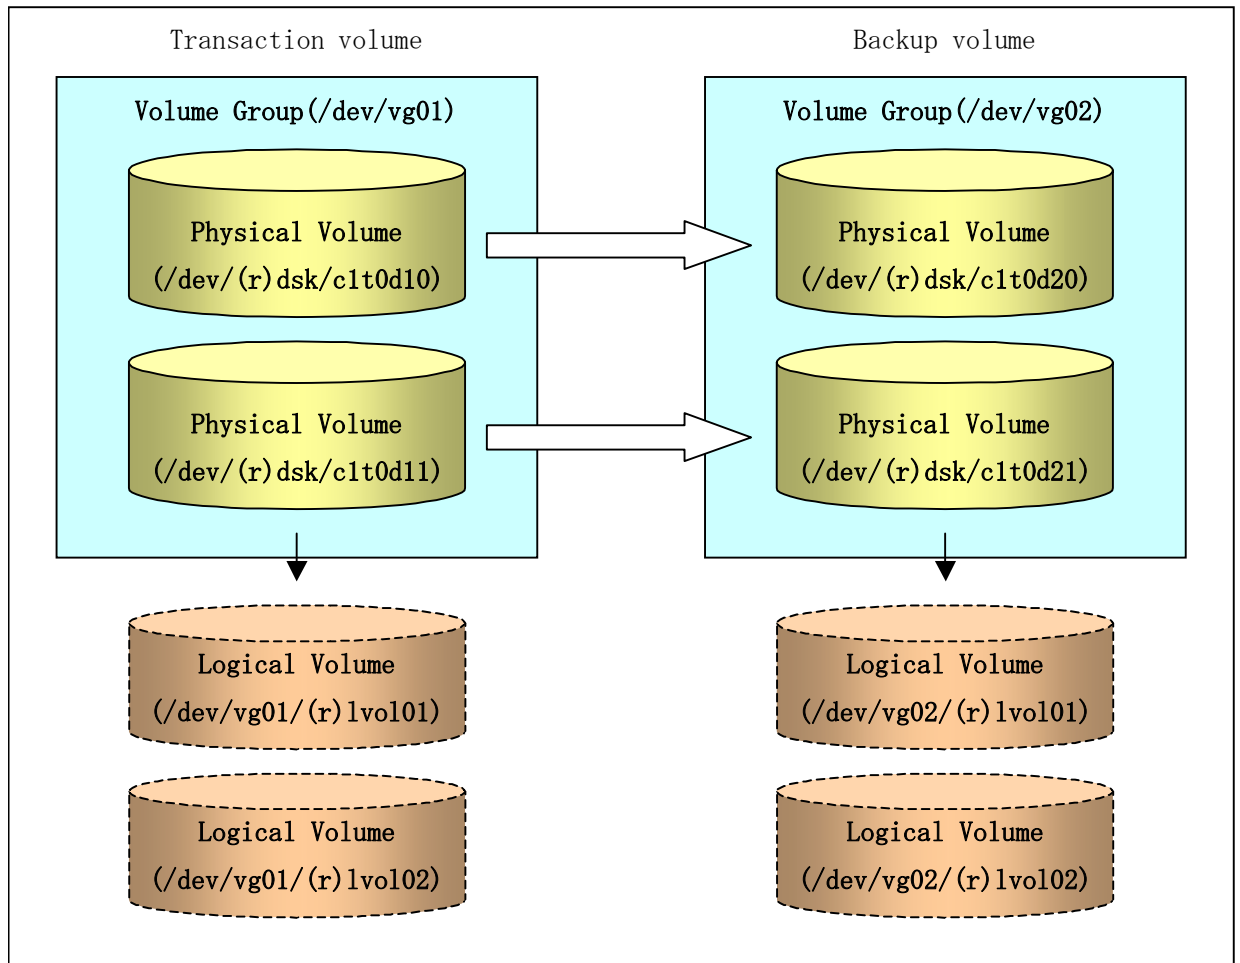
Figure: Sample configuration for operation in units of physical disks



4.8.2.1 Operational configuration

In order to use a replica volume after replication if the original volume is a volume group, the replica volume must have the same logical volume configuration as the original volume, and must be a volume group of the same physical size.

Figure: Transaction volume and backup volume



4.8.2.2 Preparations

4.8.2.2.1 Saving the volume group configuration information file

The volume group configuration information file must be restored during backup post-processing.

Confirm that the volume group configuration file has been saved. It is usually saved to the following location:

```
/etc/lvmconf/<volume-group-name>.conf
```

4.8.2.2.2 Configuring the transaction volume and backup volume

When configuring the transaction volume and backup volume, specify all of the disks in the volume group.

Example:

```
# /opt/FJSVswsts/bin/swstdevinfofet -t /dev/dsk/c1t0d10
swstdevinfofet completed
# /opt/FJSVswsts/bin/swstdevinfofet -t /dev/dsk/c1t0d11
swstdevinfofet completed
# /opt/FJSVswsts/bin/swstdevinfofet -b /dev/dsk/c1t0d20
swstdevinfofet completed
# /opt/FJSVswsts/bin/swstdevinfofet -b /dev/dsk/c1t0d21
swstdevinfofet completed
#
```

4.8.2.2.3 Preparing a device map file

In the case of a backup of an LVM volume, a device map file must be created. This is because a backup volume with the same volume structure as the transaction volume must be specified.

[Example of a device map file]

```
# Transaction volume      Backup volume
/dev/dsk/c1t0d10          /dev/dsk/c1t0d20
/dev/dsk/c1t0d11          /dev/dsk/c1t0d21
```

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

4.8.2.3 Backup in units of physical disks

Perform the operation by synchronizing all physical disks in the volume group.

Perform the required pre-processing or post-processing work for each volume group before and after (respectively) the backup. Disable pre-processing and post-processing when operating individual physical disks.

Example of snapshot backup

```
(Perform preprocessing for the transaction and backup volumes.)
# /opt/FJSVswsts/bin/swstbackup /dev/dsk/c1t0d10 -Xdevmap /acm/devmap.txt
/dev/dsk/c1t0d10 swstbackup completed
# /opt/FJSVswsts/bin/swstbackup /dev/dsk/c1t0d11 -Xdevmap /acm/devmap.txt
/dev/dsk/c1t0d11 swstbackup completed
```

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

Example of synchronous backup

```
(Perform pre-processing for the backup volume.)
# /opt/FJSVswsts/bin/swststartsync /dev/dsk/clt0d10 -Xdevmap /acm/devmap.txt
/dev/dsk/clt0d10 swststartsync completed
# /opt/FJSVswsts/bin/swststartsync /dev/dsk/clt0d11 -Xdevmap /acm/devmap.txt
/dev/dsk/clt0d11 swststartsync completed
(After state of equivalency upkeep)
(Perform preprocessing for the transaction volume.)
# /opt/FJSVswsts/bin/swstbackup /dev/dsk/clt0d10
/dev/dsk/clt0d10 swstbackup completed
# /opt/FJSVswsts/bin/swstbackup /dev/dsk/clt0d11
/dev/dsk/clt0d11 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: Backup pre-processing and post-processing

	Pre-processing	Post-processing
Transaction volume	<ol style="list-style-type: none"> 1. Ensure data integrity by preventing access to all logical volumes in the volume group. 2. If file systems are included, unmount all file systems in the volume group. 	<ol style="list-style-type: none"> 1. If file systems are included, remount the volumes that were unmounted by pre-processing.
Backup volume	<ol style="list-style-type: none"> 1. Prevent access to all logical volumes in the volume group. 2. If file systems are included, unmount all file systems in the volume group. 	<ol style="list-style-type: none"> 1. Reconfigure the volume group 2. If file systems are included, remount the volumes that were unmounted by pre-processing.

Reconfiguring the volume group

Reconfigure the volume group as follows:

1. Deactivate the volume group.

```
# /usr/sbin/vgchange -a n /dev/vg02
#
```

2. Restore volume group configuration information.

```
# /usr/sbin/vgcfgrestore -n /dev/vg02 /dev/rdisk/c1t0d20
# /usr/sbin/vgcfgrestore -n /dev/vg02 /dev/rdisk/c1t0d21
#
```

3. Reactivate the volume group.

```
[When cluster operation is enabled]
# /usr/sbin/vgchange -c y /dev/vg02

# /usr/sbin/vgchange -a e /dev/vg02

[When cluster operation is disabled]
# /usr/sbin/vgchange -a y /dev/vg02

#
```

Reconfigure shared mode volume groups using the following procedure:

1. Stop the volume group (on all nodes involved in the transaction).

```
# /usr/sbin/vgchange -a n /dev/vg02
#
```

2. Restore the volume group configuration information (on the node where the volume group was created).

```
# /usr/sbin/vgcfgrestore -n /dev/vg02 /dev/rdisk/c1t0d20
# /usr/sbin/vgcfgrestore -n /dev/vg02 /dev/rdisk/c1t0d21
#
```

3. Mark the volume group as shareable (on the node where the volume group was created).

```
# /usr/sbin/vgchange -S y -c y /dev/vg02
#
```

4. Start the volume group (on all nodes involved in the transaction).

```
# /usr/sbin/vgchange -a s /dev/vg02
#
```

4.8.2.4 Restoration in units of physical disks

Perform this operation by synchronizing all physical disks in the volume group.

Perform the required pre-processing or post-processing work for each volume group before and after the restoration respectively. Disable pre-processing and post-processing when using individual physical disks.

Example of restoration

```
(Perform preprocessing for the transaction and backup volumes.)  
# /opt/FJSSwsts/bin/swstrestore /dev/dsk/clt0d10  
/dev/dsk/clt0d10 swstrestore completed  
# /opt/FJSSwsts/bin/swstrestore /dev/dsk/clt0d11  
/dev/dsk/clt0d11 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: Restoration processing

	Pre-processing	Post-processing
Backup volume	1. Ensure data integrity by preventing access to all logical volumes in the volume group.	Post-processing is not required.
Restoration destination volume	1. Prevent access to all logical volumes in the volume group. 2. If file systems are included, unmount all file systems in the volume group.	1. Reconfigure the volume group 2. If file systems are included, remount the volumes that were unmounted during pre-processing.

Reconfiguring the volume group

Reconfigure the volume group as follows:

1. Deactivate the volume group.

```
# /usr/sbin/vgchange -a n /dev/vg01  
#
```

2. Restore the volume group configuration information.

```
# /usr/sbin/vgcfgrestore -n /dev/vg01 /dev/rdisk/clt0d10  
# /usr/sbin/vgcfgrestore -n /dev/vg01 /dev/rdisk/clt0d11  
#
```


3. Reactivate the volume group.

```
[When cluster operation is enabled]
# /usr/sbin/vgchange -c y /dev/vg01

# /usr/sbin/vgchange -a e /dev/vg01

[When cluster operation is disabled]
# /usr/sbin/vgchange -a y /dev/vg01

#
```

Reconfigure shared mode volume groups using the following procedure:

1. Stop the volume group (on all nodes involved in the transaction).

```
# /usr/sbin/vgchange -a n /dev/vg01

#
```

2. Restore volume group configuration information (on the node where the volume group was created).

```
# /usr/sbin/vgcfgrestore -n /dev/vg01 /dev/rdisk/c1t0d10
# /usr/sbin/vgcfgrestore -n /dev/vg01 /dev/rdisk/c1t0d11

#
```

3. Mark the volume group as shareable (on the node where the volume group was created).

```
# /usr/sbin/vgchange -S y -c y /dev/vg01

#
```

4. Start the volume group (on all nodes involved in the transaction).

```
# /usr/sbin/vgchange -a s /dev/vg01

#
```

4.9 VxVM volume operation

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

Since disk group consistency needs to be maintained, all the physical disks in the disk group must be synchronized for the backup operation.



Point

Refer to "device information display command (stgxfwcmdispdev)" or "Checking devices in the same logical group" in the "ETERNUS SF AdvancedCopy Manager User's Guide" for information on how to check the physical volumes that require synchronization.



Note

An understanding of the basic operation of transaction volumes is required to use VxVM volumes.



Note

Snapshot type operations are recommended for operations that are in physical disk units. During total copy or differential copy when the synchronous type of operation is used, VxVM commands or similar commands that cause disk access to the copy destination disk cannot be executed.



Note

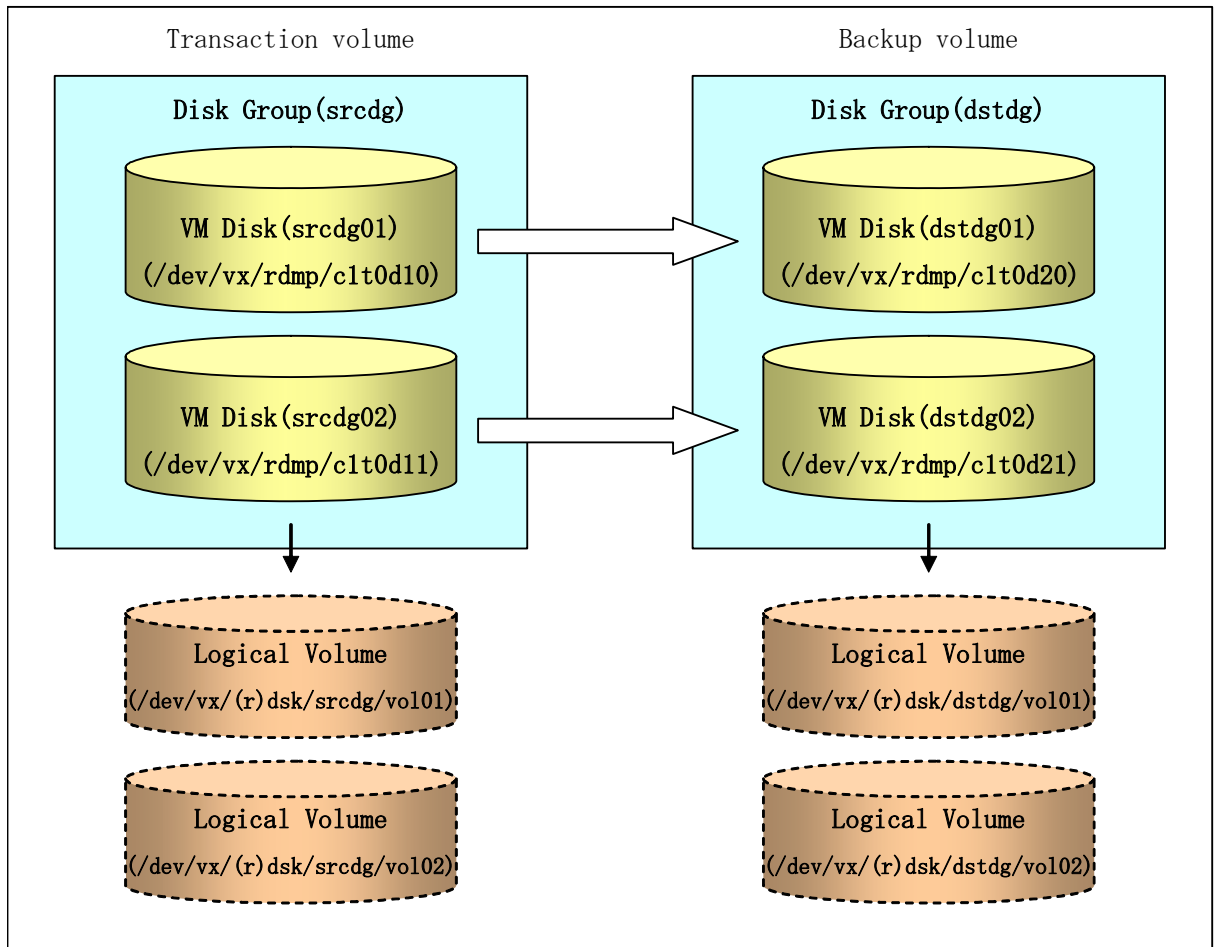
For operation in a cluster configuration, the device name (/dev/(r)dsk/c#t#d#) of the physical disks that comprise the volume group must be the same at all servers that comprise the cluster, and the ETERNUS disk indicated by the device name must also be the same.

4.9.1 Operational configuration

Note the following points when configuring disk groups for transaction or backup volumes:

- The quantity, size and type of VM disks is the same.
- The logical configuration is the same.

Figure: Transaction volume and backup volume



4.9.2 Preparations

4.9.2.1 Checking the disk group configuration information file

The disk group must be reconfigured as part of backup post-processing. Check that the configuration information file required for reconfiguring the disk group has been saved.

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

4.9.2.2 Transaction volume / backup volume settings

Set all the physical disks in the disk group when configuring the transaction volume / backup volume settings.

Example:

```
# /opt/FJSVswsts/bin/swstdevinfoaset -t /dev/vx/dmp/clt0d10
swstdevinfoaset completed
# /opt/FJSVswsts/bin/swstdevinfoaset -t /dev/vx/dmp/clt0d11
swstdevinfoaset completed
# /opt/FJSVswsts/bin/swstdevinfoaset -b /dev/vx/dmp/clt0d20
swstdevinfoaset completed
```

```
# /opt/FJSVswsts/bin/swstdevinfoaset -b /dev/vx/dmp/c1t0d21
swstdevinfoaset completed
#
```

4.9.2.3 Preparing a device map file

Since VxVM volume backup operations require specification of a backup volume that has the same volume configuration as the transaction volume, a device map file must be created.

[Example of a device map file]

```
# Transaction volume name      Output destination backup volume name
/dev/vx/dmp/c1t0d10           /dev/vx/dmp/c1t0d20
/dev/vx/dmp/c1t0d11           /dev/vx/dmp/c1t0d21
```

Refer to "Preparing a device map file" for details of device map files.

4.9.3 Executing backup

Synchronize all the physical disks in the disk group for this operation.

The required pre-processing and post-processing are executed in disk group units before and after the backup operation. Do not run pre- and post-processing when using individual physical disks.

Example of snapshot type backup

```
(Pre-processing performed for the transaction volume / backup volume)
# /opt/FJSVswsts/bin/swstbackup /dev/vx/dmp/c1t0d10 -Xdevmap /acm/devmap.txt
/dev/vx/dmp/c1t0d10 swstbackup completed
# /opt/FJSVswsts/bin/swstbackup /dev/vx/dmp/c1t0d11 -Xdevmap /acm/devmap.txt
/dev/vx/dmp/c1t0d11 swstbackup completed
#
(Post-processing performed for the transaction volume / backup volume)
```

Example of synchronous type backup

```
(Pre-processing performed for the backup volume)
# /opt/FJSVswsts/bin/swststartsync /dev/vx/dmp/c1t0d10 -Xdevmap /acm/devmap.txt
/dev/vx/dmp/c1t0d10 swststartsync completed
# /opt/FJSVswsts/bin/swststartsync /dev/vx/dmp/c1t0d11 -Xdevmap /acm/devmap.txt
/dev/vx/dmp/c1t0d11 swststartsync completed
(After equivalency maintenance status)
(Pre-processing performed for the transaction volume)
```

```
# /opt/FJSVswsts/bin/swstbackup /dev/vx/dmp/clt0d10
/dev/vx/dmp/clt0d10 swstbackup completed
# /opt/FJSVswsts/bin/swstbackup /dev/vx/dmp/clt0d11
/dev/vx/dmp/clt0d11 swstbackup completed
#
(Post-processing performed for the transaction volume / backup volume)
```

The pre-processing and post-processing performed before and after backup is as follows:

Table: Backup pre-processing and post-processing

	Pre-processing	Post-processing
Transaction volume	1. Access to all logical volumes in the disk group is stopped to guarantee consistency. 2. If file systems are included, all file systems in the disk group are unmounted. 3. If the disk group has not been imported, the disk group is imported.	1. If a file system is included, the volume that was unmounted by the pre-processing is mounted.
Backup volume	1. Access to all logical volumes in the disk group is stopped. 2. If file systems are included, all file systems in the disk group are unmounted. 3. The disk group is exported. 4. The physical disks in the disk group are set to offline.	1. The physical disks set to offline by the pre-processing are set to online. 2. The disk group is reconfigured. 1. If a file system is included, the volume that was unmounted by the pre-processing is mounted.

Reconfiguring a disk group

Use the following procedure to reconfigure a disk group:

1. Analyze the restore precommit.

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

Installing volume manager disk header for clt0d20 ...
Installing volume manager disk header for clt0d21 ...
-
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:
vxconfigstore -c dstdg ==> to commit the restoration.
vxconfigstore -d dstdg ==> to abort the restoration.
#
```

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

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

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



Note

If the disk group is registered as a cluster system resource, perform online/offline processing for the disk group resource instead of disk group import/export processing. If the mount point is registered as a cluster system resource, perform online/offline processing for the mount resource instead of file system mount/unmount processing.



Note

If the system performs disk exchange or similar, there may be more than one configuration information backup for a single disk group. In this case, the disk group ID displayed after execution of the above commands must be specified for execution instead of using a disk group.



Note

After these operations, the volumes in the disk group are synchronized in the background. The synchronization processing may take some time, depending on the volume configuration.

Note that, the volumes can still be used while synchronization is being performed.

4.9.4 Restoring

Synchronize all the physical disks in the disk group for this operation.

The required pre-processing and post-processing are executed in disk group units before and after the restore operation. Do not run pre and post-processing when operating individual physical disks.

Restoration example

```
(Pre-processing performed for the transaction volume / backup volume)
# /opt/FJVSswsts/bin/swstrestore /dev/vx/dmp/clt0d10
/dev/vx/dmp/clt0d10 swstrestore completed
# /opt/FJVSswsts/bin/swstrestore /dev/vx/dmp/clt0d10
/dev/vx/dmp/clt0d10 swstrestore completed
#
(Post-processing performed for the transaction volume / backup volume)
```

The pre-processing and post-processing performed before and after restoration are as follows:

Table: Restoration pre-processing and post-processing

	Pre-processing	Post-processing
Backup volume	<ol style="list-style-type: none"> 1. Access to all logical volumes in the disk group is stopped to guarantee consistency. 2. If the disk group has not been imported, the disk group is imported. 	None required.
Restore destination volume	<ol style="list-style-type: none"> 1. Access to all logical volumes in the disk group is stopped. 2. If file systems are included, all file systems in the disk group are unmounted. 3. The disk group is exported. 4. The physical disks in the disk group are set to offline. 	<ol style="list-style-type: none"> 1. The physical disks set to offline by the pre-processing are set to online. 2. The disk group is reconfigured. 3. If a file system is included, the volume that was unmounted by the pre-processing is mounted.

Reconfiguring a disk group

Use the following procedure to reconfigure a disk group:

1. Analyze the restore precommit.

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

Installing volume manager disk header for clt0d10 ...
Installing volume manager disk header for clt0d11 ...
-
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 changes required to restore the copy destination disk group configuration.

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

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



Note

If the disk group is registered as a cluster system resource, perform online/offline processing for the disk group resource instead of disk group import/export processing. If the mount point is registered as a cluster system resource, perform online/offline processing for the mount resource instead of file system mount/unmount processing.



Note

If the system performs disk exchange or similar, there may be more than one configuration information backup for a single disk group.

In this case, the disk group ID displayed after execution of the above commands must be specified for execution instead of a disk group.



Note

After these operations, the volumes in the disk group are synchronized in the background. The synchronization processing may take some time, depending on the volume configuration.

Note that, the volumes can still be used while synchronization is being performed.

Chapter 5 Backup Process That Uses the Suspend/Resume Function

This chapter provides details of the backup process that uses the Suspend/Resume function.

5.1 Overview

The Suspend/Resume function sets Suspend/Resume for the equivalency maintenance status by using Equivalent Copy (EC). Using this function for differential copying from the Suspend state enables higher speed synchronous backup operation.

AdvancedCopy Manager provides two backup functions:

Snapshot high speed backup where One Point Copy (OPC) is used to create a backup and Synchronous high speed backup where EC is used to create a backup .

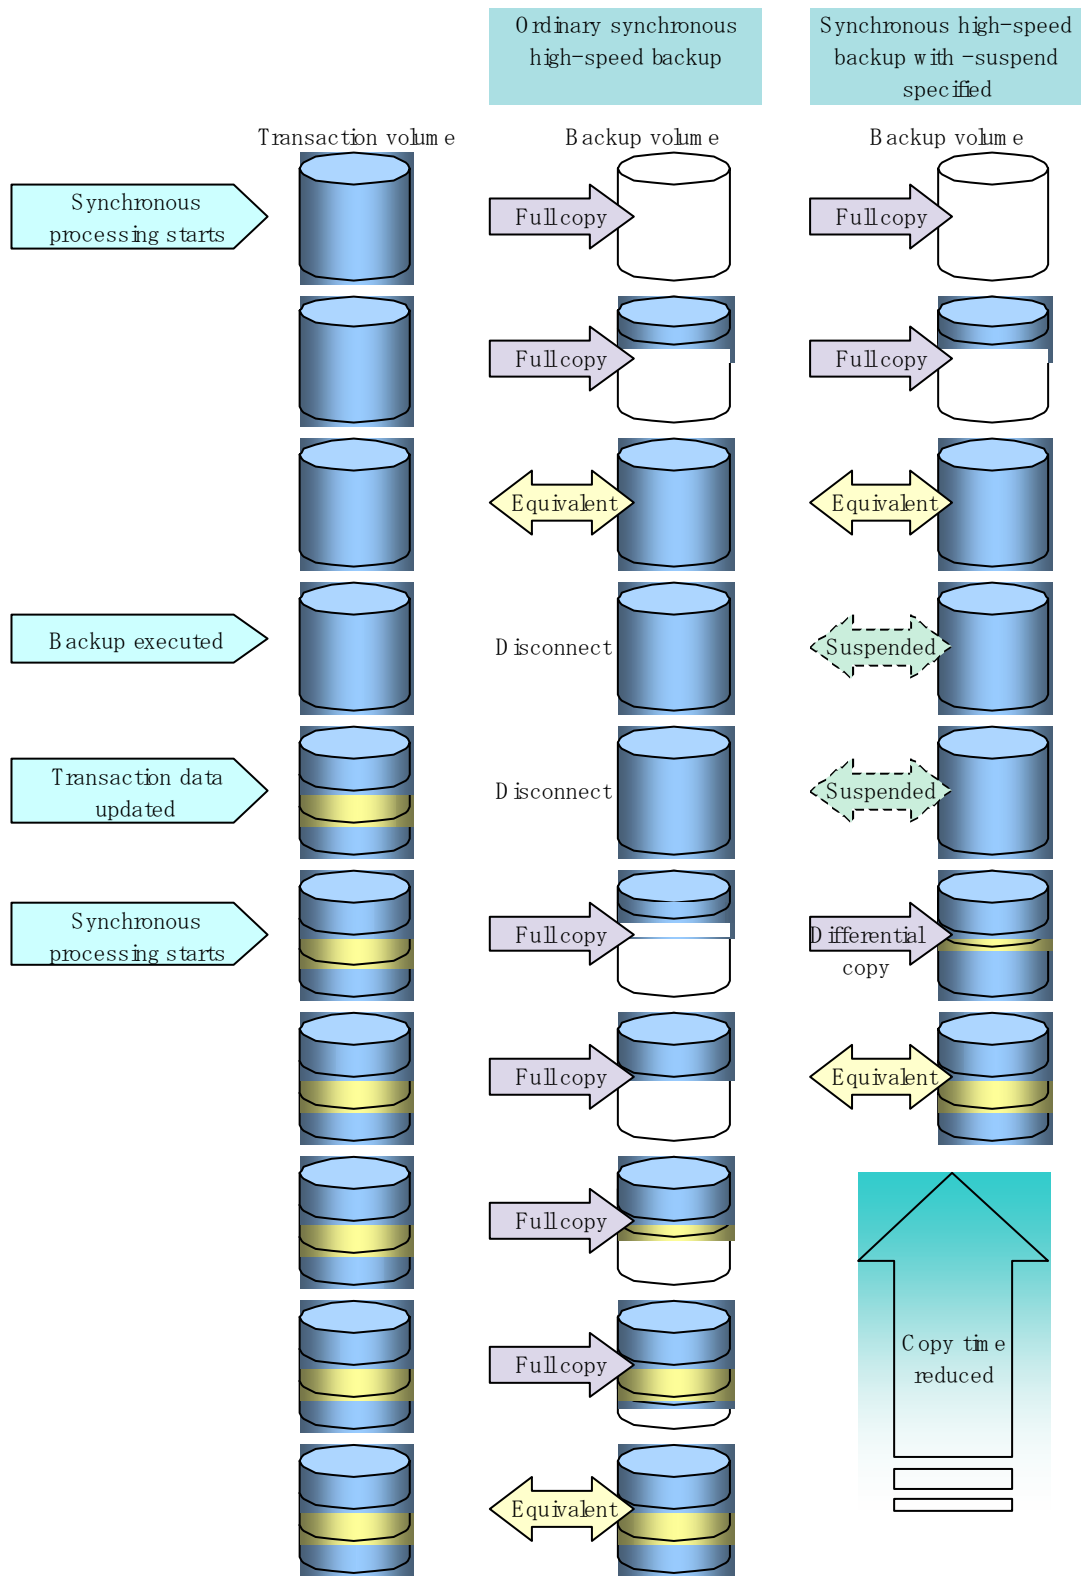
The Suspend/Resume function can be thought of as a kind of synchronous high speed backup that uses EC.

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, the preparation time and backup time is significantly reduced.

A comparison between ordinary synchronous high speed backup and synchronous high speed backup using the Suspend/Resume function is shown below.

Figure: Comparison between ordinary synchronous high speed backup and synchronous high speed backup using the Suspend/Resume function



5.2 Backup configuration

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 (see "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" specified is executed, the EC session between the transaction volume and backup volume is maintained in the Suspend state. In the Suspend state, the backup volume cannot be used as the backup destination of another transaction volume.

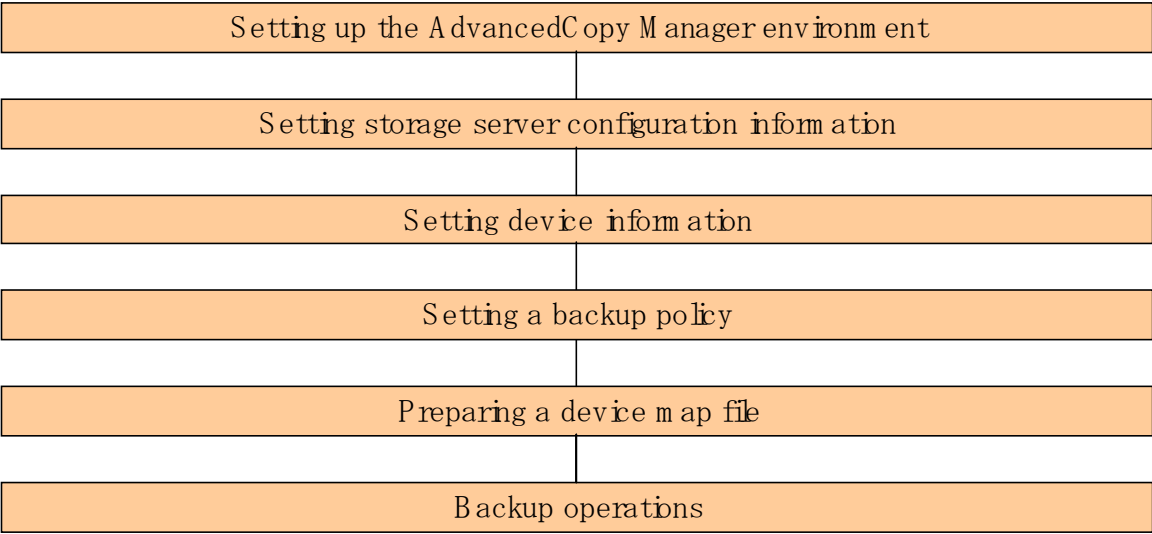


For the backup process that uses the Suspend/Resume function, each transaction volume requires a number of backup volumes equivalent to the number of saved generation number + one (1).

5.3 Flow of Operations

The flow of backup operations using the Suspend/Resume function is shown below.

Figure: Flow of backup operations using the Suspend/Resume function



5.4 Preparations

This section describes how to prepare for the backup process that uses the Suspend/Resume function.

5.4.1 Configuring up the AdvancedCopy Manager environment

Install AdvancedCopy Manager and configure its operating environment according to the "ETERNUS SF AdvancedCopy Manager Installation Guide" and "ETERNUS SF AdvancedCopy Manager Operator's Guide".

5.4.2 Configuring Storage Server information

Set the configuration information on the Storage Server processing the backup operations.

This process can be executed using the GUI and command operations.

For details on the operation using the GUI, refer to "Register a New Storage Server" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For details on the operation using the command, refer to the Storage_Server configuration information configuration command ([swstsvrset](#)). It is also possible to use the [Storage Server configuration information display command \(swstsvrdisp\)](#) to refer to the Storage Server configuration information that was set.

5.4.3 Configuring device information

Configuring device information involves registering a transaction volume as the backup source and a backup volume as the backup destination. This process can be executed using the GUI and command operations.

For details on the operation using the GUI, refer to "Define the Role of a Volume" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For details on registering the transaction volume/backup volume for using the command, refer to the [Device information configuration command \(swstdevinfo\)](#). It is also possible to use the [Device usage status display command \(swstdevdisp\)](#) to refer to the volume information that was registered.



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 Configuring a backup policy

Before a backup policy can be set you must register as many usable backup volumes as are necessary for the backup, according to the specified backup policy.

This process can be executed using the GUI and command operations.

For details on the operation using the GUI, refer to "Set Backup Policy" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For details on the operation using the command, refer to the following:

- Backup policy settings: Backup policy configuration command (swstbkpolset)
- Checking the backup policy that was set: Backup policy display command (swstbkpoldisp)
- Deletion of the backup policy: Backup policy deletion command (swstbkpoldel)

5.4.5 Preparing a device map file

In its backup operations, AdvancedCopy Manager automatically selects a backup destination volume from a group of devices registered as backup volumes, according to whether the device has the same capacity as the transaction volume.

If there is a suspended backup volume at the time of backup destination selection and the start of the synchronous processing, the associated device as defined in the device map file is selected first and used as the backup destination.

To easily manage 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 better management of backup volume resources.



To manage multiple generations of backups, multiple device map files must be prepared.

5.5 Backup Operations

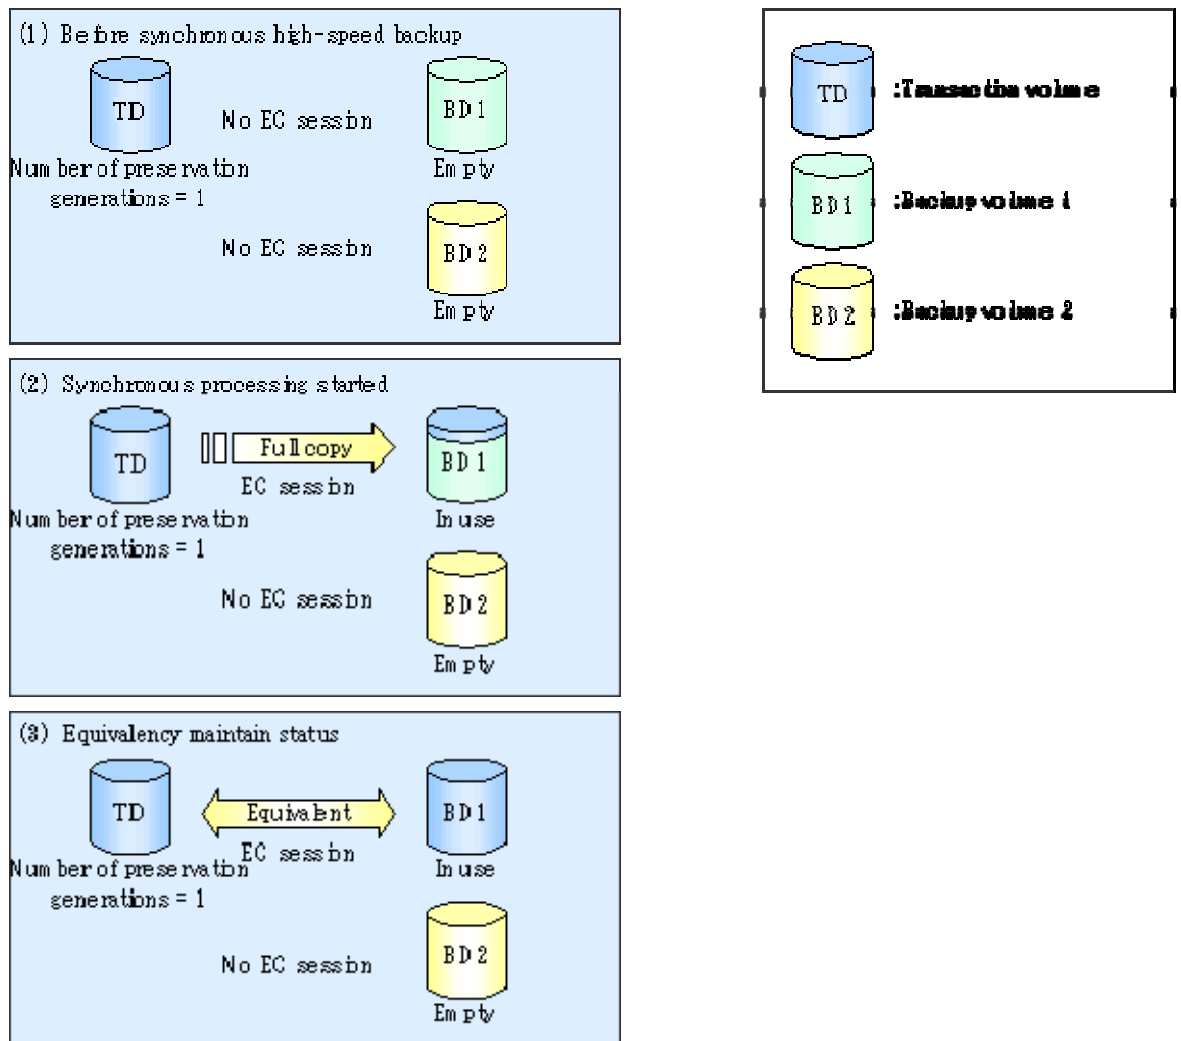
For synchronous high speed backup using the Suspend/Resume function, execute the backup execution command (swstbackup) with "-suspend" specified after synchronous processing has started and the equivalency maintenance status has been attained., then suspend the synchronous processing. When history information is deleted because the number of generations has been exceeded or when the history information deletion command is executed, the backup volume enters the Suspend state without any history information.

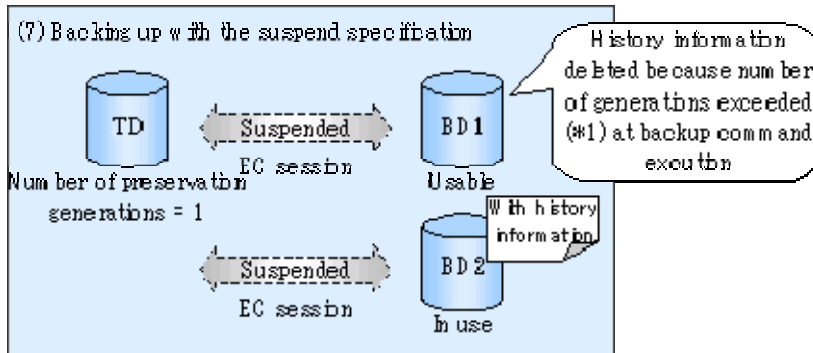

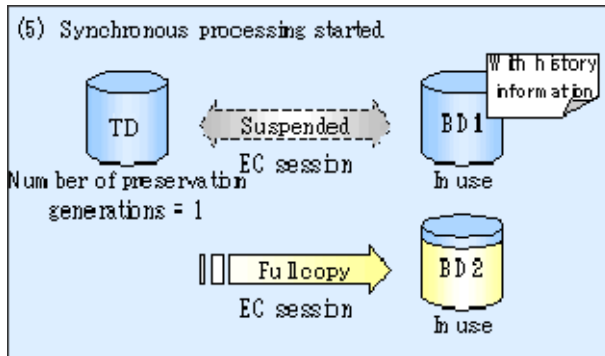
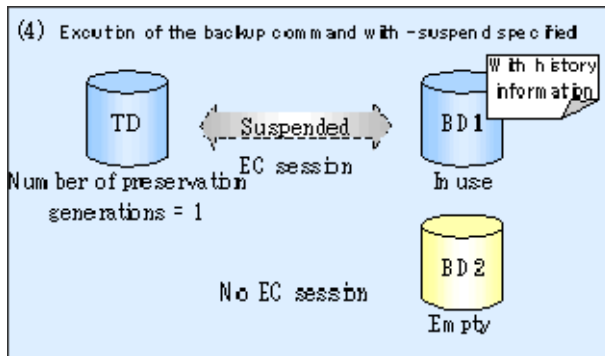
If new synchronous processing starts at this time, then the suspended backup volume without history information is selected first, and differential copying starts. Only data that has been updated since the Suspend state was last set is copied, allowing the equivalency maintenance status to be attained quickly.

Note that synchronous high-speed backup that uses the Suspend/Resume function can be operated from the GUI client. Refer to "Perform a Snapshot Backup" in the "ETERNUS SF AdvancedCopy Manager User's Guide" for details.


The flow of synchronous high speed backup operations using the Suspend/Resume function is shown below.


Figure: Synchronous high speed backup operations using the Suspend/Resume function





*1 "Number of generations exceeded" means that the oldest history information is automatically deleted. Existing history information has a certain number of saved generations when the backup command is executed.





5.5.1 Starting synchronous processing

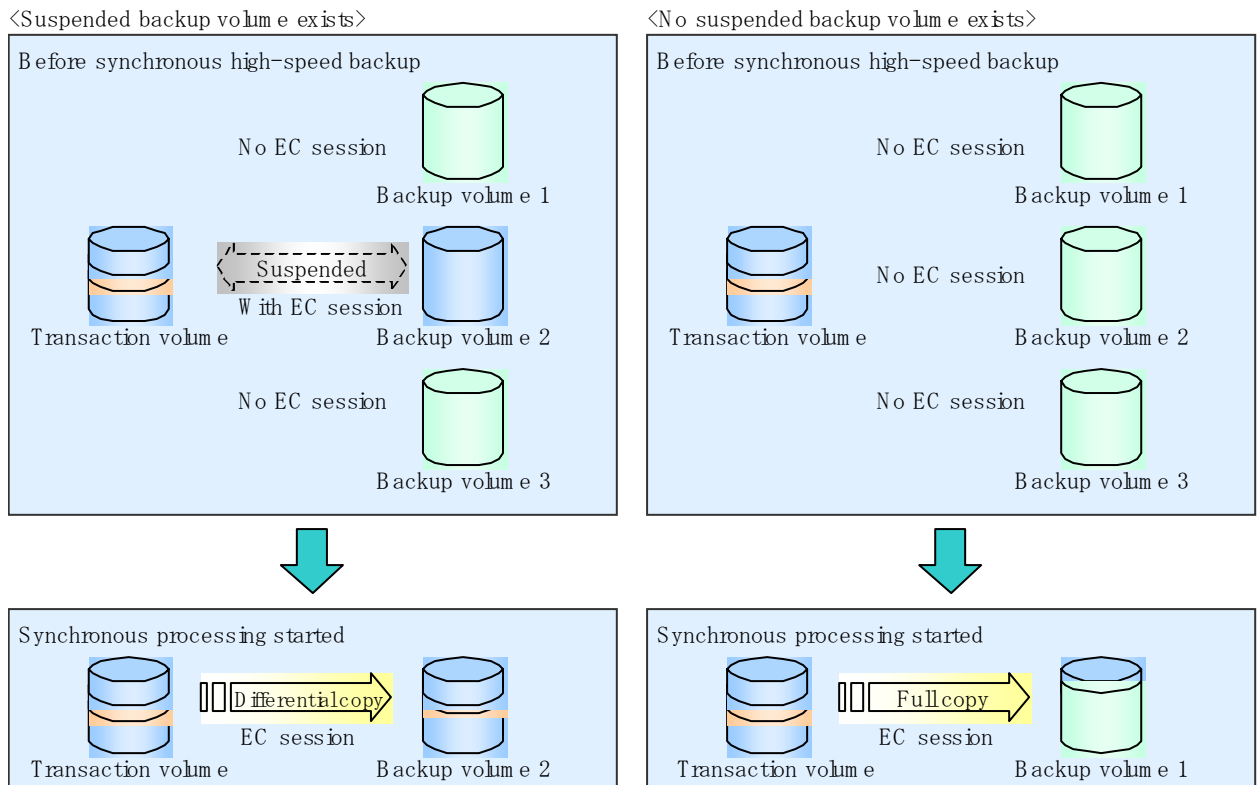
For synchronous high speed backup, execute the synchronous processing start command (swststartsync) at a command prompt to start the synchronous processing. If there is a suspended backup volume at the start of the synchronous processing, then the following devices are selected first, and EC resumes (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 has no history information

If there is no suspended backup volume at the start of the synchronous processing, EC (ie, a full copy) starts.

The process of backup volume selection at the start of the synchronous processing is shown below.

Figure: 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 the backup synchronous processing progress display command (swstsyncstat) at a command prompt.

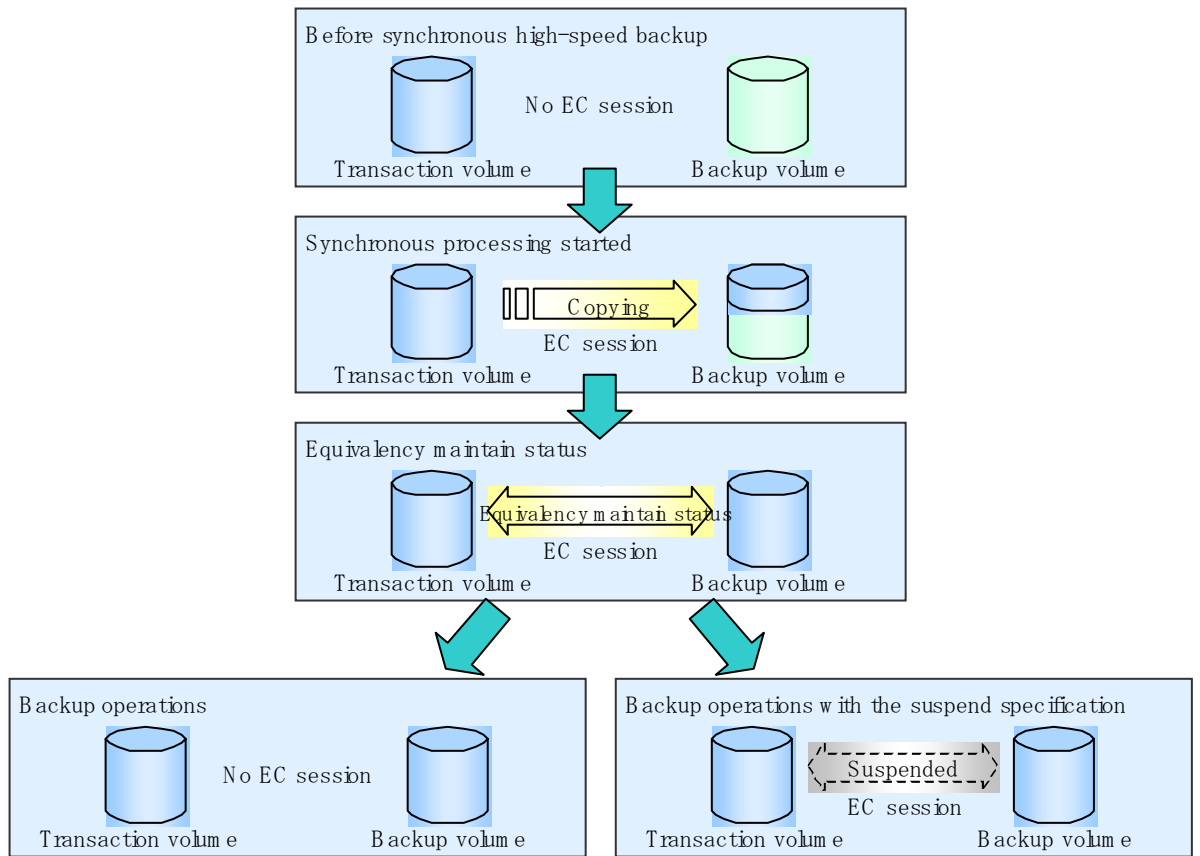
5.5.2 Backing up

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 reach the equivalency maintenance status, execute the backup execution command (swstbackup) with "-suspend" specified.

In synchronous high speed backup, executing the backup command without "-suspend" specified releases the EC session. Executing it with "-suspend" specified suspends the EC session instead of releasing it.

Ordinary backup operations and backup operations with the suspend specification are shown below.

Figure: Ordinary backup operations and backup operations with the suspend specification



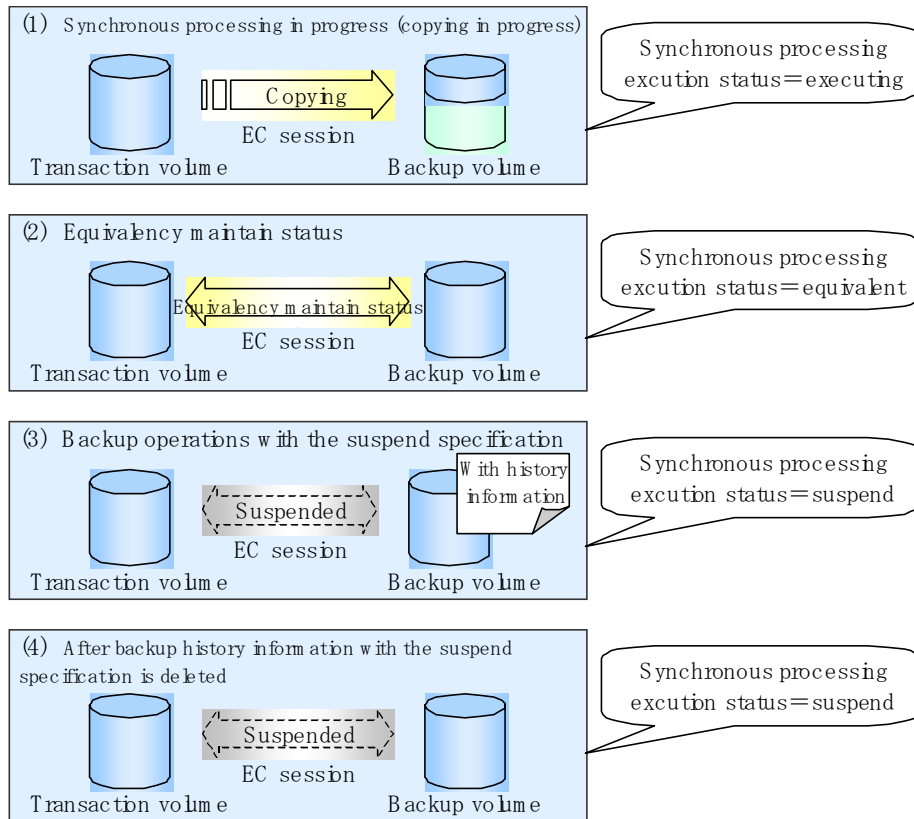
5.5.3 Displaying the synchronous processing execution status

To check the status of synchronous high speed backup that uses the Suspend/Resume function, execute the synchronous processing progress display command (swstsyncstat) at a command prompt. 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 (ie, copying) is in progress
- equivalent: equivalency maintenance status
- suspend: suspended processing

The command displays the synchronous processing execution status, as shown below.

Figure: Ordinary backup operations and backup operations with the suspend specification



5.5.4 Canceling synchronous processing

To cancel synchronous processing (during an EC session), execute the synchronous processing cancel command (swstcancelsync) at a command prompt. The statuses of the different types of synchronous processing cancelled are explained below.

- executing: synchronous processing in progress (ie, copying in progress)
- equivalent: equivalency maintenance status
- suspend: suspended processing

After the backup command with `-suspend` 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 occurring (by running multiple EC sessions), then only the EC sessions that correspond 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.



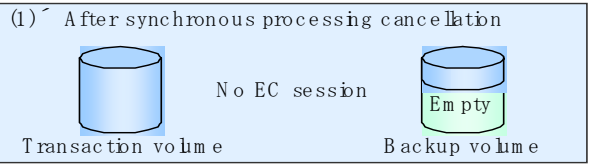

Point

- Cancelling 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.

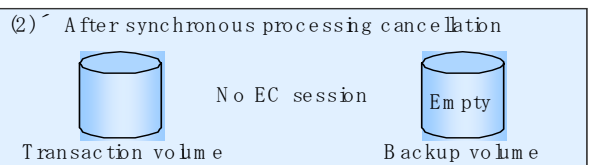
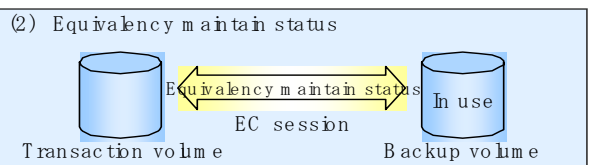
Operational states occurring when synchronous processing is cancelled are shown below:

Figure: Operational states caused by synchronous processing cancellation

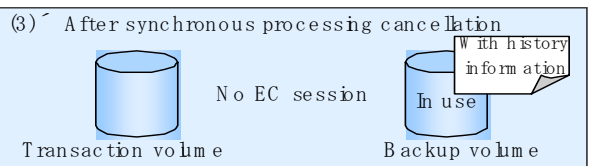
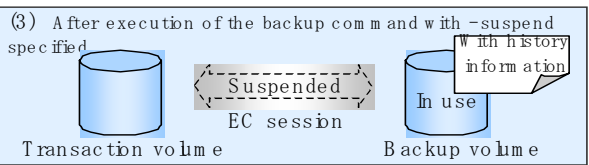
〈Case of the backup volume being cleared because copy processing is cancelled〉




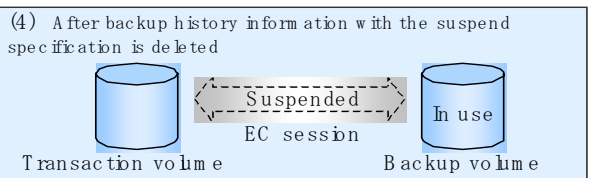
〈Case of the backup volume being cleared because of cancellation under the equivalency maintain status〉



〈Case of the backup volume remaining in use after cancellation in the Suspend state〉



〈Case of the backup volume being cleared because of cancellation in the Suspend state〉



5.5.5 History Management

To display backup history information, execute the history information display command (swsthistdisp).

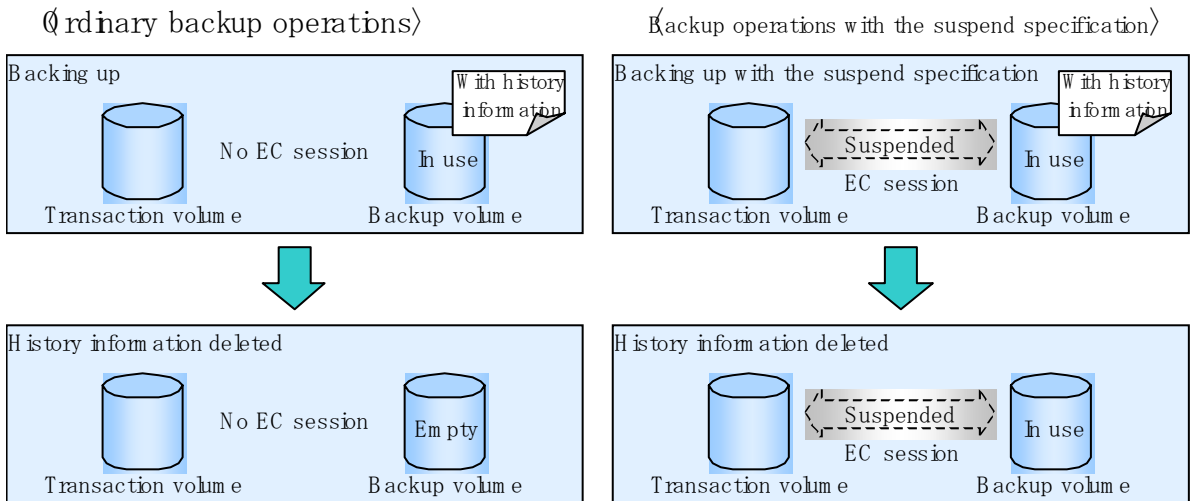
To delete backup history information, execute the history information deletion command (swsthistdel) at a command prompt.



- 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 the device usage status display command (swstdevdisp) at a command prompt. 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 the 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: Deleting ordinary backup history information and backup history information with the suspend specification



5.5.6 Restoring

Before restoration, the synchronous processing cancel command (`swstcancelsync`) at a command prompt 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 the synchronous processing progress display command (`swstsyncstat`) at a command prompt.

To restore data from a backup volume, execute the restore execution command (`swstrestore`) at a command prompt.

To cancel restoration or to delete management information associated with restoration that ended abnormally, execute the restore cancellation command (`swstcancelrest`) at a command prompt.



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 by the QuickOPC Function

This chapter provides details of the backup operation using the QuickOPC function.

Note: This function was previously referred to as Differential OPC.

6.1 Overview

For a snapshot high speed backup using One Point Copy (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, a disk array that supports the QuickOPC function is required.

The QuickOPC function can be used only for an intra-cabinet copy (OPC). It cannot be used for an inter-cabinet copy (ROPC).


The QuickOPC function also cannot be used for backing up SDX objects in units of logical volumes.



Point

When the QuickOPC function is used, the hardware records any 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: 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 by noting the following in addition to the information given in Operation Design for backup management.



Note

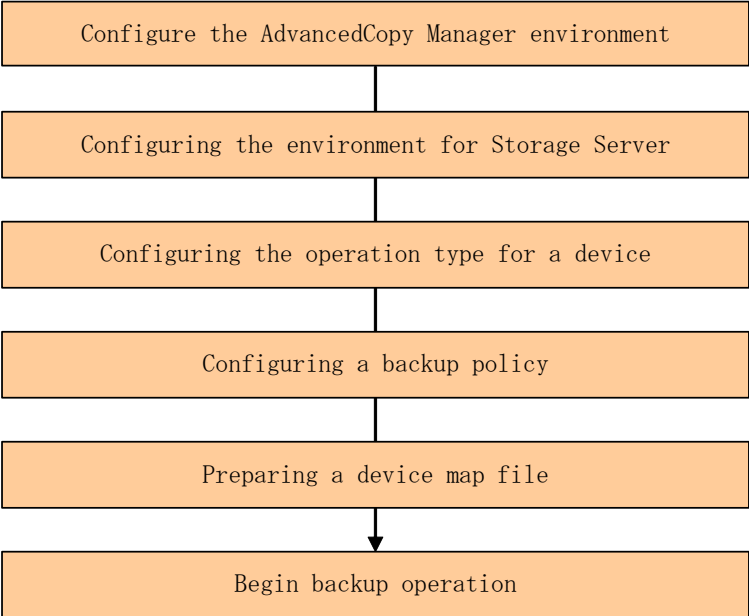
Note the following when making decisions on backup policy configuration and backup volume preparation:

- In a backup operation using the QuickOPC function, an OPC session remains active even after the physical copy has been completed. 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.
- The QuickOPC function can be used only for an intra-cabinet copy (OPC). It cannot be used for an inter-cabinet copy (ROPC). Prepare a backup volume in the same cabinet as the transaction volume to be backed up.

6.3 Operation Flowchart

A flow of operation using the QuickOPC function is shown below.

Figure: Flow of operation using the QuickOPC function



6.4 Preparation

This section provides details of the preparation for a backup operation using the QuickOPC function.

6.4.1 Configuring the AdvancedCopy Manager environment

For the sequence of processes from the installation to the operating environment configuration of AdvancedCopy Manager, see the "ETERNUS SF AdvancedCopy Manager Installation Guide" and the "ETERNUS SF AdvancedCopy Manager Operator's Guide".

6.4.2 Configuring the environment for Storage Server

Set the configuration information of the Storage Server used for backup.

This process can be executed using the GUI and command operations.

For details on the operation using the GUI, refer to "Modify Storage Server Information" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For details on the operation using the command, refer to the Storage Server configuration information configuration command (swstsvrset). It is also possible to use the Storage Server configuration information display command (swstsvrdisp) to refer to the Storage Server configuration information that was set.

6.4.3 Configuring 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 executed using the GUI and command operations.

For details on the operation using the GUI, refer to "Define the Role of a Volume" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For details on registering the transaction volume/backup volume for using the command, refer to the Device information Configuration command (swstdevinfo). It is also possible to use the Device usage status display command (swstdevdisp) to refer to the volume information that was registered.

Volumes in the tracking state can be checked using the tracking status display command (swstrkstat) at a command prompt.



- The configuration of the transaction volume in the tracking state cannot be changed.
- The configuration of the backup volume in the tracking state cannot be changed.



The "-b" and "-u" options can be specified together in the device use status display command to check for backup volumes, whose histories were deleted, in the suspend or tracking state.

6.4.4 Configuring a backup policy

When a backup policy is set, you will 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 executed using the GUI and command operations.

For details on the operation using the GUI, refer to "Set Backup Policy" in the "ETERNUS SF AdvancedCopy Manager User's Guide".

For details on the operation using the command, refer to the following:

- Backup policy settings: Backup policy Configuration command (swstbkpolset)
- Checking the backup policy that was set: Backup policy display command (swstbkpoldisp)
- Deletion of the backup policy: Backup policy deletion command (swstbkpoldel)

6.4.5 Preparing a device map file

In the backup operation, AdvancedCopy Manager automatically selects a device that 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 use a defined backup destination device, 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 provided the device is online or available.



If multiple generations are to be managed, a device map files must be prepared for each generation.

6.5 Operation

6.5.1 Executing backup

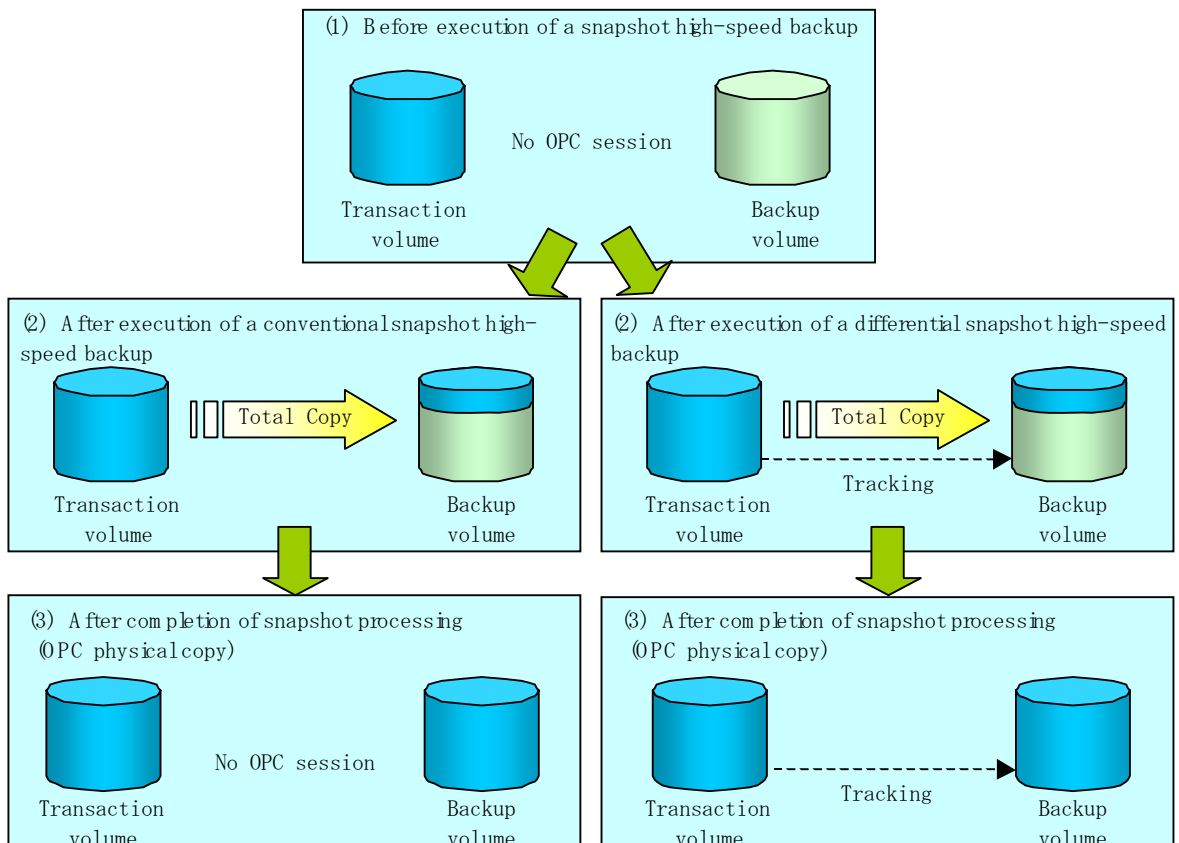
Execute the differential snapshot high speed backup by specifying the "-T" option in the backup execution command (swstbackup) at a command prompt.

Executing the backup command by specifying the "-T" option starts the QuickOPC. When the QuickOPC starts, the physical copy and tracking processing begins. When the physical copy is complete, only tracking processing is active. When the backup command is later executed by specifying the "-T" option, only data that has changes since OPC was started previously is actually copied.

Note that differential snapshot high-speed backup can be operated from the GUI client. Refer to "Perform a Snapshot Backup" in the "ETERNUS SF AdvancedCopy Manager User's Guide" for details.

The figure below shows the conventional snapshot high speed backup and the differential snapshot high speed backup.

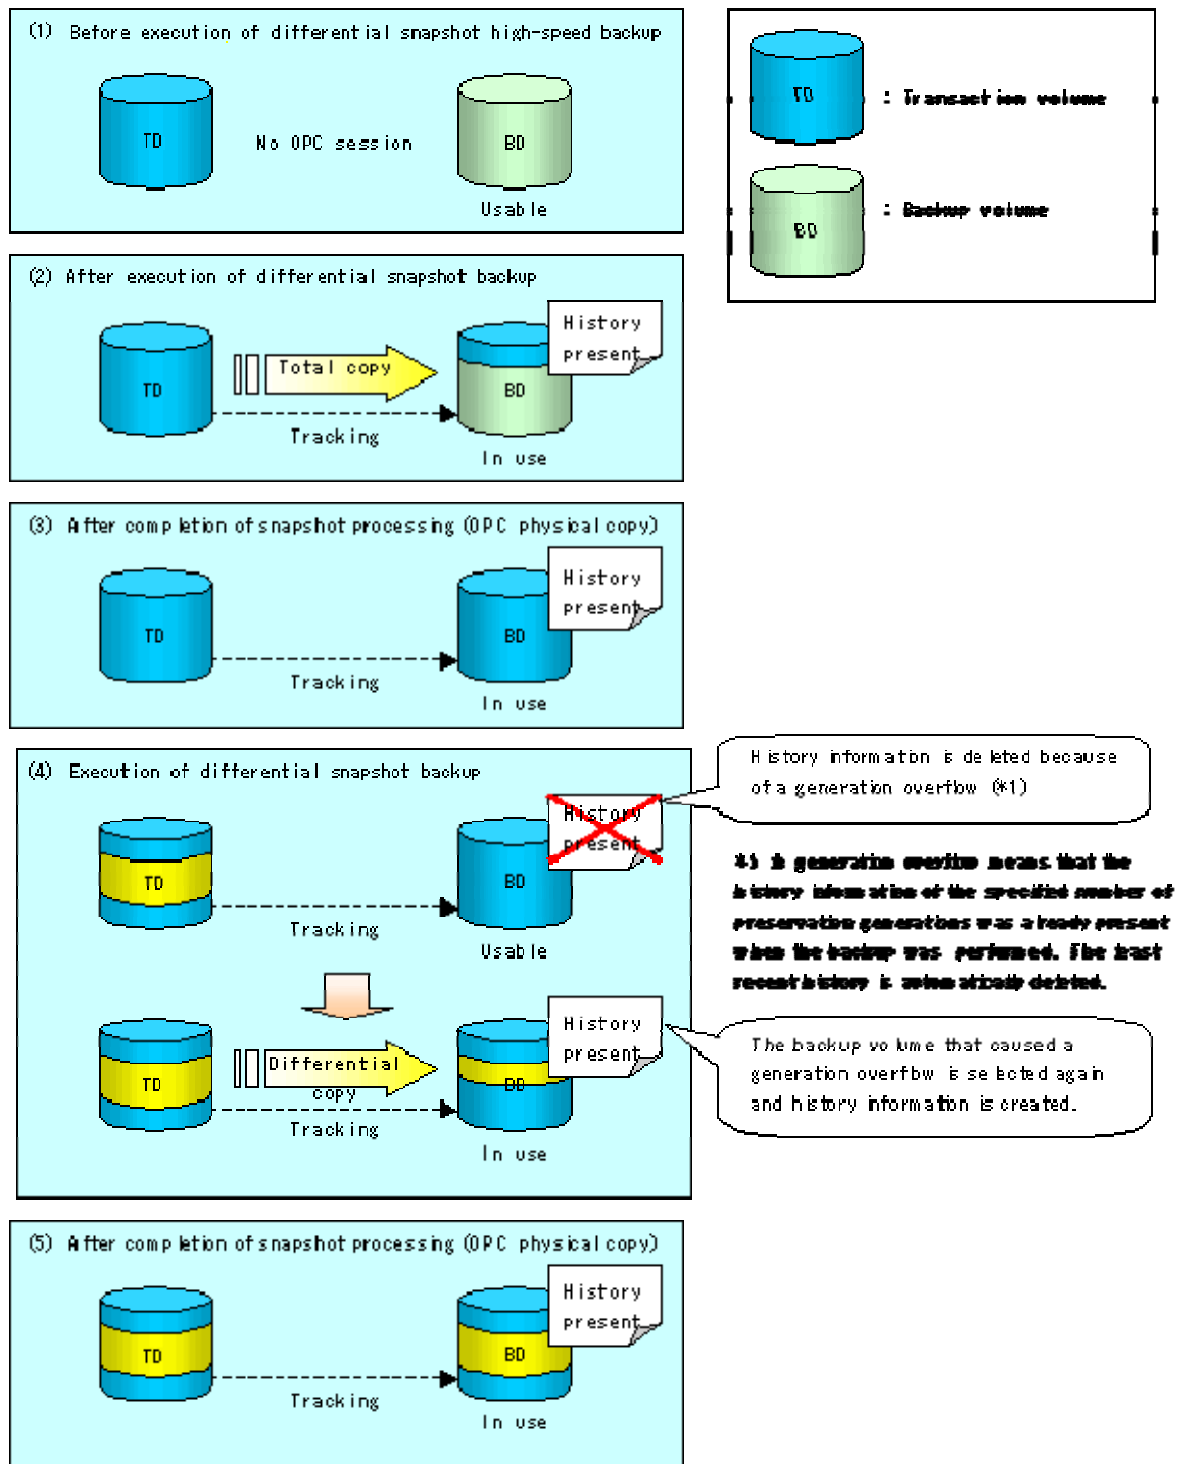
Figure: Conventional snapshot high speed backup and differential snapshot high speed backup



After the 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: Flow of differential snapshot high speed backup operation



6.5.2 History management

Backup history information can be referenced using the history information display command (swsthistdisp) at a command prompt.

Backup history information can be deleted using the history information deletion command (swsthistdel) at a command prompt.

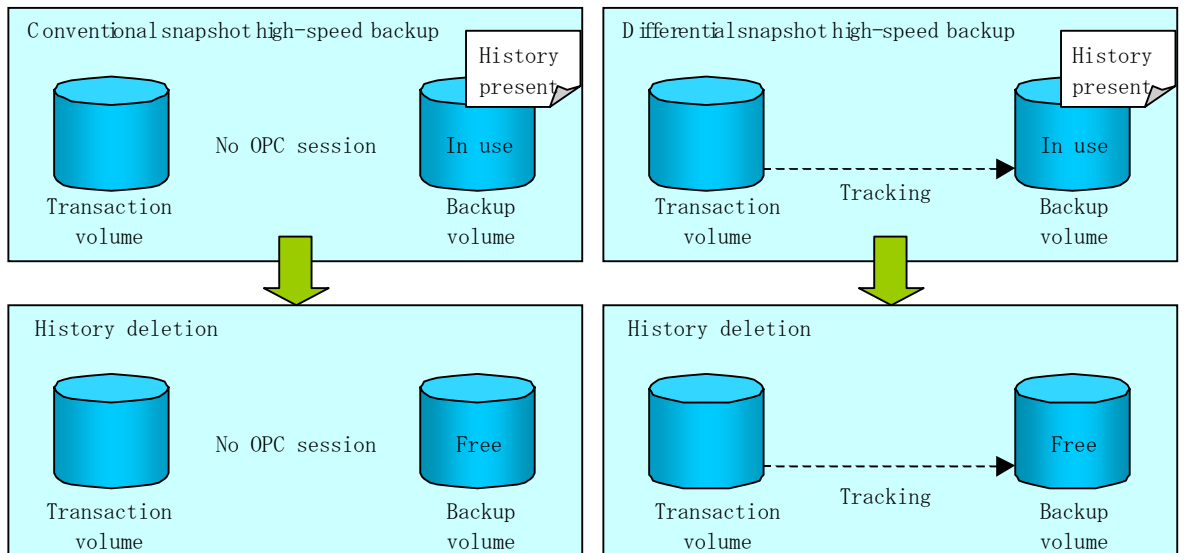
Even if the history information created by the differential snapshot high speed backup is deleted, the tracking state is not released (that is, the physical copy, if it is being executed, is not stopped).

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 column by the device usage status display command (swstdevdisp) at a command prompt. 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 column by the device use 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 column.

The figure below shows a comparison of history deletion between the conventional snapshot high speed backup and the differential snapshot high speed backup.

Figure: Comparison of history deletion between snapshot high speed backup and differential snapshot high speed backup




6.5.3 Restoration

This section explains how to use the restore execution command (swstorestore) for restoration.

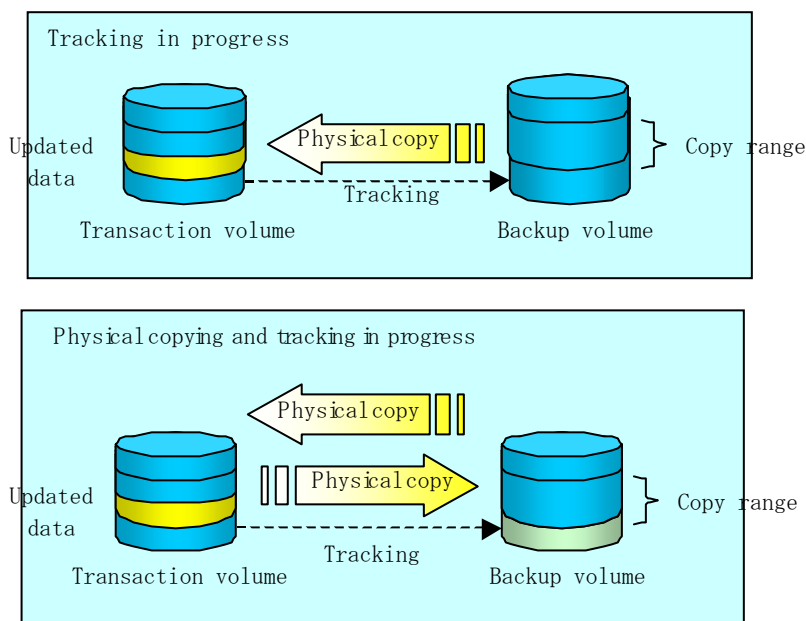
In a differential snapshot high speed backup, tracking processing from the transaction volume to the backup volume is performed.

Figure: Tracking processing for differential 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 most recent backup, the time taken for the physical copy can be reduced for the restoration, as well as for the backup.


Figure: 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 steps must be performed in first:

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 the tracking cancel command (swstcanceltrk) at a command prompt.

Figure: Restoration to a different 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 Canceling restoration

Use the restore cancellation command (swstcancelrest) to cancel the restoration.



Point

Canceling the restoration does not release tracking from the transaction volume to the backup volume.

6.5.5 Displaying the execution status of tracking processing

The execution status of tracking processing by the QuickOPC function can be checked with the tracking status display command (swstrkstat) at a command prompt.

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 (that is, the logical copy) is displayed as a percentage in the Update column.

6.5.6 Stopping tracking processing

Use the tracking cancel command (swstcanceltrk) at a command prompt to stop tracking processing.



If the command is issued while the physical copy and tracking are in progress, the tracking is stopped but the physical copy is not stopped.

Chapter 7 Replication

This chapter provides details of AdvancedCopy Manager replication in an HP-UX system.

The HP-UX version of AdvancedCopy Manager's Agent enables the replication described in this chapter by linking with AdvancedCopy Manager's Manager running under the Windows, Solaris or Linux Operating Systems.

In addition, the replication unit in the HP-UX system is a physical disk (LU: Logical Unit). However, in the case of a physical disk which constitutes an LVM (a Logical Volume Manager), it is managed and operated by each volume group.

7.1 Overview

This chapter explains the operation of the replication function.

Using the advanced copy function (OPC/EC or ROPC/REC function) of an ETERNUS disk array, the replication function performs high speed replication between volumes, regardless of the volume capacities.

Replication that uses the One Point Copy (OPC) or Remote One Point Copy (ROPC) function of a Fujitsu ETERNUS storage system disk array unit is called snapshot replication. Replication that uses the Equivalent Copy (EC) or Remote Equivalent Copy (REC) function is called synchronized replication.

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.


7.1.1 Snapshot replication processing

Snapshot replication creates replicas by copying from a source volume to a replica volume using the OPC or ROPC function of a Fujitsu ETERNUS disk array.

The snapshot replication performs the following two types of processing:

1. When a copy is created, the snapshot processing (OPC or ROPC) from the source volume to the replica volume is executed with the replication creation command (swsrpmake) at a command prompt. The replication creation command unmounts or mounts the source volume. This processing determines the original data (for details, refer to "Pre-processing and Post-processing of Replication"). Refer to steps 1. to 4. in Figure 7.1 below.
2. If the user wants to recreate a replica, only the replication creation command (swsrpmake) is required. If the previous snapshot processing is in progress, it is stopped and a new snapshot process is started.

Figure 7.1 Snapshot replication processing



Snapshot replication is completed when the replication creation command is issued. The actual data is internally copied by the OPC or ROPC function of the Fujitsu ETERNUS storage system disk array unit.



Note


The ETERNUS disk array that supports the ROPC function is required in order to use it in the Advanced Copy snapshot replication function.

The following types of snapshot replication are available, depending on the AdvancedCopy Manager functions used:

- Ordinary snapshot replication
- QuickOPC snapshot replication
- SnapOPC replication
- SnapOPC+ replication

Ordinary snapshot replication

All data at a specific point in time (logical copy) is copied to the copy destination disk area.



The copy destination disk area must be the same or larger than the copy source disk area.


The copy time is the time taken to copy all the data.

QuickOPC Snapshot type replication

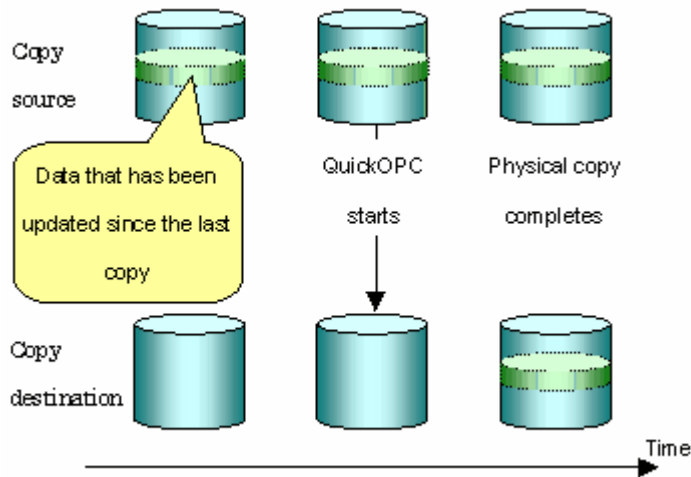
If the ETERNUS disk array supports the QuickOPC function, QuickOPC snapshot replication can be performed.

Initially, all the data at a specific point in time (logical copy) is copied to the copy destination disk area.

Second and subsequent copy operations copy only the data that has been updated since the previous copy operation. This enables great reductions in the time taken for physical copying.



Second and subsequent copies



Point

The copy destination disk area must be the same or larger than the copy source disk area.

The duration for second and subsequent copy operations is the time taken for copying the differential data.



Note

To execute snapshot replication, it is necessary to use an ETERNUS disk array that supports the QuickOPC function.

The QuickOPC function can be used only for copying within a cabinet (OPC). The function cannot be used for copying between cabinets (ROPC).

The QuickOPC function also 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 into the copy source disk area after a certain point in time (a logical copy) to be copied to the destination disk area.



Point

The copy destination disk area can be smaller than the copy source disk area. The copy time is the time needed to copy data that has been updated.




Note

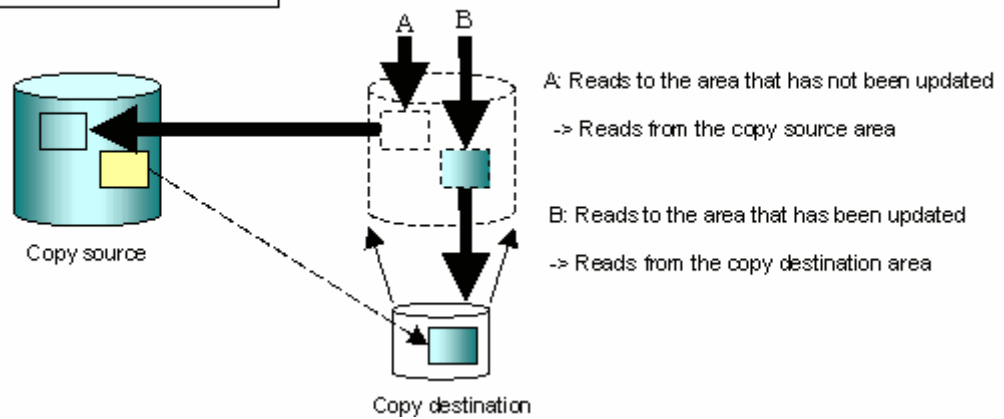
SnapOPC can only be used for replication (OPC) within a cabinet, not for inter-cabinet replication (ROPC).

The copy destination disk for SnapOPC is referred to as a Snap Data Volume.

Mechanism of SnapOPC



Reading from the copy source



Snapshot type replications that use SnapOPC are referred to as "SnapOPC snapshot type replications".

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 the "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.

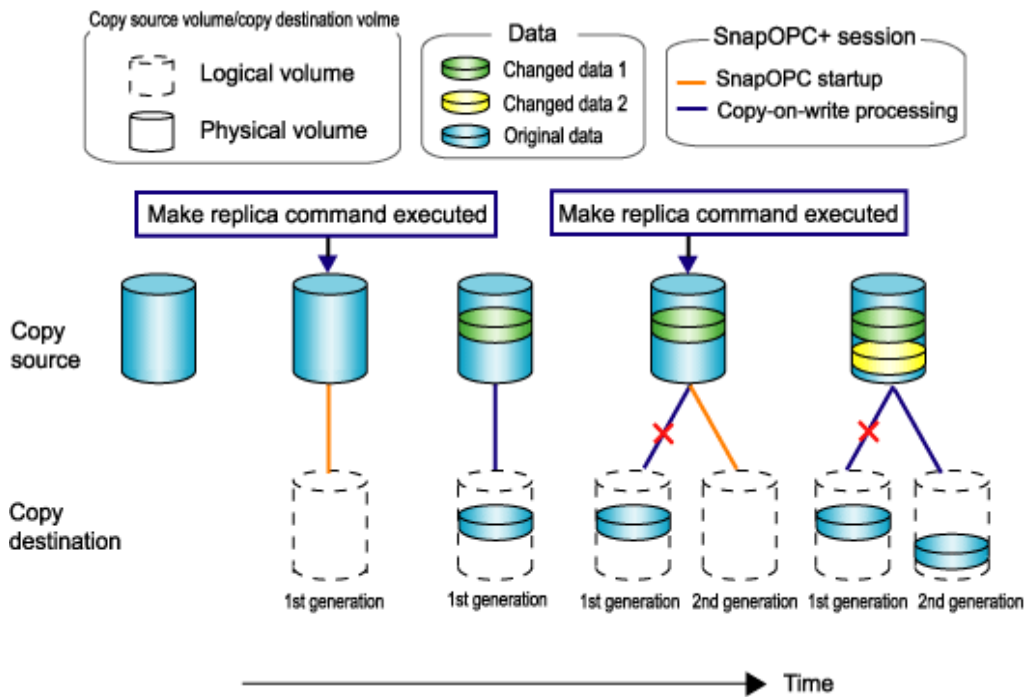
These points suggest that SnapOPC should be used in temporary areas for tape backups, for systems where access performance is not an important consideration.

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 flow of SnapOPC+ processing.

Figure: 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 data updated 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 to 8, 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 pre-update data for the data updated immediately prior and for a hardware management area. 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.




Note

The replication function cannot be used to restore data when SnapOPC+ is used. If restoration is required, the data must be copied manually from the copy destination volume.

The following figure shows how SnapOPC+ is used to perform 1:N replication operations between a copy source volume (1) and multiple snap generations (N: 1 to 8).

Figure: 1:3 replication operations using SnapOPC+



Due to the functionality 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: 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	Y	R 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	N	Y
Maximum number of sessions(*)	8	8
Restoration using the replication function	Y	N
Recommended usage	<ul style="list-style-type: none"> ● Tape backup temporary area ● Backup of file servers or similar that have a small amount of updates 	<ul style="list-style-type: none"> ● Tape backup temporary area ● Backup of file servers or similar that have a small amount of updates

*: The maximum number of sessions is the maximum number of sessions that can be set for one copy source volume.

Y = Yes

N = No

R = Recommended

7.1.1.1 Snap Data Volume extension function

This functionality extends physical capacity of the Snap Data Volume as additional area is allocated from the Snap Data Pool when data is updated and the capacity of the Snap Data Volume is exceeded. This functionality is set in advance.

If there is no need to extend the area of the Snap Data Volume it can be used without setting up a Snap Data Pool.

Snap Data Volume

The Snap OPC/SnapOPC+ copy destination volume is called the “Snap Data Volume”. Of this copy destination volume (Snap Data Volume), the capacity that can be seen in the host is called the “logical capacity”, and the capacity that has been physically configured is called the “physical capacity”.

The size of the logical capacity is equal to or greater than the size of the copy source capacity. Control information is set for the Physical capacity (including a conversion table) to manage the backup data and the storage capacity for backup data. This information requires 0.1% of the logical capacity.



Note

Where there is insufficient capacity in the Snap Data Volume access to the source volume is not possible, refer to “Troubleshooting when an insufficiency of free space has occurred in the Snap Data Volume or Snap Data Pool”.

Snap Data Pool

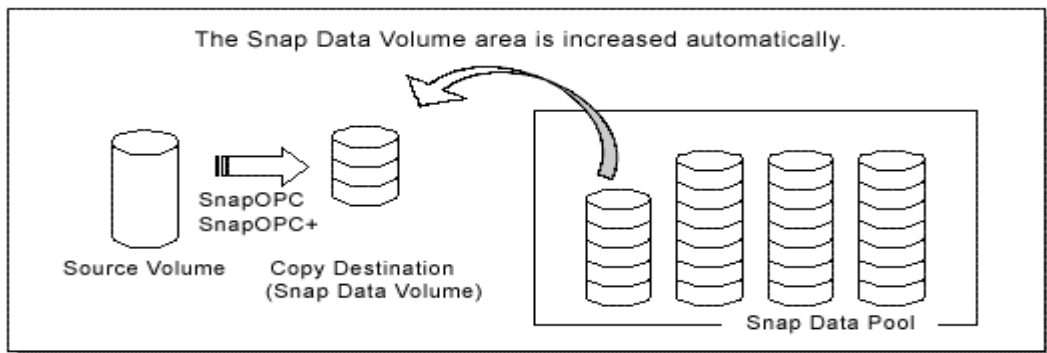
The Snap Data Volume pool function of ETERNUS enables a Snap Data Pool to be set. When there is insufficient capacity in the Snap Data Volume and a Snap Data Pool has been set up, area is added from the Snap Data Pool automatically.



Note

When there is insufficient capacity in the Snap Data Pool access to the copy destination (Snap Data Volume) is not possible, refer to “Troubleshooting when an insufficiency of free space has occurred in the Snap Data Volume or Snap Data Pool” .

Figure: Snap Data Volume extension image



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 a Fujitsu ETERNUS storage system disk array unit.

Synchronized replication uses two methods to create a copy: full copying and incremental copying. In the case of full copying, 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 subsequent copies once a previous replica already exists. Copies are created by reflecting the data that has been updated after a previous has been made (full copying and incremental copying).

The following steps define the procedure for synchronized replication:

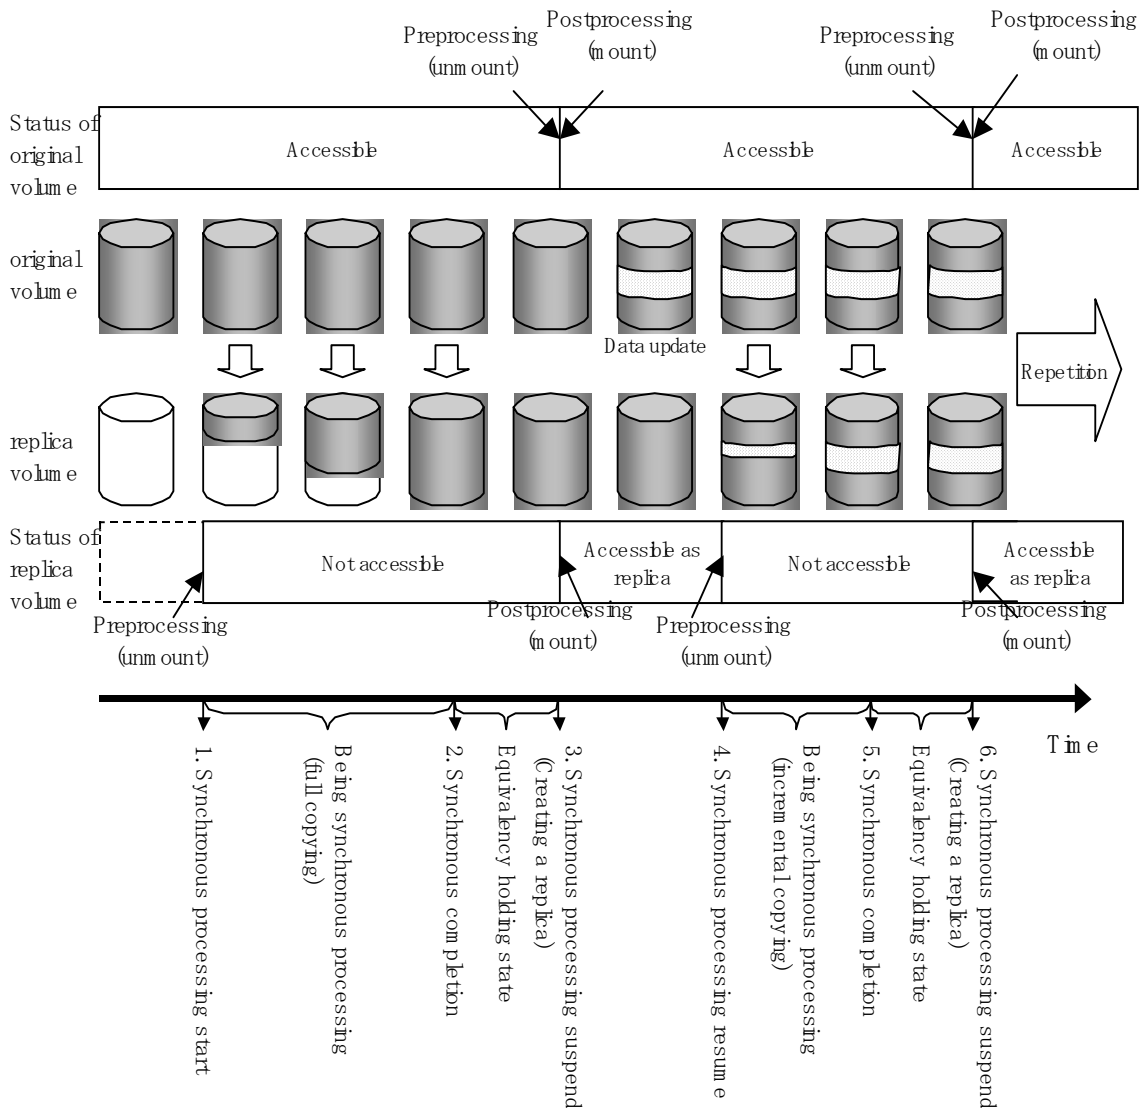
1. Full copying starts at the start of synchronized processing from the source volume to the replica volume, using the synchronous processing start command (swsrpstartsync) at a command prompt. (refer to step (1) in Figure 7.2)
2. Full copying is complete. The source volume and replica volume both enter the equivalency maintenance status. From this point on, updating of the source volume is sequentially reflected in the replica volume, so that data equivalence is always maintained (this is the equivalency maintenance state). (Refer to the states in steps (2) and (3) in Figure 7.2)
3. Synchronized processing is temporarily stopped with the replication creation command (swsrpmake) to create a copy of the source volume. The replication creation command unmounts or mounts the source volume. This processing determines the original data (for details, refer to "Pre-processing and Post-processing of Replication"). After execution of the replication creation command, the replica volume can be accessed (this is the copy established status). Refer to the statuses in steps (3) and (4) in Figure 7.2.

4. To create another copy, the synchronized processing that has been temporarily stopped is restarted with the synchronous processing start command. In the case of this copy, only the updated (ie the incremental) data on the source volume since the previous copy was made until the restart of synchronized processing is copied to the replica volume. Refer to step (4) in Figure 7.2.

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 copying of the incremental data has been completed, the status returns to equivalency maintenance. Refer to the states in steps (5) and (6) in Figure 7.2.
6. When this status has been set, the copy is recreated with the replication creation command. Refer to step (6) in Figure 7.2.

Figure 7.2 Synchronized replication processing



Note

- When using the EC or REC function in synchronized replication, disk array devices must support the EC or REC function.
- The replication creation command (swsrpmake) cannot be executed unless both the source volume and replica volume are in equivalency status.



Note

If the backup volume is part of a volume group, do not execute the `vgchange (1M)` command under synchronous processing (EC/REC) execution. If the backup volume becomes deactivated during the execution of synchronous processing, it is not only impossible to make it active again but it will be in a state such that operations including stopping synchronous processing, the backup execution command, and the backup synchronous processing cancel command, cannot be performed.

When you have accidentally deactivated the backup volume, reactivate it after performing a forced stop of the synchronous processing using `GRmgr` etc.

7.2 Operational configuration

This section explains the operation of replication.

Configure replication with the following procedure:

- Determining the server that performs replication.
- Determining the replication target.

Refer to "Notes on the configuration of SnapOPC/SnapOPC+ replication", when the SnapOPC/SnapOPC+ replication is used.

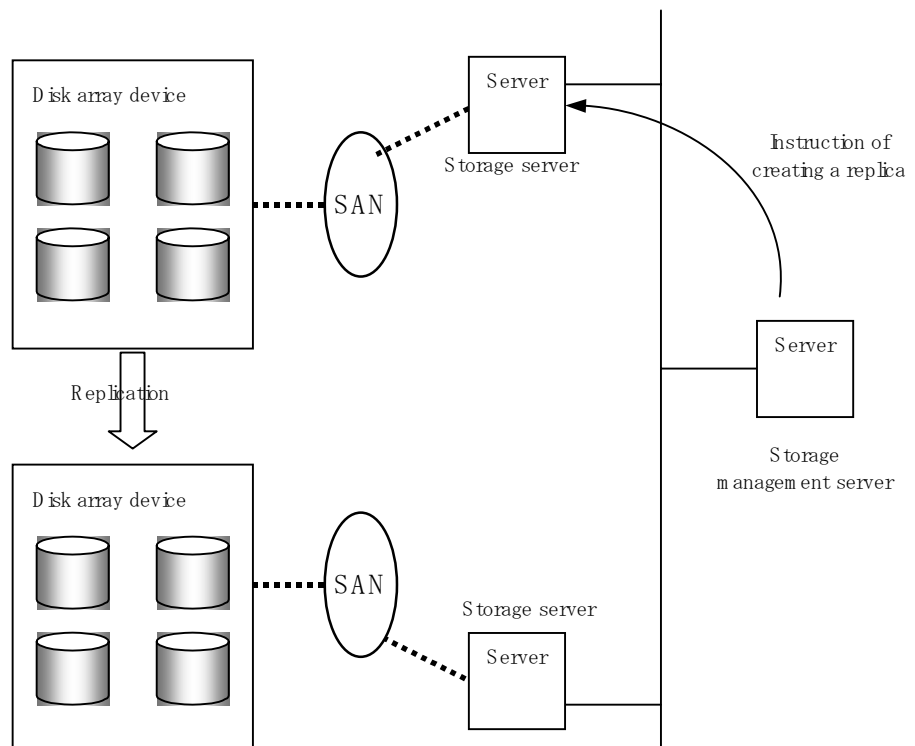
7.2.1 Determining the server that performs replication

Determine the server that performs replication.

The following servers are used in the replication process:

- Storage Management Server
Multiple Storage Servers are unified and operated centrally. The Storage Management Server can also be used as a Storage Server.
However, a Storage Management Server requires AdvancedCopy Manager to run under the Windows, Solaris or Linux operating systems.
- Storage Server
AdvancedCopy Manager operations are performed.
AdvancedCopy Manager can run under the HP-UX OS in this case.

Figure 7.3 Configuration of the servers that perform replication





Note

The following requirements apply when system to system replication is performed using the remote copying function (either ROPC or REC) of a Fujitsu ETERNUS storage system disk array unit:

- The ROPC or REC function must be installed in both systems.
- Both 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 INIT side to the TARG side. For bi-direction 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.



Point

- 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 the source and replica volumes:

1. Copy direction

Determine the direction of copying. Specify one of the following directions:

- bi-directional: Copying is bidirectional. Copying from the replication destination to the replication source can be used for restore.
- uni-directional: Copying is performed only from the source volume to the replica volume. During this operation, copying from the replica to the source volume is disabled.

2. Operational servers (used only for server-to-server replication)

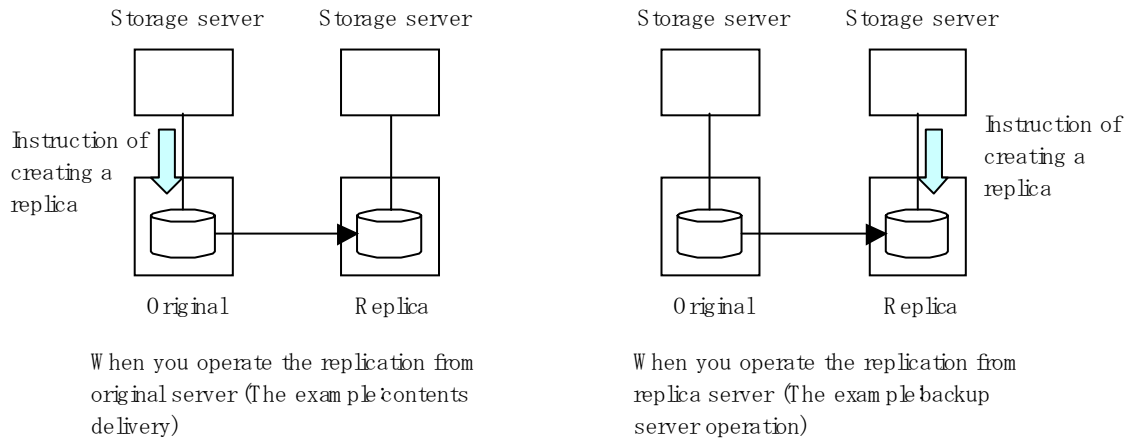
In the case of replication between servers, you can specify whether to allow replication to be performed to either the target source volume or replica volume, on either the copy source server or copy destination server, or on both servers, as follows:

- "Original server": Sets the Storage Server with the source volume that was connected as the operational server.
- "Replica server": Sets the Storage 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 "Creating groups".

Figure 7.4 Servers used in performing replication



In this manual, a Storage Server used to perform replication is called the operational server. An original server or a replica server is specified as the operational server. An operational server can use all of the replication functions. In contrast, any other type of server can use only the information display function and operational release function.



Replication is not supported for the system disk or the disk where AdvancedCopy Manager is installed.



Refer to "General notes" for notes on determining a replication target.

7.2.2.1 Notes on the replication volume

There are points to be careful of for the different combinations of source volume and replica volume that can be used for replication.

The disk size of a source volume and a replica volume

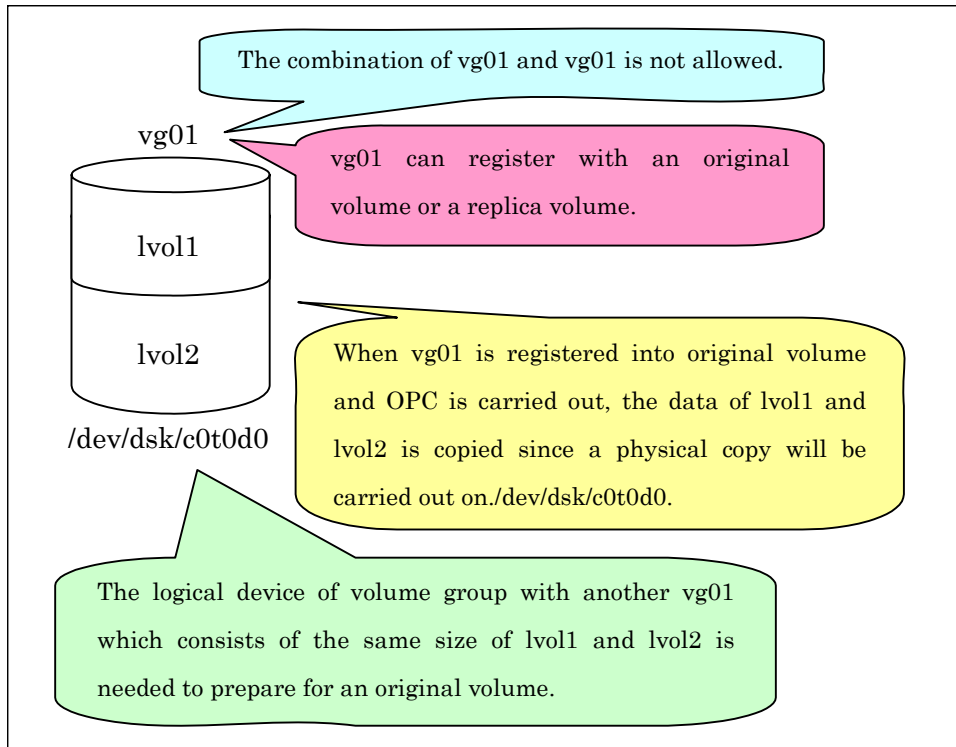
When the size of the physical disk (or one which constitutes a logical disk when this has been specified) of an source volume differs from that of a replica volume, ensure that the replica volume is at least as large as the source volume, otherwise some data will be lost in the replication.

Using a volume group for a source volume or a replica volume

When the source volume is a volume group, then when using the replica volume after replication, the replica volume needs to be a volume group of the same logical volume configuration as the source volume, and with the same physical disk size (refer to Figure 7.5).

In addition, when the configuration of the volume group cannot be supported by AdvancedCopy Manager, it cannot be registered with a replication management function. For details about logical disk configurations which cannot be supported, refer to "Managing a Device on AdvancedCopy Manager".

Figure 7.5 Source volume and replica volume for two or more logical disks

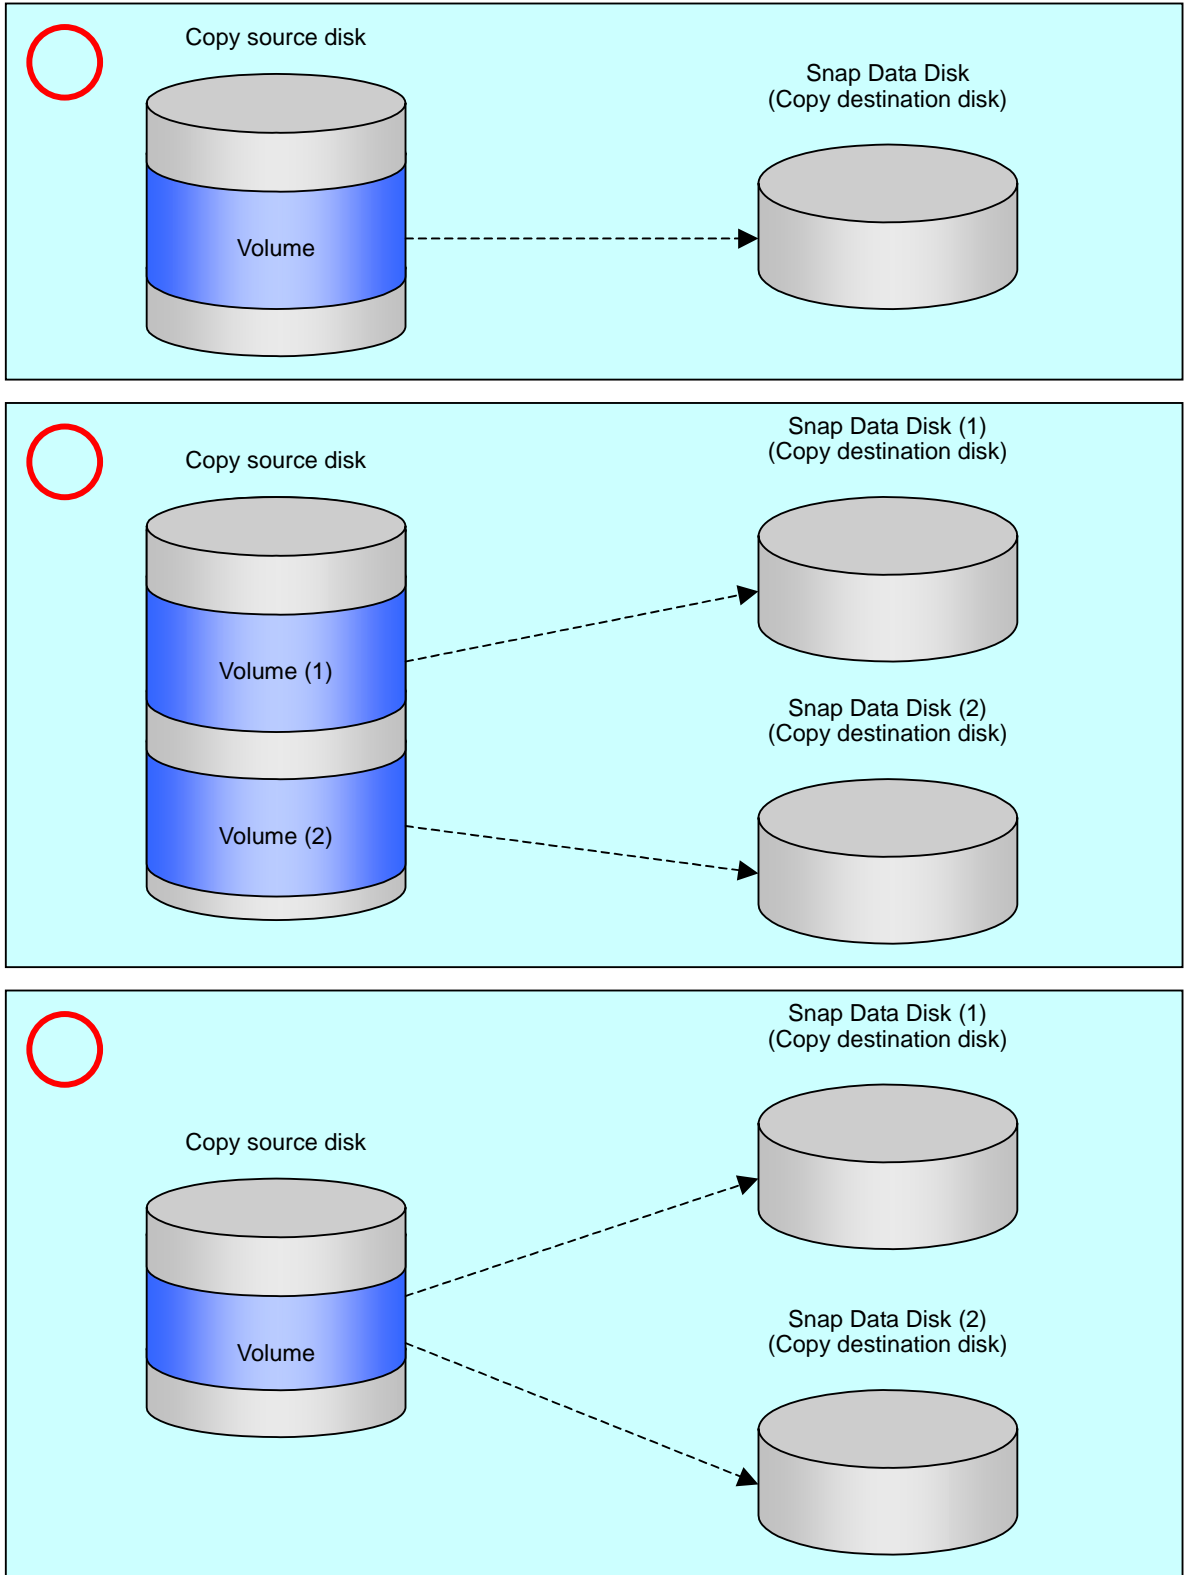


7.2.3 Notes on the configuration of SnapOPC/SnapOPC+ replication.

7.2.3.1 Notes on session configurations



Only one session can be set up for each Snap Data Volume, the SnapOPC/SnapOPC+ copy destination disk.

Figure: When one session is set for a Snap Data Volume



Accordingly, multiple sessions cannot be set up for a single Snap Data Volume, as shown in the following figure:


Figure: When multiple sessions are set for a Snap Data Volume



The following restrictions apply:

- It is not possible to copy from the Snap Data Volume to a disk other than the copy source disk while SnapOPC is running. It is not possible to copy from the Snap Data Volume to a different disk while SnapOPC+ is running.

Figure: When copying from a Snap Data Volume to a disk other than the copy source



7.2.3.2 Creating a Snap Data Disk

Create a Snap Data Volume by performing the following steps:

7. Calculate the physical capacity for the Snap Data Volume.
8. Define and initialize the Snap Data Volume.
9. Connect to the host.
10. Create partitions and file systems.

Calculate the physical capacity for the Snap Data Volume.

The formula for estimating the physical capacity of Snap Data Volumes is as follows:

$$\text{Physical capacity} = (\text{number of updated blocks for the copy source volume}) \times (\text{safety factor})$$

Strictly speaking, both the number of blocks that have been updated on the Snap Data Volume before the SnapOPC is taken and the management area used by the hardware (about 0.1% of the logical capacity) should also be taken into account, 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 the update amount measurement command (`swstestupdate`).

Make this measurement using the following procedure:

- 1) Start measuring the update amount by configuring up a pseudo SnapOPC session on the copy source volume.

```
# /opt/FJSVswsts/bin/swstestupdate start /dev/dsk/c1t1d1
/dev/dsk/c1t1d1 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/dsk/c1t1d1
Volume-Name      Update
/dev/dsk/c1t1d1 644333
#
```

- 4) After the measurements are complete, cancel the pseudo SnapOPC session.

```
# /opt/FJSVswsts/bin/swstestupdate stop /dev/dsk/c1t1d1
/dev/dsk/c1t1d1 swstestupdate completed
#
```

Define and initialize the Snap Data Volume.

Define and initialize the Snap Data Volume using ETERNUSmgr. 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 system (the Server Connection Guide) for information on this procedure.

Create partitions and file systems.

Create partitions, slices, 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 the updates to the copy destination volume as much as possible, to avoid filling up the Snap Data Volume.

7.2.3.3 Creating a Snap Data Pool

Use ETERNUSmgr to create a Snap Data Pool.

Note the following points when creating a Snap Data Pool.



Note

- In principle, the copy capacity should be a suitable estimate, and operations should be such that copying can be completed within the calculated Snap Data Volume capacity. Endeavour to design operations such that area is not normally added to the Snap Data Volume from the Snap Data Pool when copying, and the Snap Data Pool is just used as a supplementary area wherever possible.
- Estimate the capacity of updates for each copy source volume and determine the Snap Data Volume capacity and the Snap Data Pool capacity. If it is difficult to estimate the capacity in advance, a value of 30 to 50% of the total copy source capacity is recommended for the Snap Data Volume capacity plus the Snap Data Pool capacity. (This is only an estimate, and the Snap Data Pool capacity must be modified to suit operating conditions.)

7.2.3.4 Snap Data Volume and Snap Data Pool monitoring

When there is insufficient capacity for the Snap Data Volume and the Snap Data Pool, access to the copy destination volume is not possible. That is, in this state it is impossible to read from or write to the copy destination volume. With SnapOPC+ it is not possible to write to the copy destination or to previous generations.

For this reason, monitoring must be performed to ensure the capacity of both the Snap Data Volume and Snap Data Pool is sufficient at all times during normal operations.

- Using a Snap Data Volume without a Snap Data Pool
Specify "stat" as the subcommand for the Snap Data Volume operation/reference command (swstsdv) to regularly monitor the capacity of the Snap Data Volume.
- Using a Snap Data Pool
Specify "poolstat" as the subcommand for the Snap Data Volume operation/reference command (swstsdv) to perform regular monitoring to check whether a Snap Data Pool capacity shortage has occurred.
If the usage rate of an area of the Snap Data Pool exceeds 50%, regardless of whether the data is encrypted or not, consider increasing the capacity and also the frequency of monitoring the capacity.
If the rate exceeds 70% increase the capacity immediately.

If a capacity shortage occurs for the Snap Data Volume or the Snap Data Pool, refer to "Troubleshooting when an insufficiency of free space has occurred in the Snap Data Volume or Snap Data Pool".

Creating an access path

Where there is insufficient capacity in the Snap Data Volume or Snap Data Pool access to the copy destination (Snap Data Volume) is not possible.

Access to the Snap Data Volume is also impossible in the following circumstances:

- After the Snap Data Volume is initialized
- After disk management information is initialized

Therefore, it is recommended that you reserve a volume (access path) that is separate from the Snap Data Volume.

Create the volume to be used as an access path on the same ETERNUS disk array as the Snap Data Volume. Since this is not used for the storage of data the capacity is not important and any size can be specified.

The following operations are possible using the access path:

- Snap Data Volume initialization
- Snap Data Volume status display
- Snap Data Pool status display

7.2.3.5 Cluster Operations

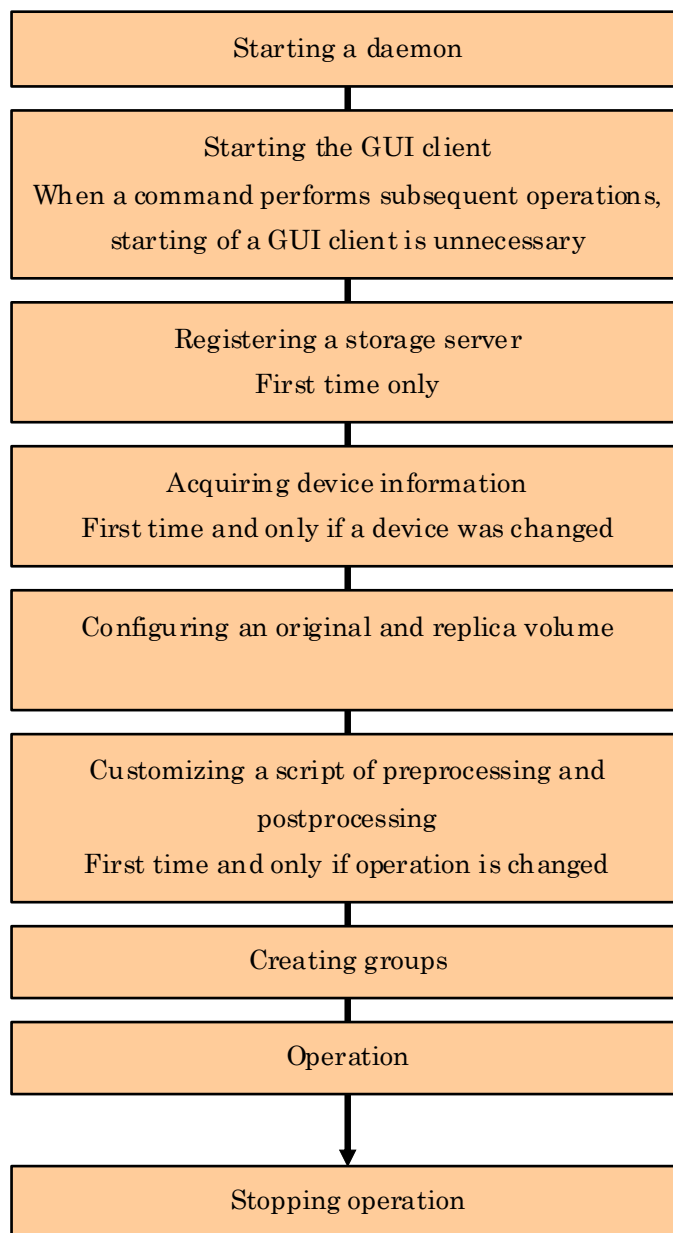
Do not use Snap Data Volumes as shared disks in cluster systems. This is to avoid the danger of the cluster system failing over when the physical capacity shortage occurs for the Snap Data Volume.

In cluster operations, use one of the following methods to prevent the Snap Data Volume from being used as a shared disk:

- a. Permit the Snap Data Volume to be referenced from all the nodes in the cluster system.
- b. Use inter-server replication between the cluster system and a non-cluster system.

7.3 Flow of Operations

The following figure shows the flow of the replication operations.



7.4 Preparations

The following preparations are required before replication can start.

7.4.1 Starting a daemon

Before replication is started, the daemon must be started on the Storage Management Server and the Storage Server. In general, the daemon is automatically started at system startup.

However, if the daemon does not start correctly for any reason, or if a daemon has stopped, it is necessary to start a daemon for each server.

For information about starting Daemons, refer to "Starting and Stopping Daemons".

7.4.2 Starting the GUI client

To start the GUI client refer to "Starting the AdvancedCopy Manager Workbench" in the "ETERNUS SF AdvancedCopy Manager User's Guide" for details.

The GUI client is not used for command-only operations.

7.4.3 Registering a Storage Server

Use the Storage Management Server to register the Storage Server that is to be managed. When the Storage Server is operated as a Storage Management Server it is not necessary to register a server.

The Storage Server can be registered using the GUI or command operations.

For information on using the GUI, refer to "Register a new Storage Server," in the "ETERNUS SF AdvancedCopy Manager User's Guide."

For information on using commands, refer to the server information addition command (stgxfwcmaddsrv) in the "ETERNUS SF AdvancedCopy Manager Operator's Guide" of the OS on which the Storage Management Server is running.

7.4.4 Fetching device information from a Storage Server

This processing can be performed using either GUI operations or command operations.

For information on using the GUI, refer to "Manage Devices," in the "ETERNUS SF AdvancedCopy Manager User's Guide."

For information on using commands, refer to the device information collection/reflection command (stgxfwcmsetdev) in the "ETERNUS SF AdvancedCopy Manager Operator's Guide" of the OS on which the Storage Management Server is running.



Note

Device information on a Storage Server must be temporarily stored to the repository.

The area required when updating the repository may not be available when information concerning devices under multiple Storage 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 the edition of the following manual that is applicable to the OS of the Storage Management Server for the procedure:

- Step 2, "Response to repository access failure" in "ETERNUS SF AdvancedCopy Manager Operation Manual."



Note

- The time required for this operation is proportional to the total number of devices defined on the Storage Server. If the number of devices is large, perform this operation while the CPU load is low. As a guideline, each operation takes about 0.5 seconds per device (or partition) under no-load conditions.
- In relation to a device that is part of an LVM, a volume group and a logical device are fetched, but the physical device which constitutes a logical device is not.

7.4.5 Configuring the source volume and replica volume

Use the replication volume information configuration command (`swsrpssetvol`) 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 the replication volume information display command (`swsrpsvolinfo`).

When specifying the volume of another Storage Server in server-to-server replication, specify the volume in the format "volume-name@storage-server-name."

7.4.6 Customizing a script of Pre-processing and Post-processing

Replication processing must be performed while volumes are unmounted. Therefore, when replication is executed, mount and unmount processing of the volume is performed by using the pre-processing and post-processing script.

The replication processing cannot be executed when the volume cannot be unmounted.

For details of the pre-processing and post-processing scripts, refer to "Pre-processing and Post-processing of Replication."

When the operation corresponds to any of the following, it is necessary to customize the scripts that are used for pre-processing and post-processing of replication:

- When replication target is a volume group.
- For operational reasons, avoid unmount and mount processing of a transaction volume in which the file system has been built.
- In order to describe special processing in pre-processing and post-processing scripts.

For information about the customization method, refer to "Pre-processing and Post-processing of Replication".



Point

When AdvancedCopy Manager is upgraded

The script may be updated after an upgrade of AdvancedCopy Manager.

Therefore, you should re-customize the script after upgrade rather than reusing the script from the previous version.



The pre-processing and post-processing are not performed in case of replicating a replication group.

Unmount all the volumes in the replication group before replication, and mount all the volumes after replication.

It is also necessary to deactivate a volume group (LVM) before it is used as a replication destination. It is then necessary to reconfigure the volume group after the copy is created.

Refer to "The reconfiguration of the volume group" in "Replication in units of physical disks" for information on the reconfiguration of a volume group.

It is also necessary to configure a disk group before a VxVM volume is used as a replication destination. It is necessary to reconfigure the disk group after the copy is created.

Refer to "The reconfiguration of the disk group" of "VxVM Volume Operation" for information on the usage and the reconfiguration of a disk group.

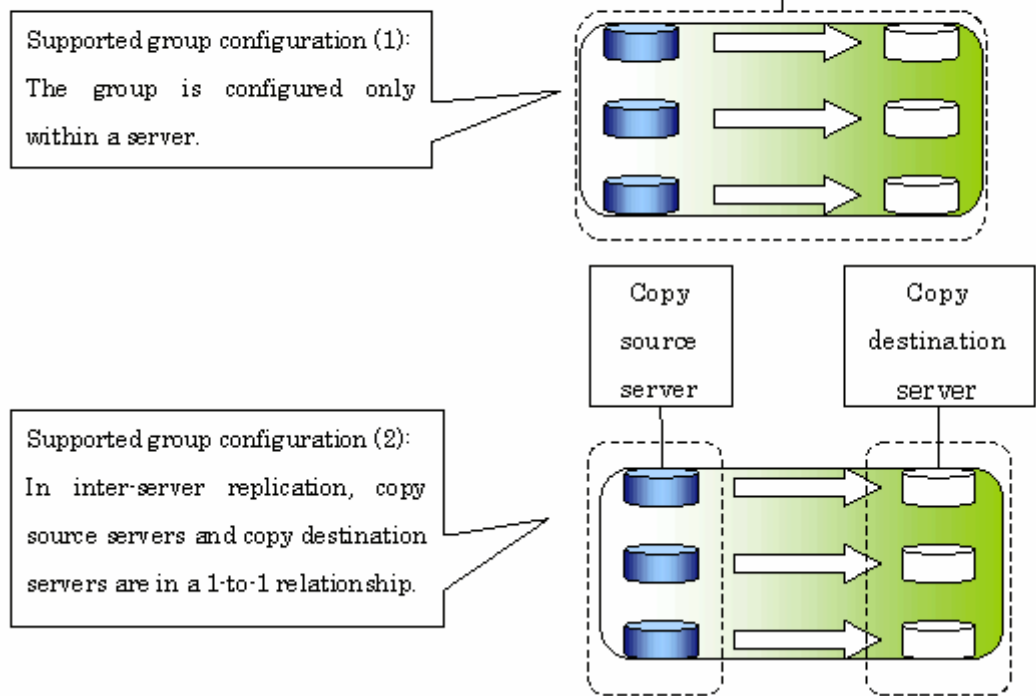
7.4.7 Creating groups

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:

1. The source servers and destination servers must all match.
2. 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" columns obtained by executing the replication volume information display command (swsrpvoinfo).
3. The replication volume information being registered must not be registered in any other group.
4. Source volumes and destination volumes must not be duplicated within the group.

[Condition 1]




[Condition 2]




[Condition 3]

Supported group configuration:
Either the copy sources or destinations are duplicated




Unsupported group configuration:
The copy sources and destinations are duplicated with another group.




[Condition 4]

Unsupported group configuration (1):
The replication source is duplicated.



Unsupported group configuration (2):
The replication destination is duplicated.



7.4.7.2 How to create groups

Groups are created using the replication volume information configuration **command** (**swsrpsetvol**).

Information for groups that have been created can be displayed using **the** replication volume information display command (**swsrpvoinfo**).

[Example]

Two groups (called GRP1) consisting of an source volume and a replica volume are created.

```
# /opt/FJSVswsrp/bin/swsrpsetvol -Xgroup GRP1 /dev/dsk/c1t2d10@SRC /dev/dsk/c1t2d20@TARG-1
swsrpsetvol completed
# /opt/FJSVswsrp/bin/swsrpsetvol -Xgroup GRP1 /dev/dsk/c1t2d11@SRC /dev/dsk/c1t2d21@TARG-1
swsrpsetvol completed
# /opt/FJSVswsrp/bin/swsrpvoinfo - L
Server Original-Volume      Size      Replica-Volume      Size      Copy
Op-Server Group
SRC      /dev/dsk/c1t2d10@SRC 4.0Gbyte /dev/dsk/c1t2d20@TARG-1 4.0Gbyte bi-direction both
GRP1
SRC      /dev/dsk/c1t2d11@SRC 4.0Gbyte /dev/dsk/c1t2d21@TARG-1 4.0Gbyte bi-direction both
GRP1
#
```

7.5 Operations

Before performing replication, refer to "Preparations," to set up the environment required for replication.

This section explains the operation method by using commands. Refer to 'Perform Replication' of "ETERNUS SF AdvancedCopy Manager User's Guide" for information about operation by using the GUI client.

7.5.1 Executing snapshot replication

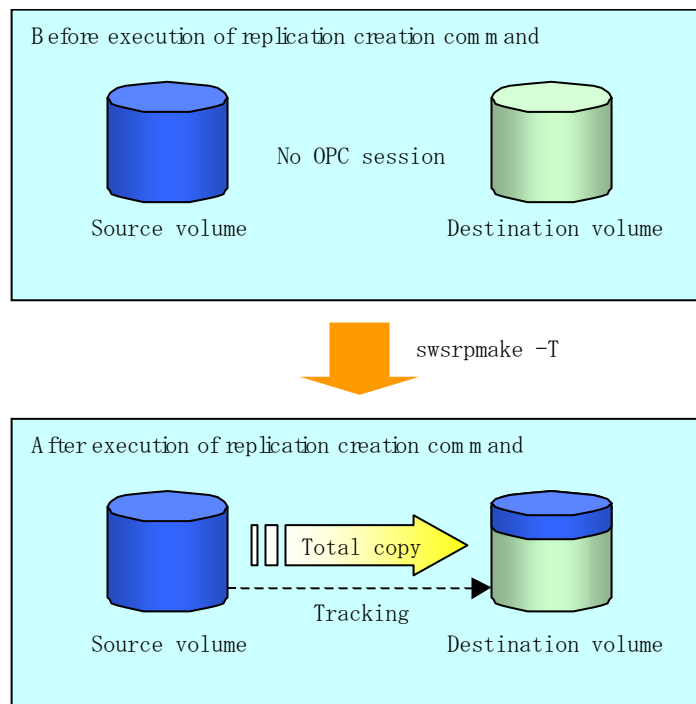
Use the replication creation command (`swsrpmake`) to perform snapshot replication. Refer to "Snapshot replication processing" for an explanation of snapshot replication.

The operation status of a physical copy can be checked by executing the operation status display command (`swsrpstat`).

QuickOPC snapshot replication

Execute QuickOPC snapshot replication by specifying the `-T` option in the replication creation command (`swsrpmake`).

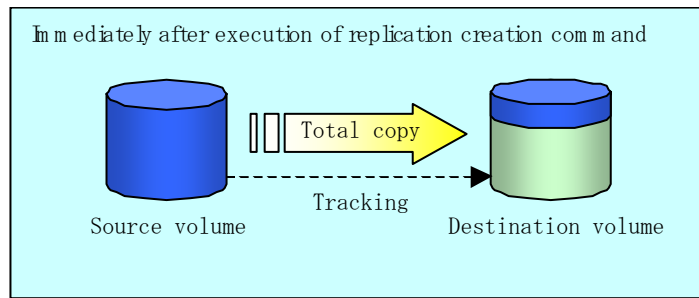
If no OPC session exists when the replication creation command is executed, the command starts snapshot processing (i.e., OPC physical copying) and tracks the processing from the source volume to the destination volume.



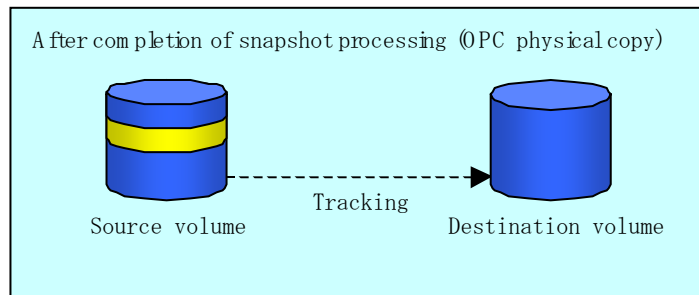
To check the execution status of physical copying, use the operation status display command (`swsrpstat`) in the same way as for an ordinary snapshot replication.

After snapshot processing (OPC physical copy) is complete, only the tracking processing is active.


To check the tracking status, use the operation status display command (`swsrpstat`) with the `-L` option specified.



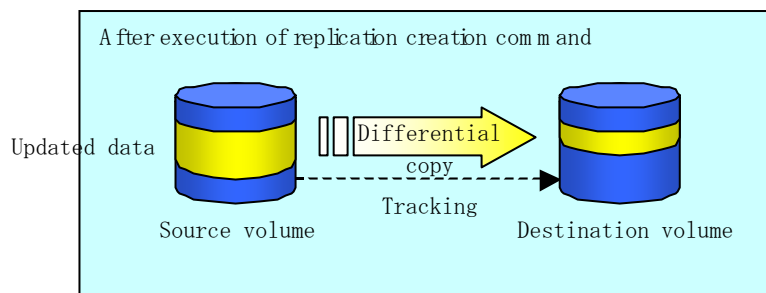
Completion of snapshot processing (OPC physical copy)



Entering the replication creation command (swsrpmake) 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.



swsrpmake -T



If you need to perform a restoration while tracking processing is being executed, you must perform a restoration by OPC (to achieve this, execute swsrpmake without the -T option)

The replication using QuickOPC is performed 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' column 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.

To check the restore execution status, execute the operation status display command (swsrpstat) with the -E option specified.

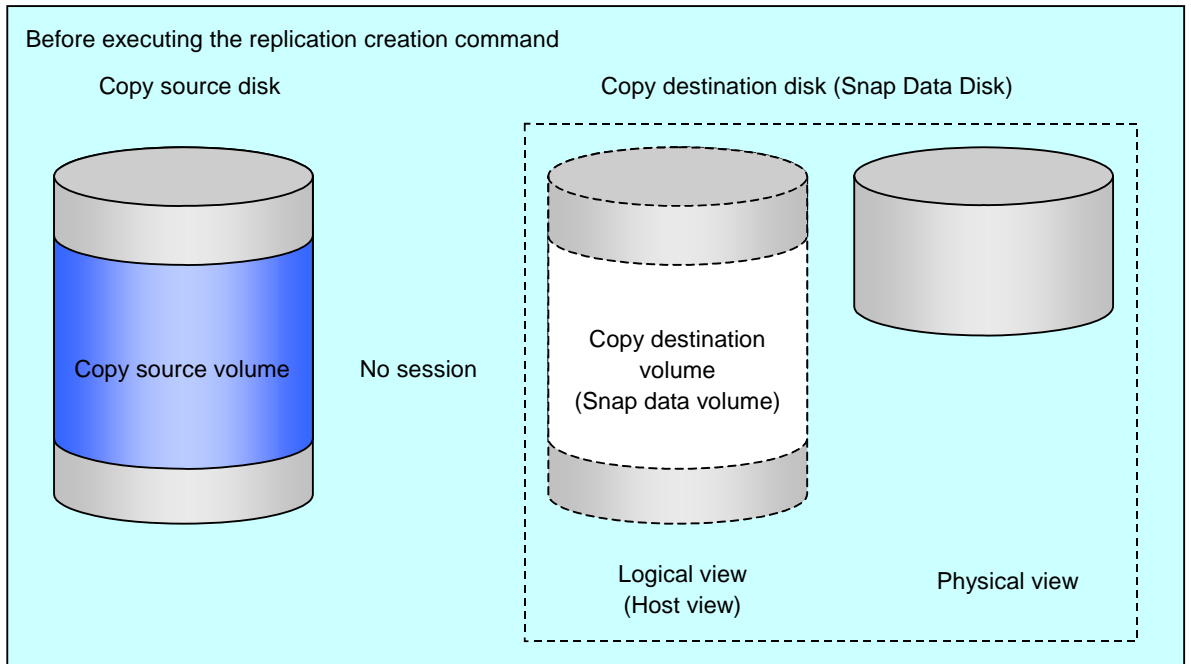
SnapOPC type replication

Execute SnapOPC type replications with the -C option specified in the replication creation command (swsrpmake).

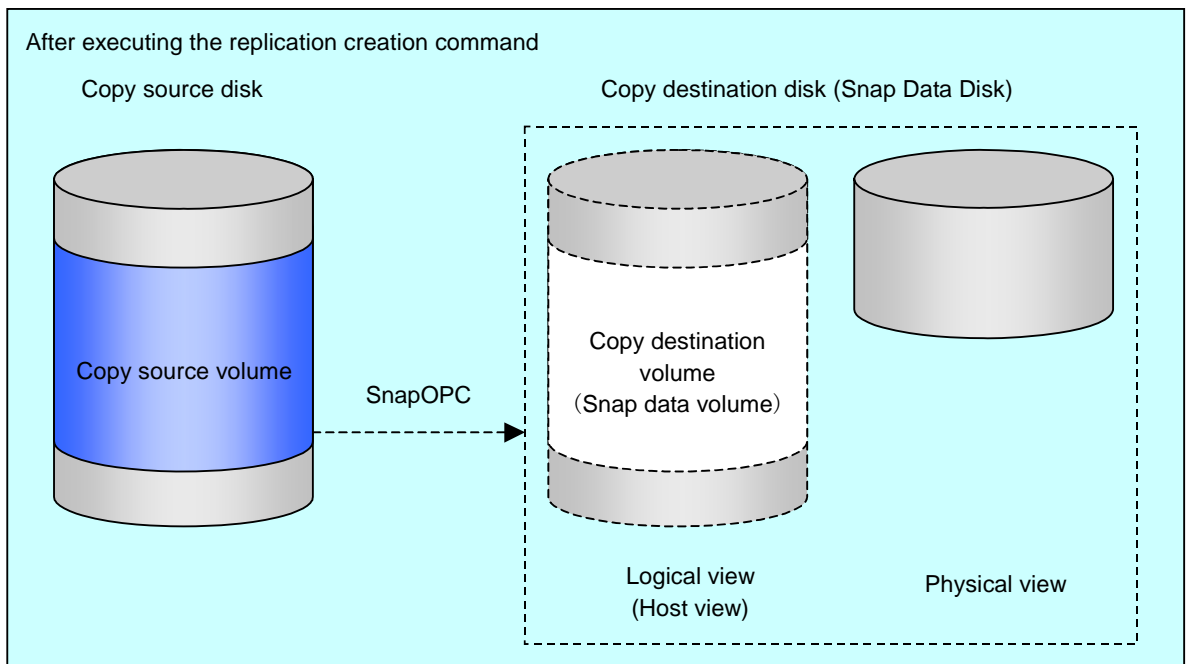
When the replication creation command is executed, a SnapOPC session will be set up between the copy source volume and the copy destination volume.

[Example]


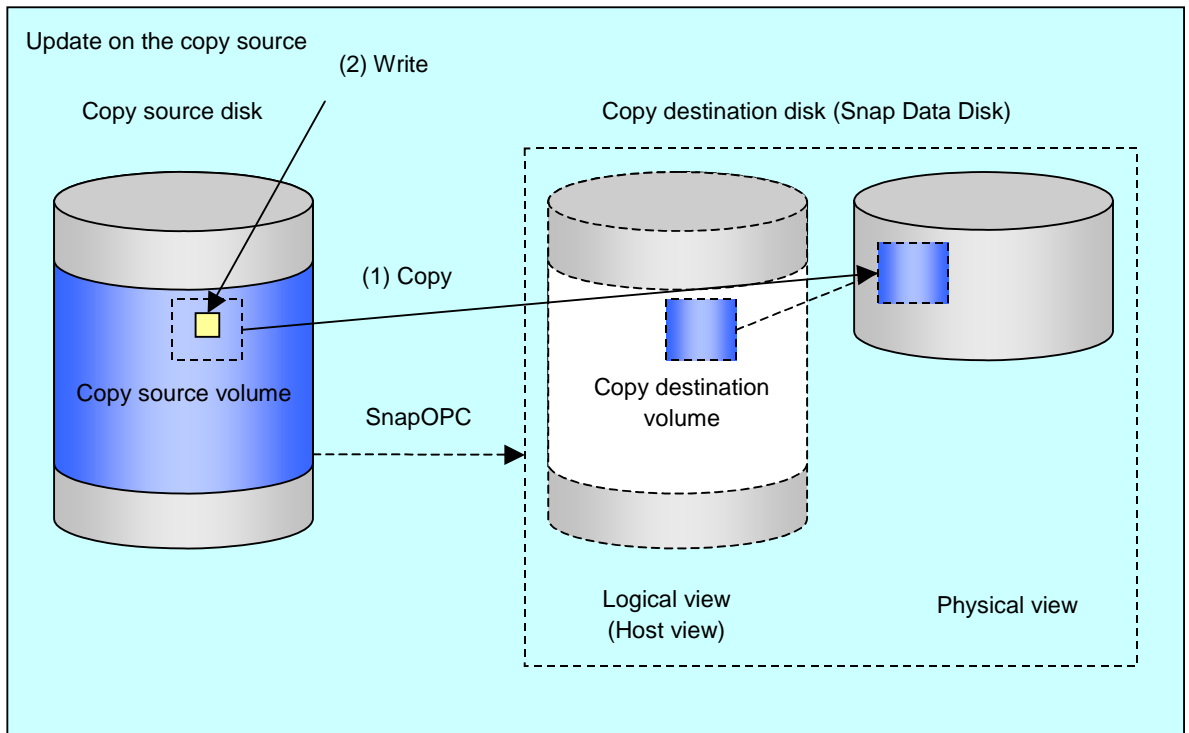
```
# /opt/FJSVswsrp/bin/swsrpmake -C /dev/dsk/c1t1d1 /dev/dsk/c1t1d2
FROM=/dev/dsk/c1t1d1@SV1,TO=/dev/dsk/c1t1d2@SV1 swsrpmake completed
#
```



swsrpmake -C



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 the SnapOPC started. This kind of copy processing is referred to as "Copy-on-Write".



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 the operation status display command (swsrpstat).

The following example shows the execution of the 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/dsk/ct1d1
Server Original-Volume      Replica-Volume      Direction Status      Execute Trk  Update
Rcv  Split Xfer Snap-Gen
SV1   /dev/dsk/ct1d1@SV1 /dev/dsk/ct1d2@SV1 regular   copy-on-write ----   off  0%
-----
#
```

If the replication creation command is executed again during SnapOPC processing, the SnapOPC session that has already been set up will be canceled, and a new session will be set up.



Note

When there is insufficient capacity in the Snap Data Volume or Snap Data Pool, the SnapOPC status changes to the error suspend status of 'failed' and replication volumes cannot be accessed (used).

The SnapOPC execution status can be checked from the Status field of the results output by the operation status display command (swsrpstat).

[Example]

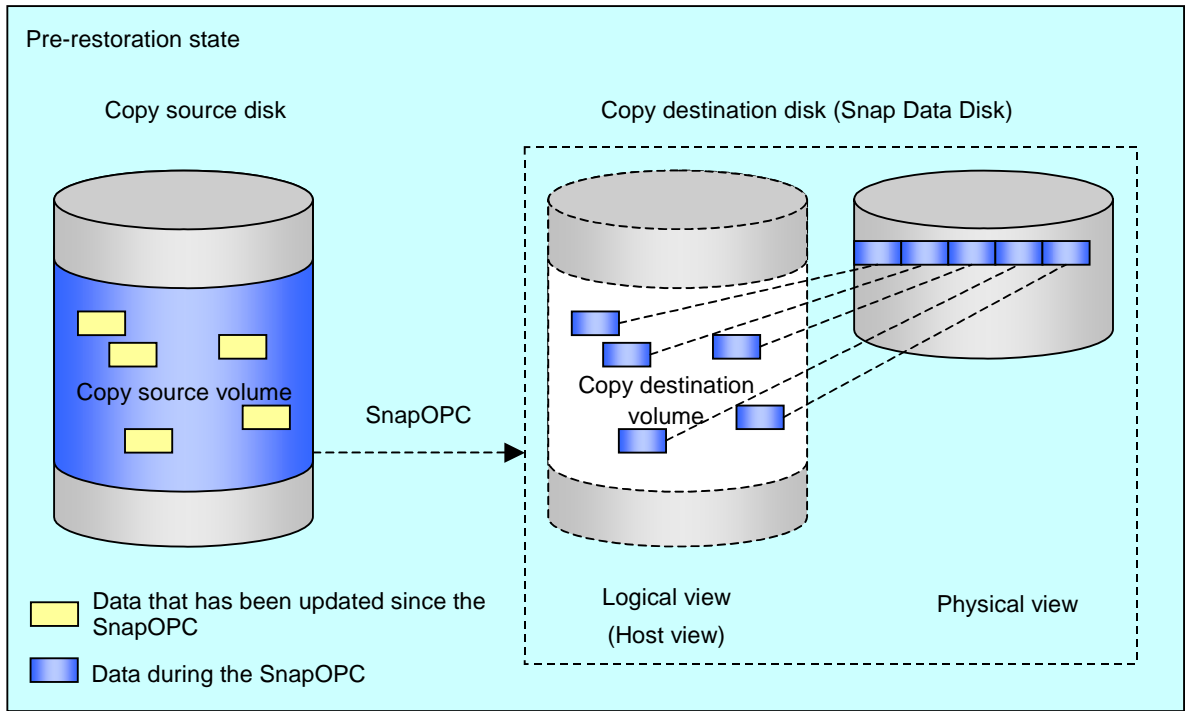
```
# /opt/FJSVswsrp/bin/swsrpstat -L /dev/dsk/ct1d1
Server Original-Volume      Replica-Volume      Direction Status Execute Trk  Update
Rcv  Split Xfer Snap-Gen
SV1   /dev/dsk/ct1d1@SV1 /dev/dsk/ct1d2@SV1 regular   failed ----   off  ----
-----
#
```

If SnapOPC is in the error suspend state of "failed", refer to "Troubleshooting when an insufficiency of free space has occurred in the Snap Data volume or Snap Data Pool" .

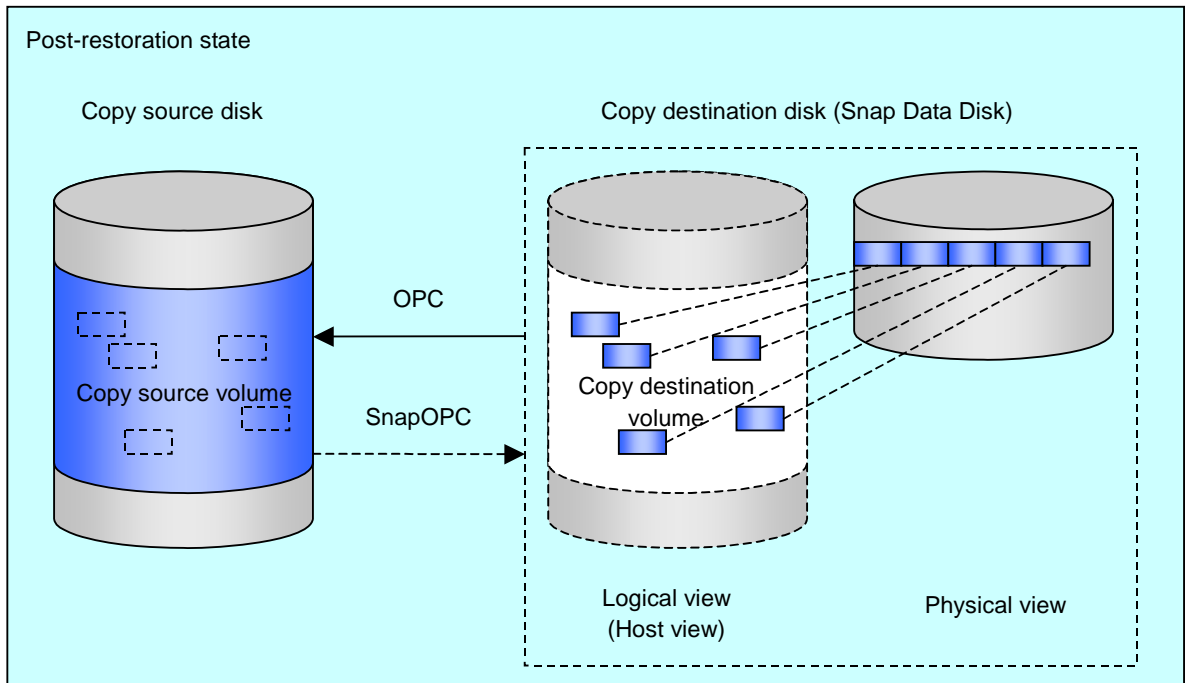
Perform restorations from Snap Data Volumes by running an OPC using the replication creation command (swsrpmake).

```
# /opt/FJSVswsrp/bin/swsrpmake /dev/dsk/ct1d2 /dev/dsk/ct1d1
FROM=/dev/dsk/ct1d2@SV1,TO=/dev/dsk/ct1d1@SV1 swsrpmake completed
#
```

When a restoration is executed, the SnapOPC session from the replication source volume to the replication destination volume is maintained, 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 replica was created is physically copied.



swsrpmake



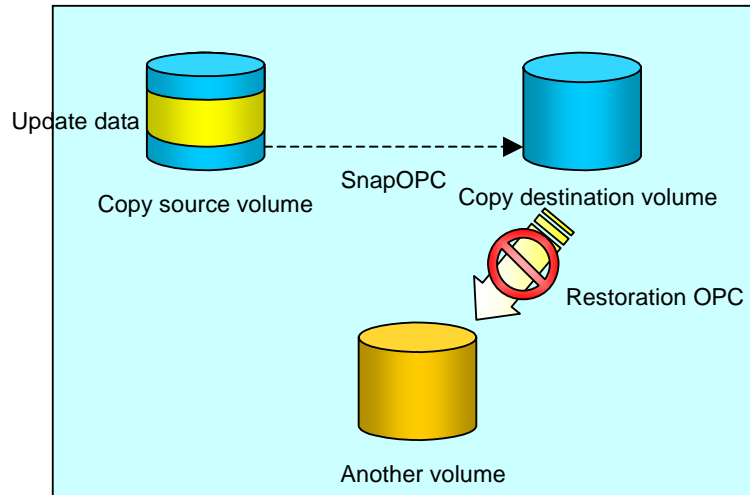
The execution status of restorations can be checked by specifying the `-E` option with the operation status display command (`swsrpstat`).

```
# /opt/FJVSwsrp/bin/swsrpstat -E /dev/dsk/clt1d1
Server Original-Volume  Replica-Volume  Direction Status Execute
SV1  /dev/dsk/clt1d1@SV1 /dev/dsk/clt1d2@SV1 reverse  snap  80%
#
```

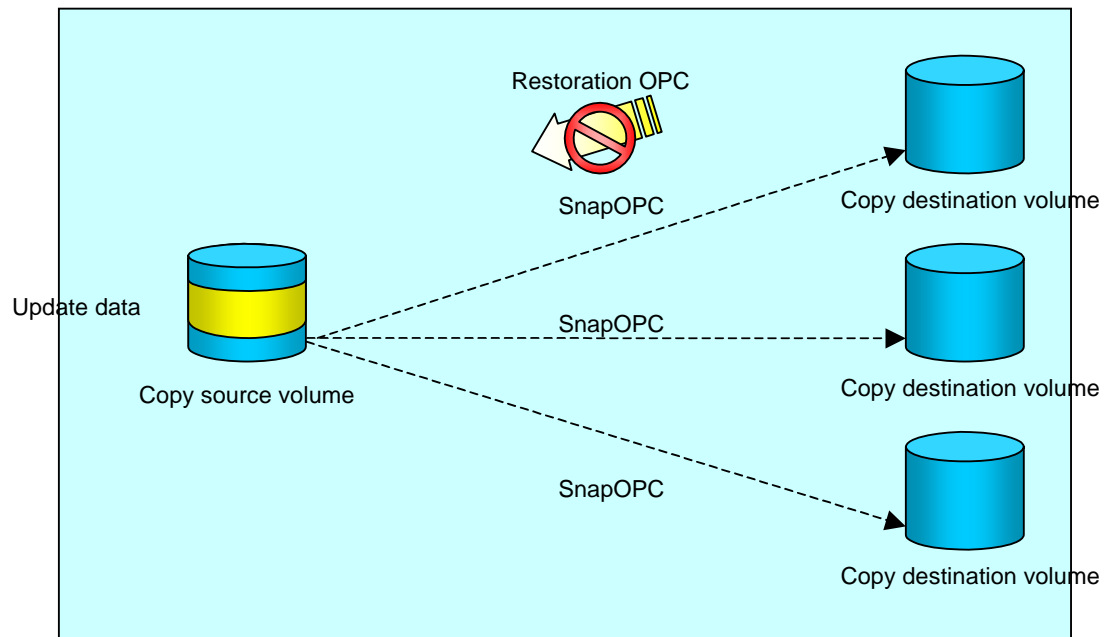


Note

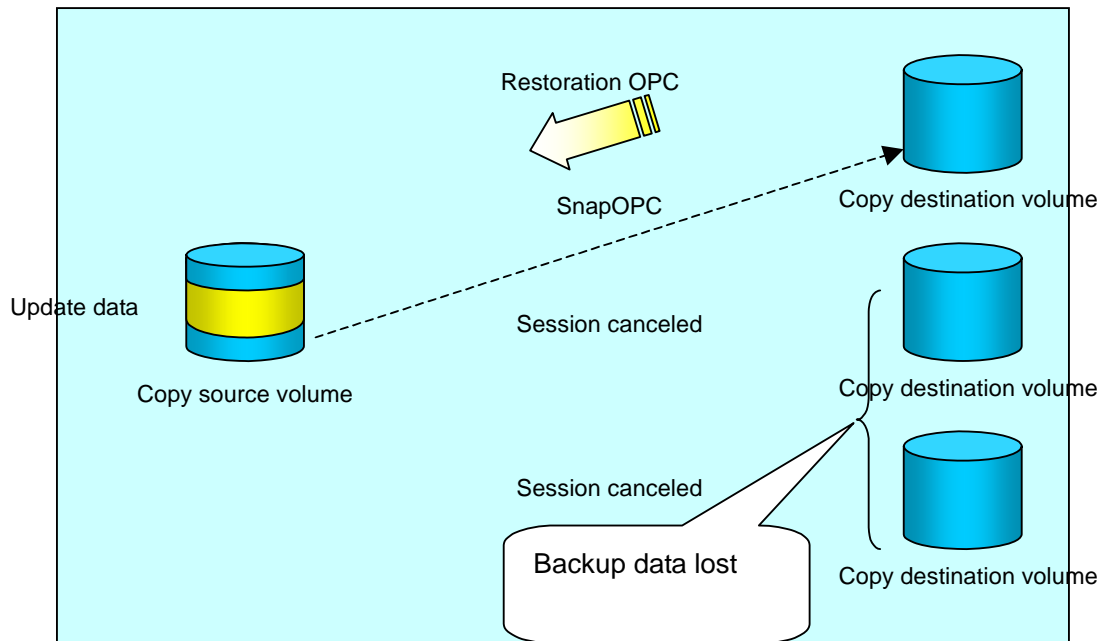
If a SnapOPC is being performed between the source volume and the destination volume, restorations to volumes other than the source volume cannot be executed. To restore to a volume other than the source volume, operating system copy functions (such as the cp command or the copy command) must be used.



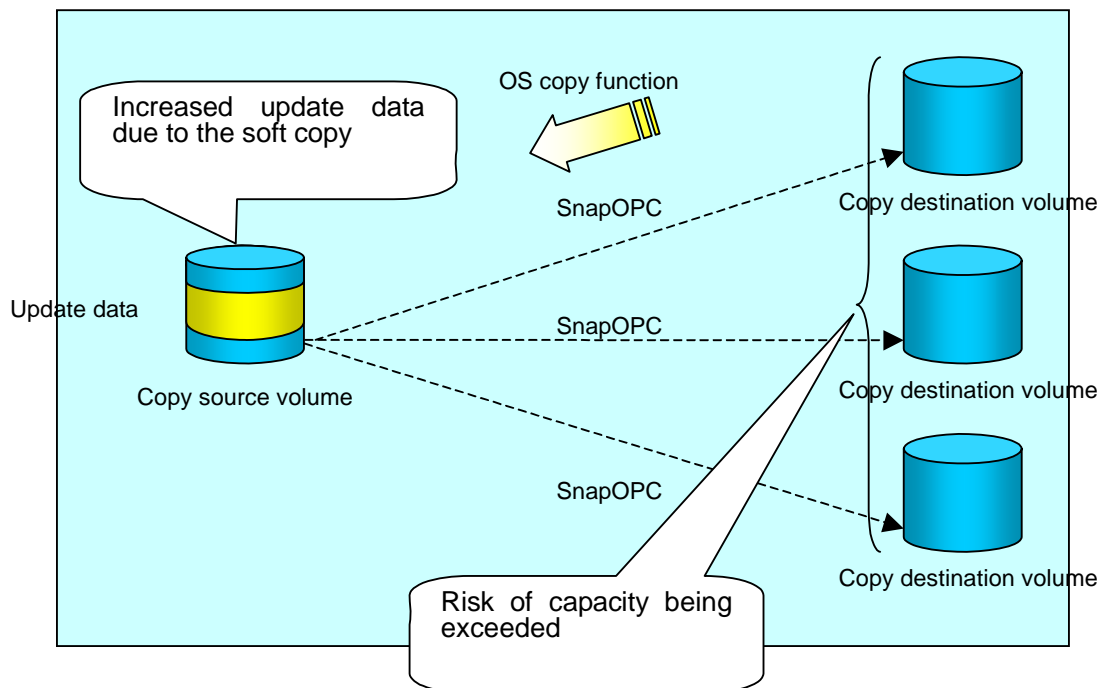
Additionally, if SnapOPCs are being performed to multiple copy destination volumes, restoration cannot be performed.



In this case, restoration using an OPC can be performed by canceling the other SnapOPCs. However, the backup data on the destination volumes whose SnapOPC sessions were canceled will be lost.



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 of a capacity shortage for the SnapOPC volume .



SnapOPC+ replication

Execute the replication creation command (swsrpmake) with the -P option specified to perform SnapOPC+ replication. This sets a SnapOPC+ session between the copy source volume and the copy destination volume. After the SnapOPC+ session is set, copy-on-write is performed between the copy source volume and the copy destination volume.

An example of executing the replication creation command with the -P option specified is shown below.

[Execution example]


```
# /opt/FJSVswsrp/bin/swsrpmake -P /dev/dsk/c1t2d0 /dev/dsk/c1t2d0@TARG-SV
FROM=/dev/dsk/c1t2d0@SRC-SV, TO=/dev/dsk/c1t2d0@TARG-SV swsrpmake completed
#
```

At this time, the logically copied copy destination volume is saved as a snap generation number 1.

The next time the replication creation command is executed with a different copy destination volume specified 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 the replication creation command with the -P option specified for the newly specified copy destination volume is shown below.

[Execution example]

```
# /opt/FJSVswsrp/bin/swsrpmake -P /dev/dsk/c1t2d0 /dev/dsk/c1t2d0@TARG-SV
FROM=/dev/dsk/c1t2d0@SRC-SV, TO=/dev/dsk/c1t2d0@TARG-SV swsrpmake completed
#
```

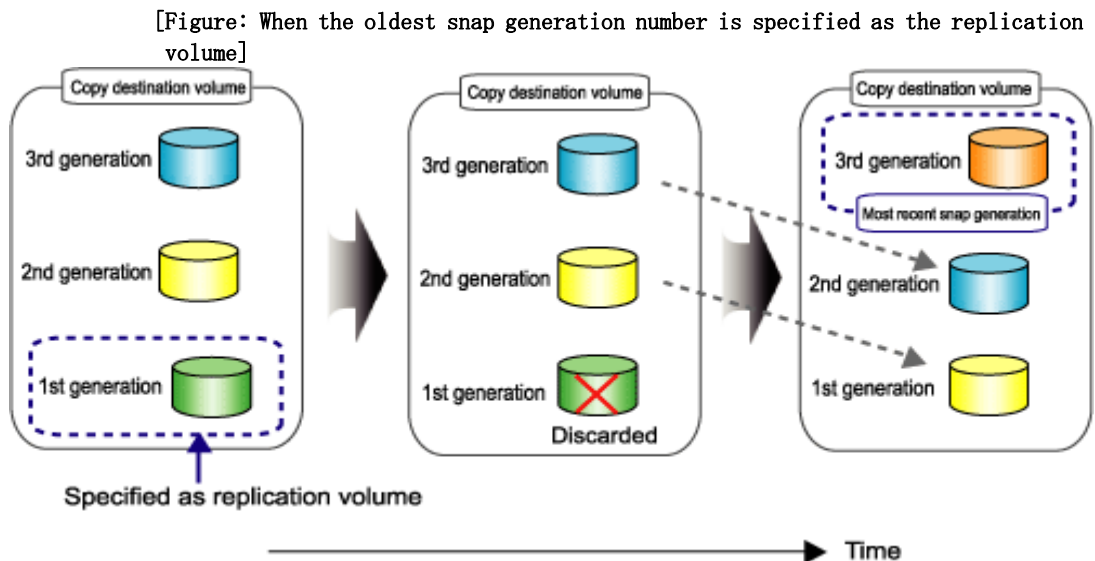
This time, the logically copied copy destination volume is saved as snap generation number 2.

Similarly, each time there is a new copy destination volume, a snap generation number is assigned, up to a maximum of 8 generations.



Note

If an earlier snap generation (other than the oldest snap generation) is specified as the copy destination volume when the replication creation command is executed, the replication creation 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).



The status of SnapOPC+ replication can be checked using the operating status display command (swsrpstat) with the -L option specified.

For the most recent snap generation, "copy-on-write (active)" is displayed in the Status column. For past snap generations, "copy-on-write (inactive)" is displayed. In the Update column, the amount of data that has finished being updated after replication creation is displayed as a percentage. In the Snap-Gen column, the snap generation number is displayed.

[Execution example]

```
# /opt/FJSVswsrp/bin/swsrpstat - L
Server Original-Volume      Replica-Volume      Direction Status
Execute Trk Update Rev Split Xfer Snap-Gen
SRC-SV /dev/dsk/c1t1d0@SRC-SV /dev/dsk/c1t1d1@TARG-SV regular
copy-on-write(inactive) ---- off 0% ---- ---- ---- 1
SRC-SV /dev/dsk/c1t1d0@SRC-SV /dev/dsk/c1t1d2@TARG-SV regular
copy-on-write(active) ---- off 5% ---- ---- ---- 2
#
```



Note

When there is insufficient capacity in the Snap Data Volume or the Snap Data Pool, the SpanOPC+ status changes to the error suspend status of 'failed', and this applies to SnapOPC+ operations carried out immediately beforehand, their status will change to 'failed' as well. Replication volumes with this status cannot be accessed (used). The status of SnapOPC+ replication can be checked in the Status column of the results output by the operating status display command (swsrpstat).

[Execution example]

```
# /opt/FJSVswsrp/bin/swsrpstat - L
Server Original-Volume      Replica-Volume      Direction Status Execute Trk
Update Rev Split Xfer Snap-Gen
SRC-SV /dev/dsk/c1t1d0@SRC-SV /dev/dsk/c1t1d1@TARG-SV regular failed ---- off
---- ---- ---- ----
SRC-SV /dev/dsk/c1t1d0@SRC-SV /dev/dsk/c1t1d2@TARG-SV regular failed ---- off
---- ---- ---- ----
#
```

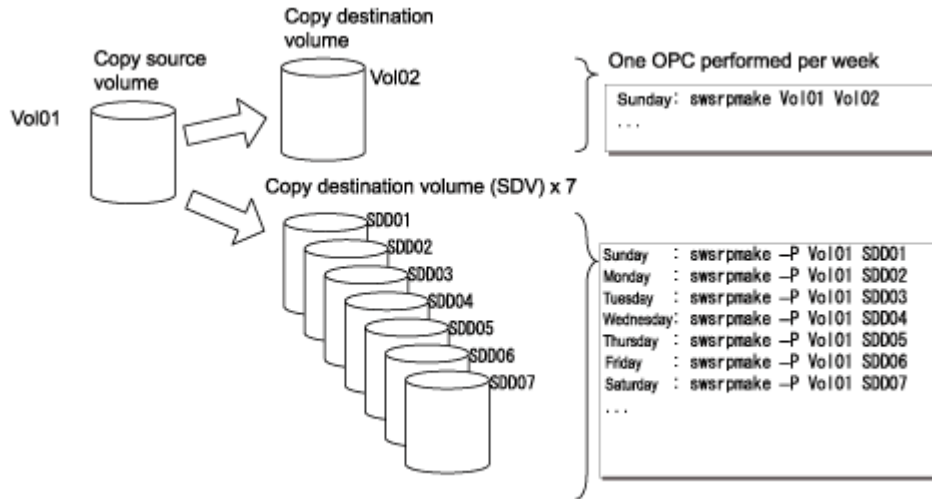
If the snapOPC+ execution status becomes error-suspended ("failed"), refer to "Action required if a Snap Data Volume or Snap Data Pool capacity shortage occurs".



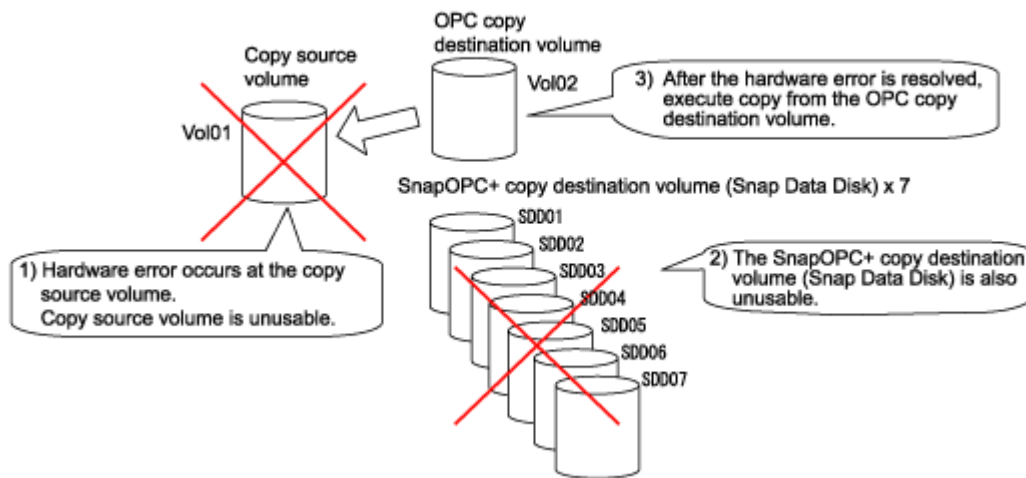
Point

Fujitsu recommends operating SnapOPC+ in conjunction with use of OPC/QuickOPC/EC(REC) to make a complete copy in case a hardware error occurs. Even if snapOPC+ and OPC/QuickOPC/EC(REC) are used in conjunction, the maximum number of sessions that can be set for one copy source volume is 8. The next example shows the operation of one OPC per week in conjunction with one SnapOPC+ per day.

[Figure: Operation example when SnapOPC+ and OPC are used in conjunction]



If a hardware error occurs at the copy source volume, perform restoration from the OPC copy destination volume.



7.5.2 Executing synchronized replication

Perform synchronized replication by using the following procedure:

1. Start synchronized processing using the synchronous processing start command (swsrpstartsync). Use the replication cancellation command (swsrpcancel) to cancel any synchronized processing that has already started.
2. After making sure that equivalency maintenance state has been established with the operation status display command (swsrpstat), temporarily stop synchronized processing with the replication creation command (swsrpmake) to create a replica of the source volume.
3. To copy the updated (that is, incremental) data, restart synchronized processing with the synchronous processing start command (swsrpstartsync).

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 the disk array.

The EC function operates in such a way that a copy is made to a destination volume in synchronization through a write to a source volume (that is, in synchronous write mode). In addition, EC functions enable the copy direction to be reversed from a suspend status.

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 the disk array.

The REC function provides three copy operation modes that can be selected for operation:

- Transfer mode
- Recovery mode
- Split mode

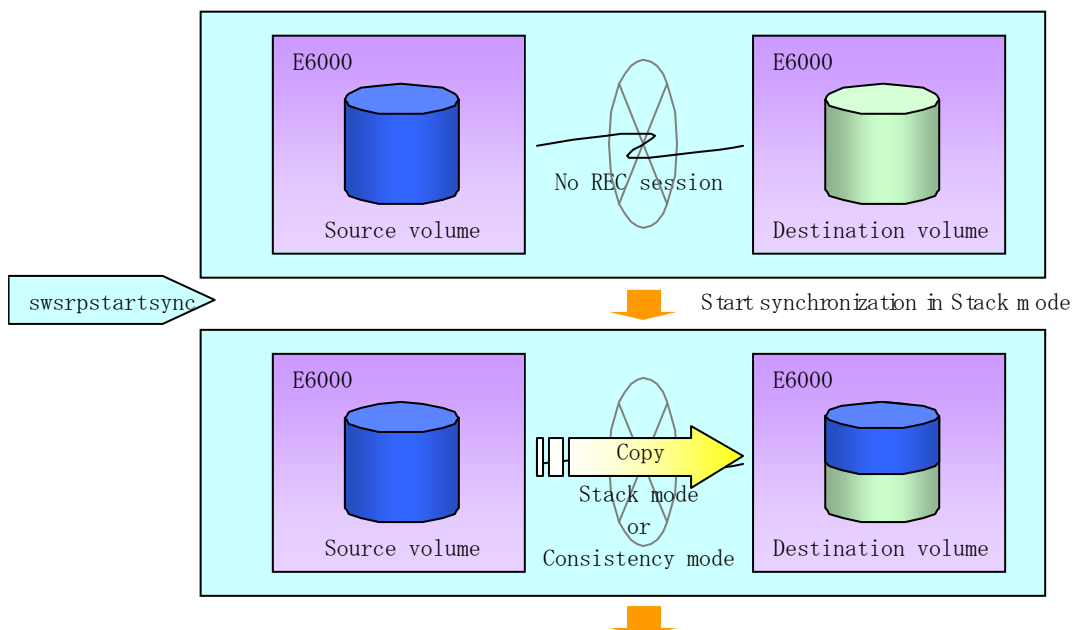
In addition, the REC function enables the copy direction to be reversed in suspended status.

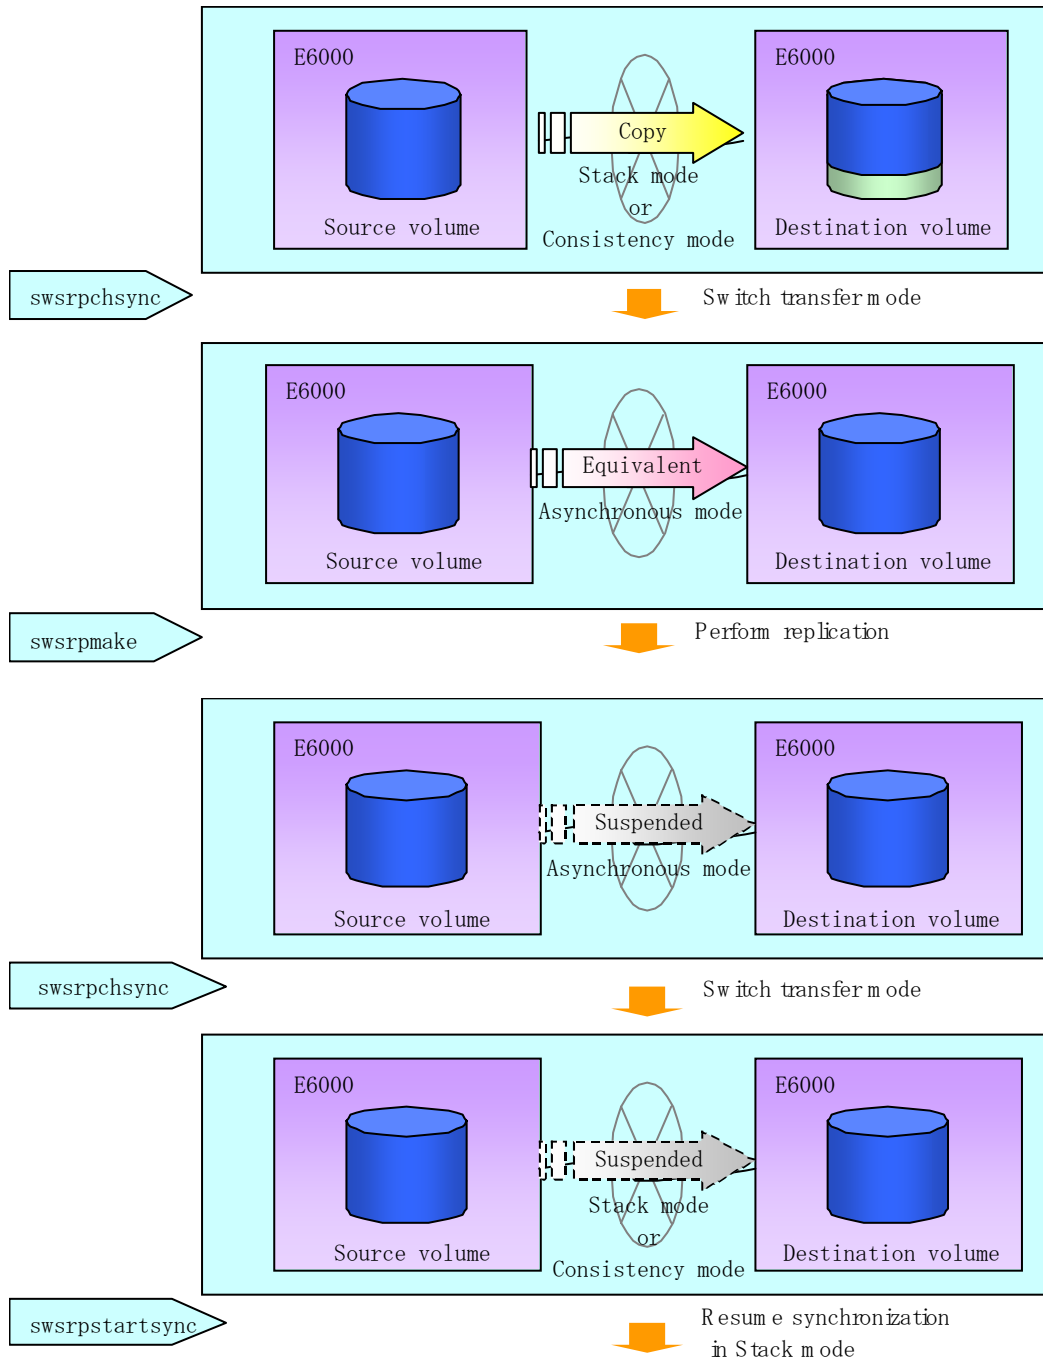
Transfer mode

Transfer mode provides the REC data transmission modes described below.

Mode	Description
Synchronous	When a write operation to a source volume occurs, this transfer mode returns the completion of the write operation to the host after copying is completed. In synchronous transfer mode, the performance of a write response depends on the performance of the network between the systems. Thus, any deterioration in network performance adversely affects the performance of a write response.
Asynchronous (sequential transfer)	This transfer mode starts sending data to a destination volume immediately after a response is made to a write operation on the source volume. The order of write operations is thus secured. If the volume of updates made to the source volume is excessive compared to transmission performance between the boxes, data to be copied remains stored, and write operations to the host are queued until the data is partly copied. To use asynchronous mode, the network must have at least the same performance as the update speed to the source volumes.
Stack	This mode stores (or stacks) data in the source system to be transferred and copies the data at irregular intervals to lower the speed of data transferred to the destination system. Updated data on the source volume is transferred to the destination volume at irregular intervals, and so the order of write operations is not guaranteed.
Consistency	This transfer mode guarantees the order in which multiple synchronous processes reflect data. Data updates in multiple synchronous processes are copied periodically and collectively, thus the order of write operations can be secured among multiple synchronous processes.

To perform a synchronous replication in Stack mode or Consistency mode, use the replication start command (`swsrpstartsync`), the replication creation command (`swsrpmake`), and the synchronization mode change command (`swsrpchsnc`). In addition, use the operation status display command (`swsrpstat`) to check the execution status. The figures below show how synchronous replication is performed in Stack mode or Consistency mode.





Note

If the REC operating mode is changed in equivalency maintenance status, after the operating mode is changed, the user must wait until the equivalency maintenance 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 (for example, a halt status).

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.

Mode	Description
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. You should 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).

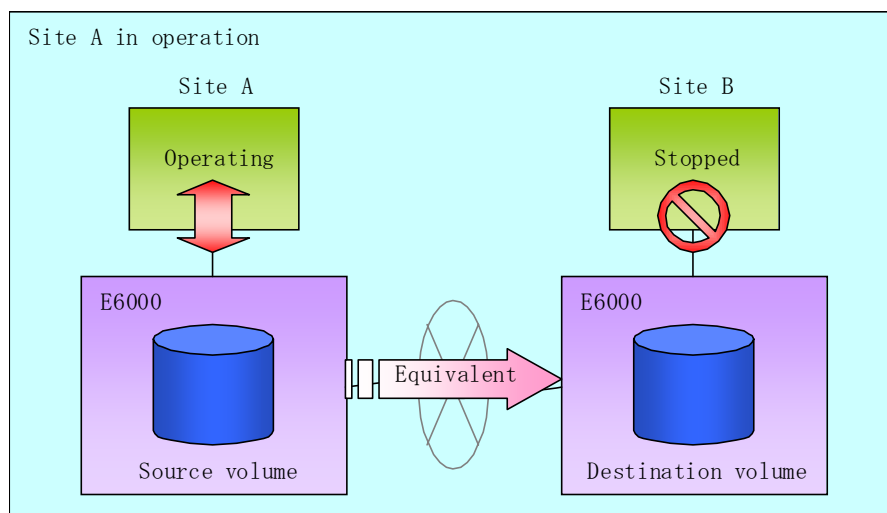
Mode	Description
Automatic Split	This split mode forcibly executes successful write operations to source volumes even if the inter-box FCRA path is fully blocked and halt status occurs. This mode enables write operations to source volumes even if the inter-box FCRA path is fully blocked, thus this mode does not affect transactions. When the inter-box FCRA path is recovered, copy processing resumes according to the recovery mode settings.
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 completely blocked. When the inter-box FCRA path is recovered, copy processing resumes according to the recovery mode settings.

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:

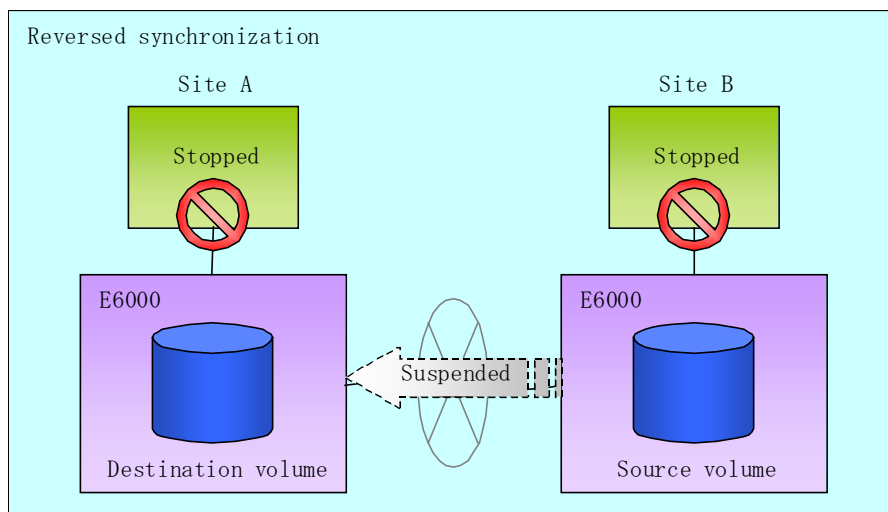
1. Assume that Site A is operating and REC is operating from Site A to Site B.



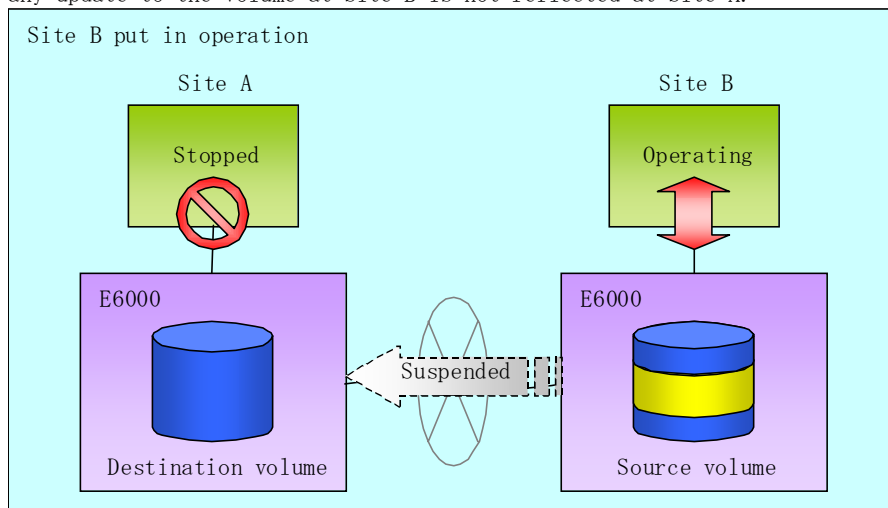
2. To switch the sites, execute the replication execution command to make a replication to Site B. Then, stop operating Site A.



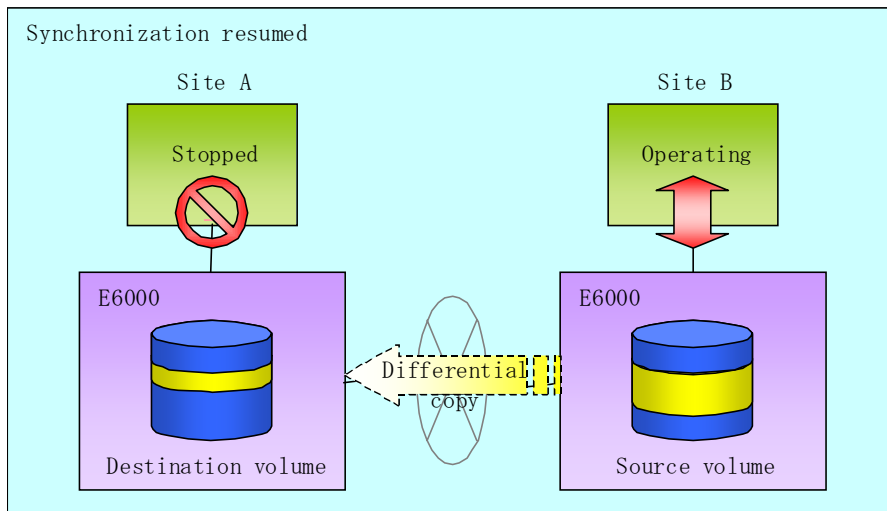
3. Execute the synchronization reverse command to reverse the copy direction.



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.



5. Start (or resume) synchronization from Site B to Site A. Updates made to the volume in Site B while synchronization is suspended are reflected into Site A with differential copies.

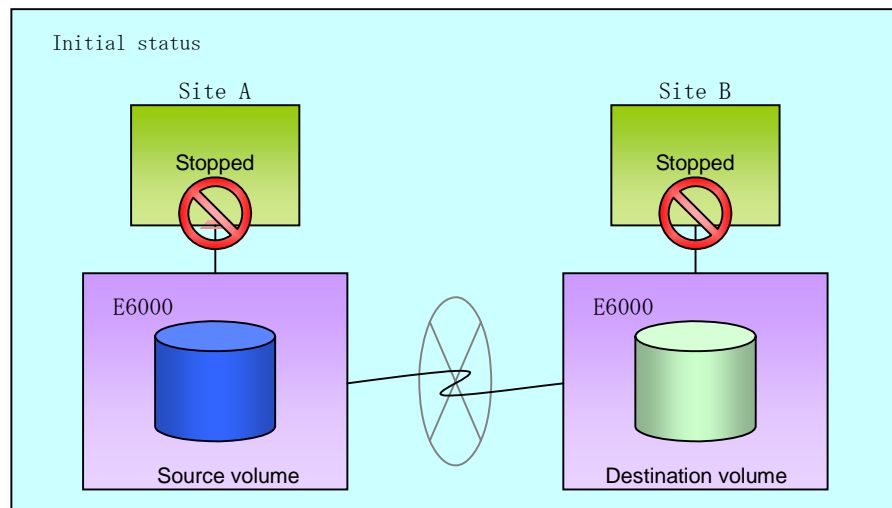


7.5.2.2.1 Initial copy skip function

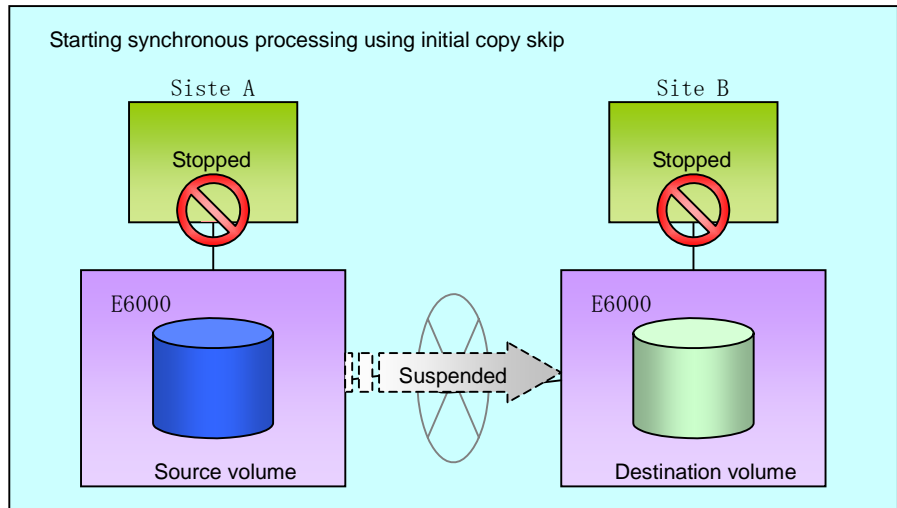
The initial copy skip function is used when the initial copy cannot be executed because of insufficient line capacity.

The following example shows the initial copy being skipped via tape transfer:

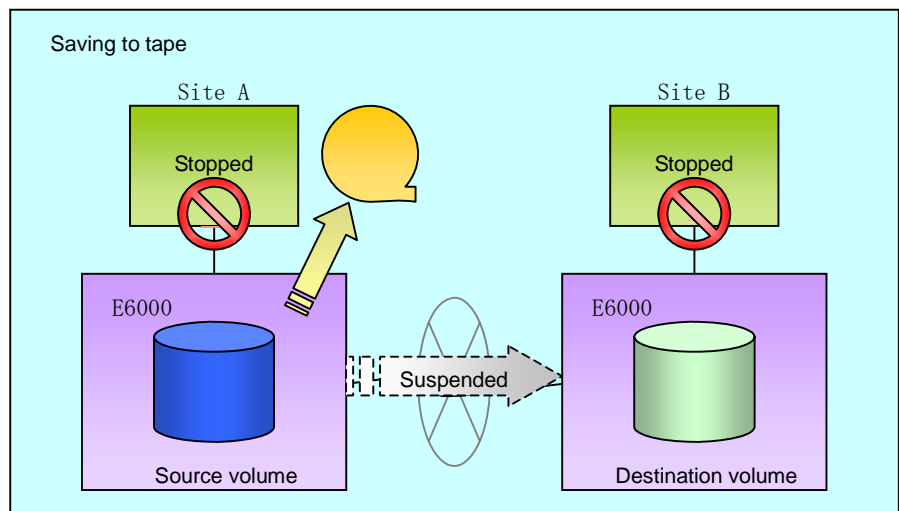
1. Suppose that operations at Site A have stopped.



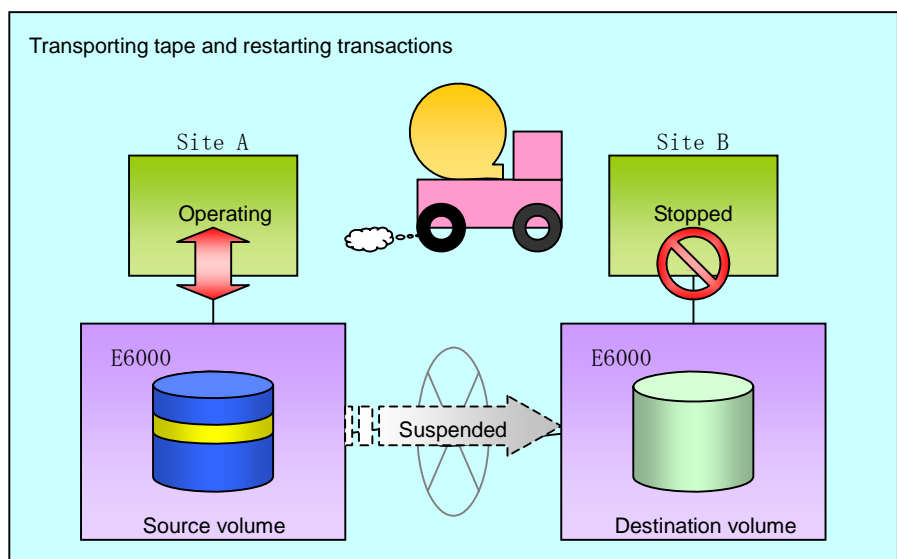
2. Next, synchronous processing begins, using the initial copy skip function. At this point, an REC session is set up, and the status is Replication Established. Data is not copied to the copy destination volume.



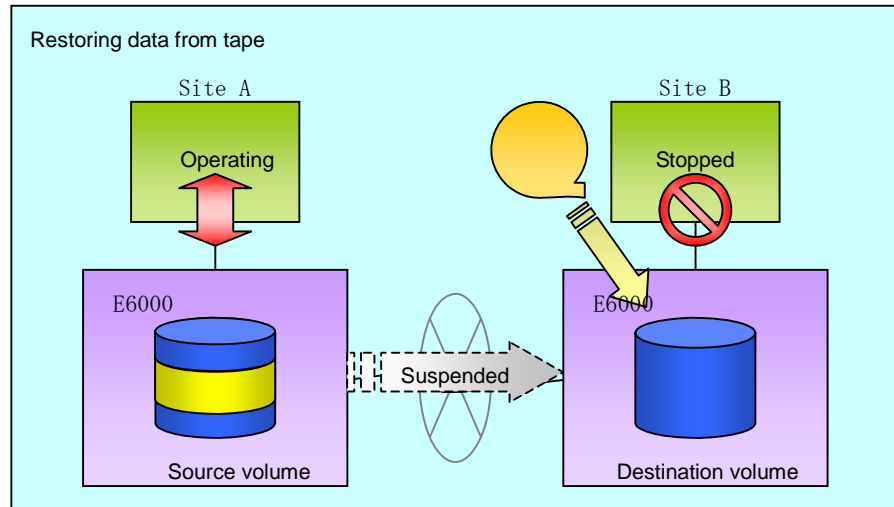
3. Next, the data on the copy source volume is backed up to tape.



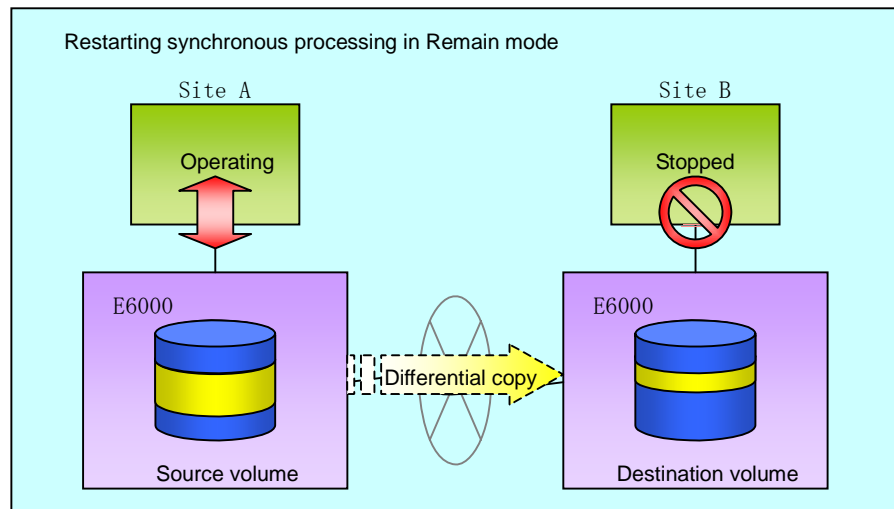
4. The tape medium is sent to Site B, and jobs at Site A restart.



- 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.



- 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.




7.5.2.3 Concurrent suspension function

The concurrent suspension function simultaneously suspends multiple EC/REC sessions for disk array systems.

By using this function, a copy with consistency can be easily taken.

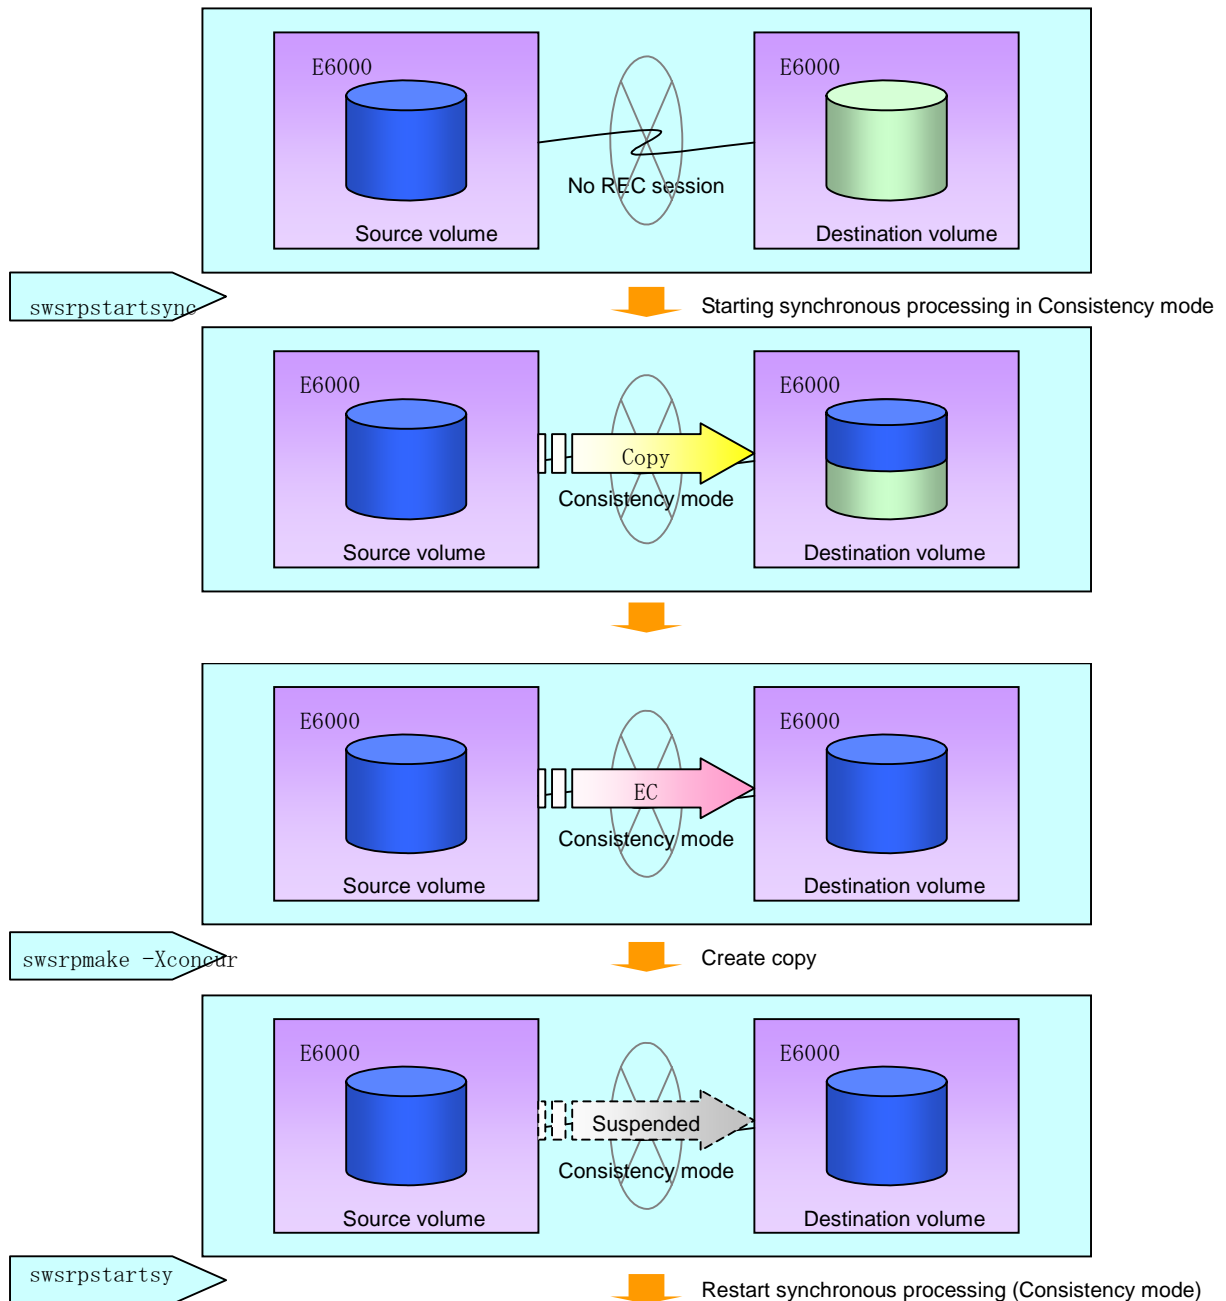
For example, it is useful for a database consisting of multiple volumes.

The behavior that takes place within a disk array system is shown below.



Replication using the concurrent suspension function is performed by specifying the `-Xconcur` option for the replication creation command (`swsrpmake`).

Additionally, if concurrent suspension 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 suspension is executed in Consistency mode is the same as the procedure for both asynchronous and synchronous modes.



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-direction copying by using the replication volume information configuration command (`swsrpsetvol`).

Restoration can be executed according to the following procedures:

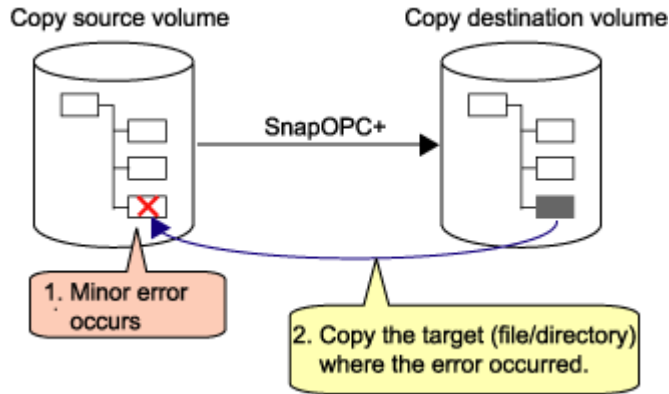
1. Execute the replication cancellation command (`swsrpcancel`) when the target volume for restoration has an EC session.
2. Execute the replication creation command (`swsrpmake`). The specification of a source volume and a replica volume is reversed when the replication is executed. The replication creation command is executed, specifying neither the `-T` option nor the `-C` option, in the backup operation that uses QuickOPC/SnapOPC.



Note

- With SnapOPC+, the replication function cannot be used to restore data from the copy destination volume to the copy source volume. If restoration is required, copy the data from the copy destination volume manually.

Figure: Restoration with SnapOPC+




7.5.4 Executing synchronous-type reversal

Use of the synchronous processing reverse command (swsrprevsync) greatly reduces the physical copy size during restoration if the differential data capacity is small, thereby reducing the time taken for AdvancedCopy.

In addition, for operations that combine multiple copy operations, restore can be performed without cancelling other copy operations. Therefore a complete copy need not be made after restoration, and the load on ETERNUS is reduced when operations are restarted.

Figure: Example of EC reverse operation



7.6 Changing the Configuration Information

When a Storage Server or device required in the replication has been changed, the information set in AdvancedCopy Manager must be changed. This section explains how to make these changes.

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 the replication volume information deletion command (swsrpdelvol) and then re-execute the replication volume information configuration command (swsrpsetvol).

7.6.2 Deleting a source volume or replica volume

To delete the source volume or the replica volume that has been set, use the replication volume information deletion command (swsrpdelvol) at a command prompt.

7.6.3 Changing the device information of a source or replica volume

When making device information and configuration change for a source or replica volume, the device configuration is changed after deletion processing of the source or replica volume, and the configuration of the volume is then performed.

Be sure to perform the following processing before changing the size of a source or replica volume or making a configuration change. If this is not done, the deletion of the source or replica volume after a device configuration change may not be possible.

1. Perform deletion processing on the original or replica volume to be changed. For details of this, refer to "Deleting a source volume or replica volume".
2. Change the device configuration.
3. Perform the additional processing of a device. For more information on this, refer to "Fetching device information from a Storage Server".
4. Perform the configuration of the new information in the source or replica volume.

7.7 Stopping the Operation

To stop replication processing that is in progress, or to change snapshot replication into synchronized replication, use the replication cancellation command (`swsrpcancel`) at a command prompt.



Note

To stop replication processing by SnapOPC+, the replication must be stopped starting with the oldest snap generation.

To forcibly stop the second or subsequent snap generations, specify the `-Xforce` option in the replication cancellation command (`swsrpcancel`). This will stop the specified snap generation as well as all earlier snap generations.

If the `-Xforce` option and the `-Xgroup` option are specified simultaneously, the SnapOPC+ sessions within the group are stopped, as well as all prior snap generations within the group.

When stopping the replication, stop the daemon on the Storage Server. Usually, it stops automatically at system shutdown.

It is also possible to stop an individual daemon if this is required for any reason. For details of this, refer to "Starting and Stopping Daemons".



Note

When the daemon stops, all functions of AdvancedCopy Manager running on the Storage Server are stopped.

Before stopping the Storage Management Server daemon, make sure that operations on all Storage Servers under management have been stopped.

7.8 LVM Volume Operation

The replication of LVM volumes can be classified into the following two modes, depending on the volume group configuration:

- Replication in units of volume groups
- Replication in units of physical disks (LU: Logical Unit)



Before starting this operation, you should understand the basic operation of ordinary volumes.

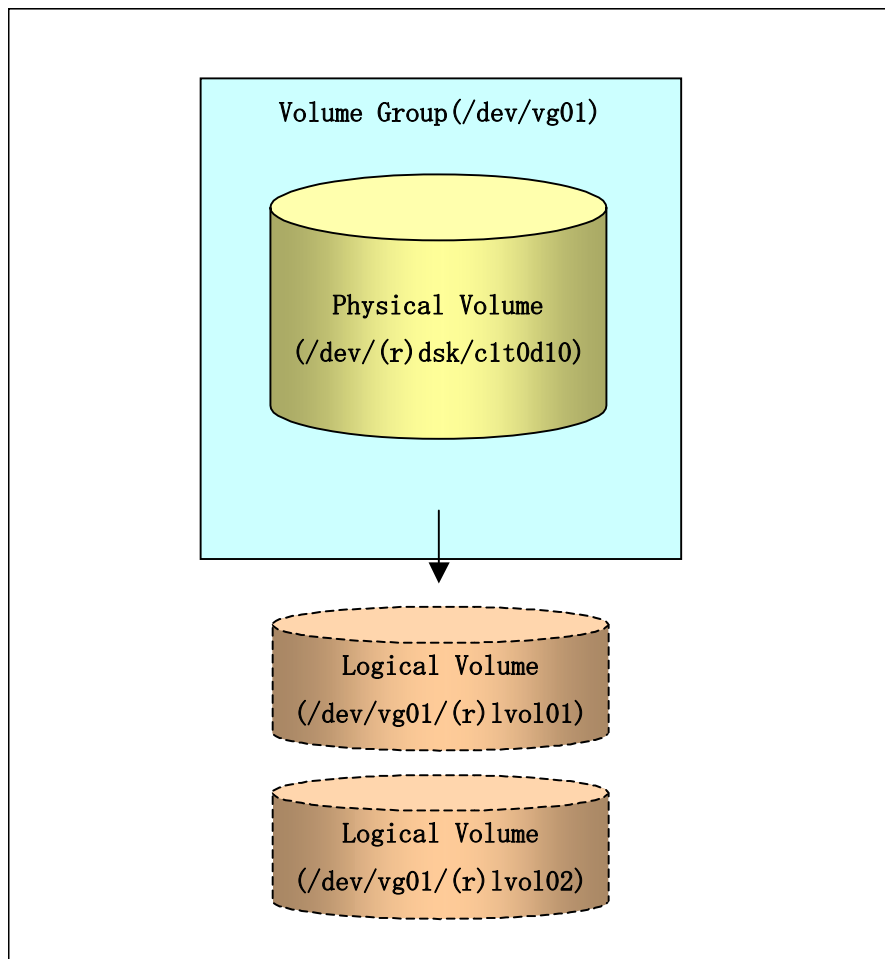
7.8.1 Replication in units of volume groups

Provided that all volume group configurations satisfy the following conditions, the replication can be performed in units of volume groups.

- One volume group has only one physical disk, and logical volumes are configured so that one physical disk includes n logical volumes.

If the above conditions are not satisfied, replication must be performed in units of physical disks.

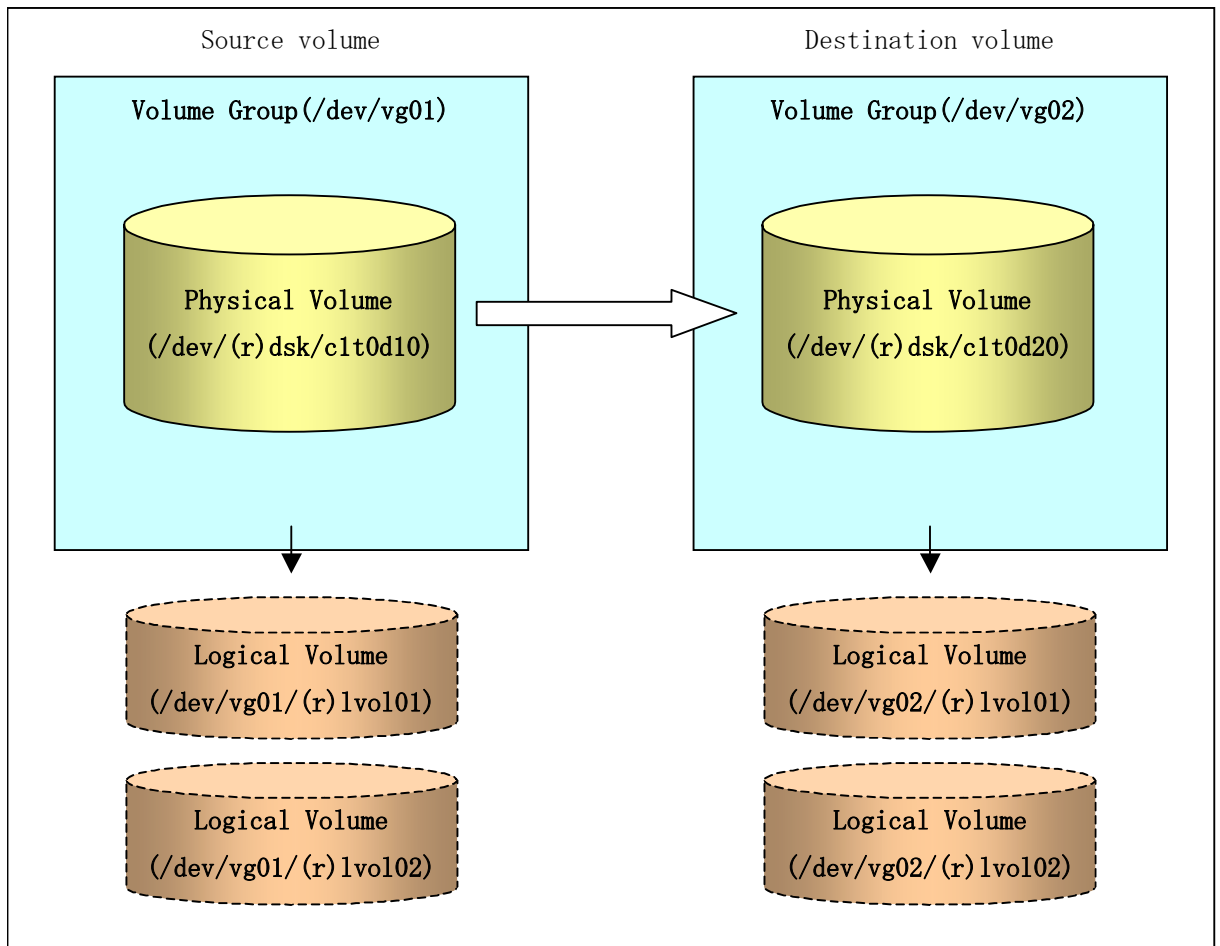
Figure Example of configuration in which replication can be performed in units of volume groups



7.8.1.1 Operation design

Ensure that the following conditions are met when designing volume groups for use as source or replica volume groups:

- All physical disks must have the same size.
- The configurations of all logical volumes must be the same.



7.8.1.2 Preparations

7.8.1.2.1 Saving the volume group configuration information file

The standard post-processing script restores a volume group configuration information file from the following file. Confirm that a volume group configuration information file has been saved.

```
/etc/lvmconf/<volume-group-name>.conf
```



Note

Replication cannot be performed if a volume group configuration information file has not been saved. In this case, target volume post-processing will result in an error during replication.

Customize the post-processing script for the target volume if volume group configuration information has been saved to another file.

7.8.1.2.2 Configuring the source volume and the replica volume

When configuring the source volume and the replica volume, specify their volume groups.

Example:

```
# /opt/FJSVswsrp/bin/swsrpsetvol /dev/vg01 /dev/vg02
swsrpsetvol completed
#
```

7.8.1.2.3 Customizing the pre-processing and post-processing scripts

If a volume group is to be replicated, the pre-processing and post-processing scripts must be customized.

Refer to "Pre-processing and Post-processing for Replication" for information on the customization procedure.



Note

If replication is attempted without customization of the scripts, pre-processing for the source volume causes an error, so that replication cannot be achieved.

7.8.1.3 Replication in units of volume groups

Execute replication with a volume group specified.

Example of snapshot replication

```
# /opt/FJSVswsrp/bin/swsrpmake /dev/vg01 /dev/vg02
FROM=/dev/vg01@SV1, TO=/dev/vg02@SV1 swsrpmake completed
#
```

Example of synchronous replication

```
# /opt/FJSVswsrp/bin/swsrpstartsync /dev/vg01 /dev/vg02
FROM=/dev/vg01@SV1, TO=/dev/vg02@SV1 swsrpstartsync completed
(After state of equivalency upkeep)
# /opt/FJSVswsrp/bin/swsrpmake /dev/vg01 /dev/vg02
FROM=/dev/vg01@SV1, TO=/dev/vg02@SV1 swsrpmake completed
#
```

7.8.2 Replication in units of physical disks

When the volume group configuration does not satisfy the conditions for operation in units of volume groups, replication can be performed by operation in units of physical disks.

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



Point

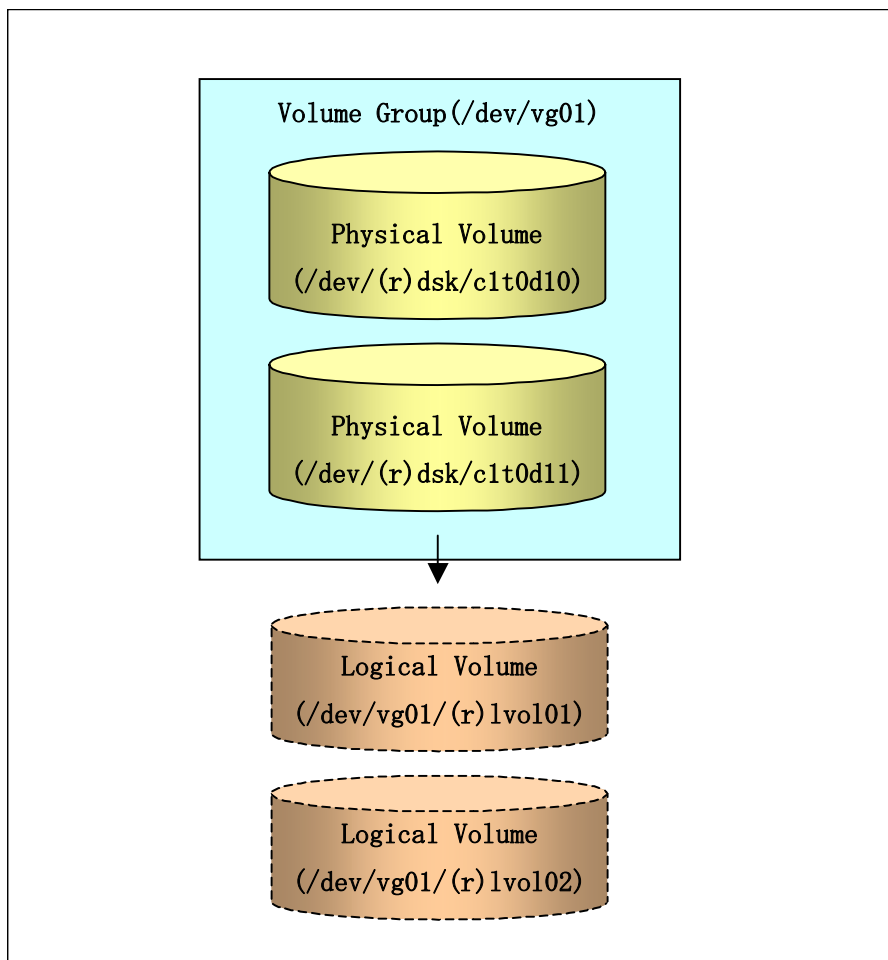
Refer to "Device information display command (stgxfwcmdispdev)", or the "Configure a Replication Group" of the ETERNUS SF AdvancedCopy Manager User's Guide for the method of checking if the physical volume is suitable for synchronization



Note

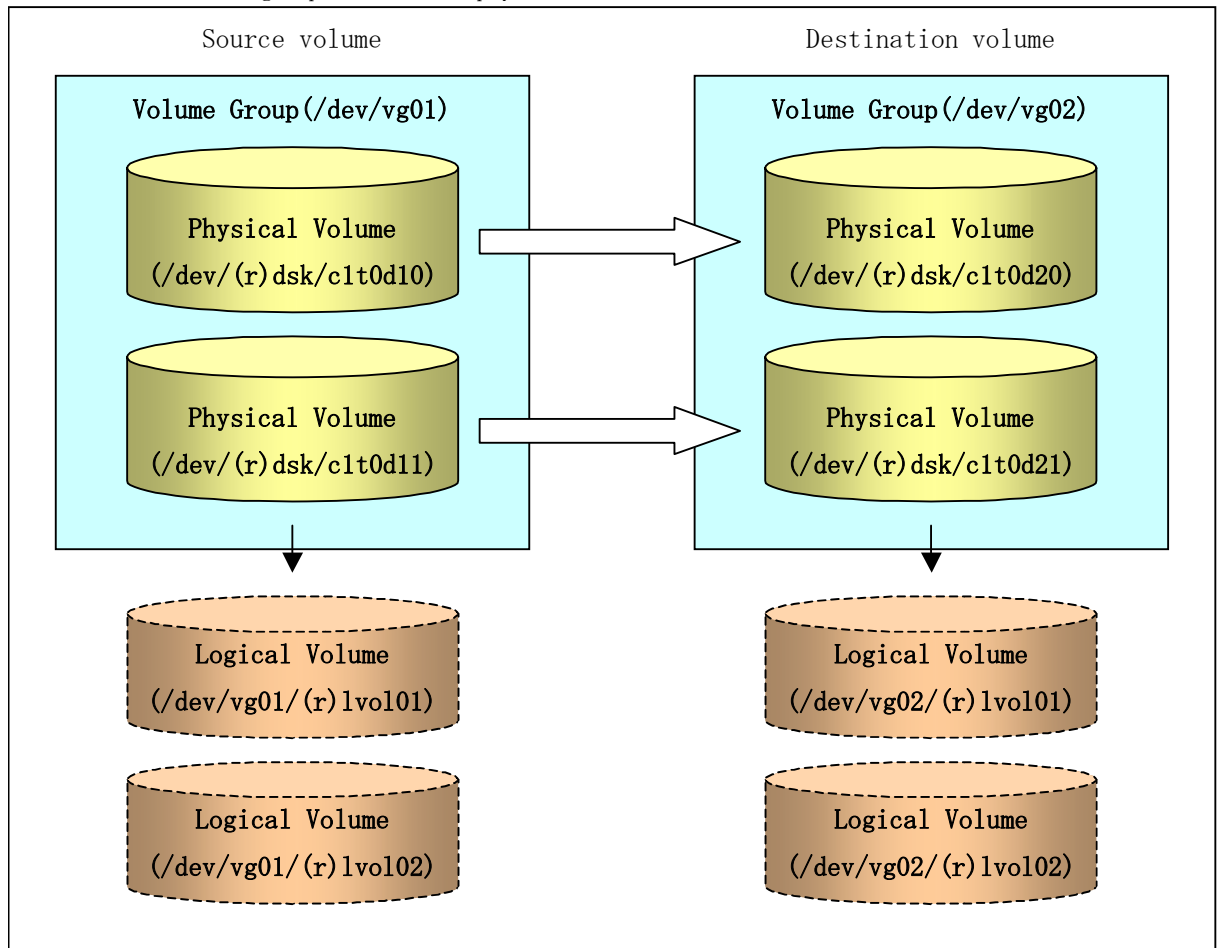
For operation in a cluster configuration, the device name (/dev/(r)dsk/c#t#d#) of the physical disks that comprise the volume group must be the same at all servers that comprise the cluster, and the ETERNUS disk indicated by the device name must be the same.

Figure Sample configuration for operation in units of physical disks



7.8.2.1 Operation design

To use the replica volume after replication in the situation when a source volume is a volume group, the replica volume must have the same logical volume configuration as the source volume and must be a volume group of the same physical size.



7.8.2.2 Preparations

7.8.2.2.1 Saving the volume group configuration information file

The volume group configuration information file must be restored during replication post-processing. Confirm that the volume group configuration file has been saved. It is usually saved to the following location:

```
/etc/lvmconf/<volume-group-name>.conf
```

7.8.2.2 Configuring the source and replica volumes

When configuring the source and replica volumes, specify all of the physical disks in the volume group.

Example:

```
# /opt/FJSVswsrp/bin/swsrpsetvol /dev/dsk/c1t0d10 /dev/dsk/c1t0d20
swsrpsetvol completed
# /opt/FJSVswsrp/bin/swsrpsetvol /dev/dsk/c1t0d11 /dev/dsk/c1t0d21
swsrpsetvol completed
#
```

7.8.2.3 Replication in units of physical disks

Perform the operation by synchronizing all physical disks in the volume group.

Perform the required pre-processing and post-processing work for each volume group before and after (respectively) the replication. Disable pre-processing and post-processing when operating individual physical disks.

Example of snapshot replication

```
(Perform preprocessing for the source and target volumes.)
# /opt/FJSVswsrp/bin/swsrpmake -f -t /dev/dsk/c1t0d10 /dev/dsk/c1t0d20
FROM=/dev/dsk/c1t0d10@SV1, TO=/dev/dsk/c1t0d20@SV1 swsrpmake completed
# /opt/FJSVswsrp/bin/swsrpmake -f -t /dev/dsk/c1t0d11 /dev/dsk/c1t0d21
FROM=/dev/dsk/c1t0d11@SV1, TO=/dev/dsk/c1t0d21@SV1 swsrpmake completed
#
(Perform postprocessing for the source and target volumes.)
```

Example of synchronous replication

```
(Perform preprocessing for the target volume.)
# /opt/FJSVswsrp/bin/swsrpstartsync -t /dev/dsk/c1t0d10 /dev/dsk/c1t0d20
FROM=/dev/dsk/c1t0d10@SV1, TO=/dev/dsk/c1t0d20@SV1 swsrpstartsync completed
# /opt/FJSVswsrp/bin/swsrpstartsync -t /dev/dsk/c1t0d11 /dev/dsk/c1t0d21
FROM=/dev/dsk/c1t0d11@SV1, TO=/dev/dsk/c1t0d21@SV1 swsrpstartsync completed
(After state of equivalency upkeep)
(Perform preprocessing for the source volume.)
# /opt/FJSVswsrp/bin/swsrpmake -f -t /dev/dsk/c1t0d10 /dev/dsk/c1t0d20
FROM=/dev/dsk/c1t0d10@SV1, TO=/dev/dsk/c1t0d20@SV1 swsrpmake completed
# /opt/FJSVswsrp/bin/swsrpmake -f -t /dev/dsk/c1t0d11 /dev/dsk/c1t0d21
FROM=/dev/dsk/c1t0d11@SV1, TO=/dev/dsk/c1t0d21@SV1 swsrpmake completed
#
(Perform postprocessing for the source and target volumes.)
```

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

	Pre-processing	Post-processing
Source volume	1. Ensure data integrity by preventing access to all logical volumes in the volume group. 2. If file systems are included, unmount all file systems in the volume group.	1. If file systems are included, remount the volumes that were unmounted by pre-processing.
Target volume	1. Prevent access to all logical volumes in the volume group. 2. If file systems are included, unmount all file systems in the volume group.	1. Reconfigure the volume group 2. If file systems are included, remount the volumes that were unmounted by pre-processing.

Reconfiguring the volume group

Reconfigure the volume group as follows:

1. Deactivate the volume group.

```
# /usr/sbin/vgchange -a n /dev/vg02
#
```

2. Restore volume group configuration information.

```
# /usr/sbin/vgcfgrestore -n /dev/vg02 /dev/rdisk/c1t0d20
# /usr/sbin/vgcfgrestore -n /dev/vg02 /dev/rdisk/c1t0d21
#
```

3. Reactivate the volume group.

```
[When cluster operation is enabled]
# /usr/sbin/vgchange -c y /dev/vg02

# /usr/sbin/vgchange -a e /dev/vg02

[When cluster operation is disabled]
# /usr/sbin/vgchange -a y /dev/vg02
#
```

Reconfigure shared mode volume groups using the following procedure:

1. Stop the volume group (on all nodes involved in the transaction).

```
# /usr/sbin/vgchange -a n /dev/vg02
#
```

2. Restore volume group configuration information (on the node where the volume group was created).

```
# /usr/sbin/vgcfgrestore -n /dev/vg02 /dev/rdisk/clt0d20
# /usr/sbin/vgcfgrestore -n /dev/vg02 /dev/rdisk/clt0d21
#
```

3. Mark the volume group as shareable (on the node where the volume group was created).

```
# /usr/sbin/vgchange -S y -c y /dev/vg02
#
```

4. Start the volume group (on all nodes involved in the transaction).

```
# /usr/sbin/vgchange -a s /dev/vg02
#
```

7.9 VxVM Volume Operation

VxVM volumes are replicated in units of physical disks that constitute each VxVM volume.

When the replication 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 operated synchronously.



Refer to "Device information display command (stgxfwcmdispdev)", or the "Configure a Replication Group" section of the ETERNUS SF AdvancedCopy Manager User's Guide for the method of checking if the physical volume allows synchronization.



Before starting this operation, ensure that you understand the basic operation of ordinary volumes.



Snapshot backup is recommended for operation in units of physical disks. In the case of synchronized backup, the command (VxVM command etc.), which disk access produces to a copy destination disk during a copy, cannot be executed.




For operation in a cluster configuration, the device name (/dev/(r)dsk/c#t#d#) of the physical disks that comprise the volume group must be the same at all servers that comprise the cluster, and the ETERNUS disk indicated by the device name must be the same.

7.9.1 Operational configuration

Design the disk groups of the source and replica volumes. Observe the following conditions when designing the disk groups:

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



7.9.2 Preparations

7.9.2.1 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.

```
# /usr/sbin/vxprint -hvmps -g srcdg > /tmp/aaa
#
```

7.9.2.2 Configuring the source and replica volumes

When configuring the source and replica volumes, specify all of the disks in the volume group.

Example:

```
# /opt/FJSVswsrp/bin/swsrpsetvol /dev/vx/dmp/clt0d10@SV1 /dev/vx/dmp/clt0d20@SV2
swsrpsetvol completed
# /opt/FJSVswsrp/bin/swsrpsetvol /dev/vx/dmp/clt0d11@SV1 /dev/vx/dmp/clt0d21@SV2
swsrpsetvol completed
#
```

7.9.3 Replication

Perform the operation of synchronizing all the physical disks in the disk group.

Perform the required pre-processing and post-processing work for each volume group before and after (respectively) the replication. Disable pre-processing and post-processing when operating individual physical disks.

Example of snapshot backup

```
(Perform preprocessing for the source and target volumes.)
# /opt/FJSVswsrp/bin/swsrpmake -f -t /dev/vx/dmp/clt0d10@SV1 /dev/vx/dmp/clt0d20@SV2
FROM=/dev/vx/dmp/clt0d10@SV1, TO=/dev/vx/dmp/clt0d20@SV2 swsrpmake completed
# /opt/FJSVswsrp/bin/swsrpmake -f -t /dev/vx/dmp/clt0d11@SV1 /dev/vx/dmp/clt0d21@SV2
FROM=/dev/vx/dmp/clt0d11@SV1, TO=/dev/vx/dmp/clt0d21@SV2 swsrpmake completed
#
(Perform postprocessing for the source and target volumes.)
```

Example of synchronous replication

```
(Perform preprocessing for the target volume.)
# /opt/FJSVswsrp/bin/swsrpstartsync -t /dev/vx/dmp/clt0d10@SV1 /dev/vx/dmp/clt0d20@SV2
FROM=/dev/vx/dmp/clt0d10@SV1, TO=/dev/vx/dmp/clt0d20@SV2 swsrpstartsync completed
# /opt/FJSVswsrp/bin/swsrpstartsync -t /dev/vx/dmp/clt0d11@SV1 /dev/vx/dmp/clt0d21@SV2
FROM=/dev/vx/dmp/clt0d11@SV1, TO=/dev/vx/dmp/clt0d21@SV2 swsrpstartsync completed
(After state of equivalency upkeep)
(Perform preprocessing for the source volume.)
# /opt/FJSVswsrp/bin/swsrpmake -f -t /dev/vx/dmp/clt0d10@SV1 /dev/vx/dmp/clt0d20@SV2
FROM=/dev/vx/dmp/clt0d10@SV1, TO=/dev/vx/dmp/clt0d20@SV2 swsrpmake completed
# /opt/FJSVswsrp/bin/swsrpmake -f -t /dev/vx/dmp/clt0d11@SV1 /dev/vx/dmp/clt0d21@SV2
FROM=/dev/vx/dmp/clt0d11@SV1, TO=/dev/vx/dmp/clt0d21@SV2 swsrpmake completed
#
(Perform postprocessing for the source and target volumes.)
```

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

	Pre-processing	Post-processing
Source volume	<ol style="list-style-type: none"> 1. Ensure data integrity by preventing access to all logical volumes in the volume group. 2. If file systems are included, unmount all file systems in the volume group. 	<ol style="list-style-type: none"> 1. If file systems are included, remount the volumes that were unmounted by pre-processing.
Target volume	<ol style="list-style-type: none"> 1. Prevent access to all logical volumes in the volume group. 2. If file systems are included, unmount all file systems in the volume group. 3. Export the disk group. 4. A disk group subordinate's physical disk is set to offline. 	<ol style="list-style-type: none"> 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/vxconfigstore -p dstdg
Diskgroup dstdg configuration restoration started .....

Installing volume manager disk header for clt0d20 ...
Installing volume manager disk header for clt0d21 ...
-
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:
  vxconfigstore -c dstdg ==> to commit the restoration.
  vxconfigstore -d dstdg ==> to abort the restoration.
#
```

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

```
# /etc/vx/bin/vxconfigstore -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 replication source volume as disk group dstdg on the replication destination server.

```
# /usr/sbin/vxdg -C -n dstdg import srcdg
#
```



Note

When the disk group name is the same in both the source server and the destination server, 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 since VxVM5.0, this is restored.

```
# vxdisk updateudid c1t0d20 c1t0d21
```



Note

In a cluster system, when a disk group or a mount resource has been defined, instead of using the import/export commands for the disk group use the online/offline processing, and instead of using the file system mount/unmount commands use the online/offline processing.



Note

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



Note

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

Chapter 8 Operational Maintenance

This chapter provides information on how to maintain a database, perform troubleshooting, and change the operation information.

8.1 Maintaining a Database

AdvancedCopy Manager holds various repositories and management information as information on the usage of its functions. They are named generically and are referred to as a database in this chapter.

The databases are important information and it is recommended that they be backed up periodically. Note that the databases referred to here differ from databases used for transactions, such as ORACLE databases.

The following table lists instances when backup of the database is needed.

Table: Databases on a Storage Server

Type	When a backup should be created	Backup method (references)	Size
Registry	A Storage Server has been added or Storage Management Server information (server name, IP address, or port number) has changed.	Backing up a registry or Making a package backup of databases	Refer to Backing up a registry.
Backup management list	A backup transaction has been completed.	Backing up a backup management list or Making a package backup of databases	Refer to Backing up a backup management list.
Volume configuration information (Backup management)	When the configuration of the volume group has changed	Preserving the volume configuration information	Refer to Preserving the volume configuration information
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 Making a package backup of databases	Check the size of the pre-processing and post-processing script files for backup management.
Pre-processing and post-processing scripts for tape-copy	A pre-processing or post-processing script has changed.	Saving a database by using the copy command or Making a package backup of databases	Check the size of the pre-processing and post-processing script files for tape-copy.
Replication management list	A replication transaction has been completed.	Backing up a replication management list or Making a package backup of databases	Refer to Backing up a replication management list.
Volume configuration information (Replication management)	When the configuration of the volume group is changed	Preserving the volume configuration information	Refer to Preserving the volume configuration information
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 Making a package backup of databases	Check the size of the pre-processing and post-processing script file for replication management.

Type	When a backup should be created	Backup method (references)	Size
Policy file for management information package backup	The policy file has changed.	Saving a database by using the copy command or Making a package backup of databases	Check the size of the policy file for management information package backup.



Point

You can make backups of databases either individually or as a whole.

If a Storage Server also operates as a Storage Management Server, the Storage Server contains a database for the Storage Management Server as well as for the Storage Server.

8.1.1 Making individual backups of databases

8.1.1.1 Maintaining a backup a 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

Back up the backup management.

For more information, refer to "Resource backup command (swstresback)."

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/<logic node name>/data
```



Point

When you reuse a directory that has previously been used to store the data, the size of this directory must be three times greater than the size indicated above.

8.1.1.1.2 Restoring a backup management list

Restore the backup management list.

For more information about this, refer to "Resource restore command (swstresrst)."

8.1.1.1.3 Preserving the volume configuration information

When the transaction volume or backup volume is a volume group, the volume configuration information file `"/etc/lvmconf/<volume_group_name>.conf"` is required to restore the volume configuration information in the processing after backup and restoration.

In order to prepare for a system crash, ensure that you maintain the above mentioned file on another storage medium at the time of creating the volume group.



Usually, a volume configuration information file `"/etc/lvmconf/<volume_group_name>.conf"` is automatically created at the time of `vgcreate (1M)` execution. If `vgcfgbackup (1M)` is used, a backup of the volume configuration information can be created by assigning arbitrary file names. For details of this, refer to the online manual.

8.1.1.2 Maintaining a replication management list

This section describes 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 Storage Management Server. For information about backing up a repository, refer to the version of the "Saving a database" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide" relevant to the OS of the Storage Management Server.



— Operation information for replication management is saved in a repository. It is recommended that a replication management list be backed up when operation information for replication management is changed with the Replication volume information setting command (swsrpsetvol) or with the Replication volume information deletion command (swsrpdelv)

8.1.1.2.2 Restoring a replication management list

Perform the following steps to restore a replication management list.

Restore the repository to the Storage Management Server. For information about restoring a repository, refer to "Recovering a database" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide" relevant to the OS of the Storage Management Server. Specify 0 for the recovery mode.

On the Storage Server, execute the Resource adjustment command (swsrprecoverres) with the `"-r"` option specified. For details of the resource adjustment command, refer to "Resource adjustment command (swsrprecoverres)".



— When the procedure above is performed, operation information for replication management of a repository is copied to a Storage Server, where it is available for use.

- When the operation information for replication management in the repository does not need to be changed and only the operation information for replication management on the Storage Server is incorrect for some reason, perform the procedure above.

8.1.1.2.3 Preserving the volume configuration information

When the source volume or replica volume is a volume group, the volume configuration information file `"/etc/lvmconf/<volume_group_name>.conf"` is needed to restore the volume configuration information in the post-processing script.

In order to prepare for a system crash, ensure that you maintain the abovementioned file on another storage medium at the time of creating the volume group.



Usually, a volume configuration information file `"/etc/lvmconf/<volume_group_name>.conf"` is automatically created at the time of `vgcreate (1M)` execution. If `vgcfgbackup (1M)` is used, a backup of volume configuration information can be created with arbitrary file names. For details of this, refer to the online manual.

8.1.1.3 Maintaining a database

For information about maintaining a database used by a Storage Management Server, refer to the version of the "Maintaining a database" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide" relevant to the OS on the Storage Management Server:

8.1.1.4 Maintaining an authentication management list

Refer to the edition of the following manual that is applicable to the Storage Management Server OS for details of the authentication facility management method.

"Maintaining an authentication management list" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide".

8.1.1.5 Maintaining a registry

This section describes how to maintain a registry. Perform this procedure on a Storage Management Server or a Storage Server.

8.1.1.5.1 Backing up a registry

Back up the following files by executing the `cp` command.

- a. File names

In the case of a non cluster

```
/etc/opt/swstorage/swnode.ini
```

In the case of a cluster

```
/etc/opt/swstorage/Logic node name/swnode.ini
```

- b. 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.

8.1.1.5.2 Restoring a registry

Restore the following file by executing the cp command:

- a. File names
In the case of a non cluster

```
/etc/opt/swstorage/swnode.ini
```

In the case of a cluster

```
/etc/opt/swstorage/Logic node name/swnode.ini
```

- b. 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 Storage Server, Storage Management Server, or tape server.

With a package backup of management information, you can make a backup of management information for all Storage Servers, Storage Management Servers, and tape servers from just one Storage Server, Storage Management Server, or tape server. Executing a package backup of management information runs backup commands (making backups of backup management resources and databases) for the databases of other Storage Servers, Storage Management Servers, and tape servers, and collects the backup data and management information created by these backup commands.

The servers 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: Storage Management Server and Storage Server



Example: Assume that the package backup of management information function has been used to back up the management information on a Storage Management Server, Storage Server A, and Storage Server B onto Storage Server B in the diagram above. In this case, Storage Server B is the management information package backup execution server. The Storage Management Server, Storage Server A, and Storage Server B are the management information package backup target servers. That is, Storage Server B operates as both the management information package backup execution server and a management information package backup target server.

The following table shows the types of management information that is handled in a package backup of management information.

Table: Types of management information handled in a package backup of management information

Function	Type of data
Database	Database
Authentication feature	Authentication management list
Registry	Registry
Backup management	Backup management list, or backup management list and recovery control file Pre-processing and post-processing scripts Tape-copy 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 Tape-copy backup volume locking 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
Tape backup management	Tape management information



Note

Storage Server management information created by Softek AdvancedCopy Manager 10.3 (V10.0L40) or earlier cannot be handled by the package backup of management information function. To create a backup in this case, back up each management information component individually as described in "Maintaining a database" in the Operator's Guide applicable to the Storage Server.

8.1.2.1 Preliminaries

Before creating a package backup of management information, you must perform the following steps:

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 information on the backup data size, refer to "Maintaining a database" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide" for the operating system running on the Storage Management Server.
 - The required size for an authentication management list is the size of the authentication management list backup data.
For information on the backup data size, refer to "Maintaining a database" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide" for the operating system running on the Storage 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 information on the backup data size, refer to "Maintaining a database" in the "ETERNUS SF AdvancedCopy Manager Operator's 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 information on the size of each component of backup data, refer to "Maintaining a database" in the "ETERNUS SF AdvancedCopy Manager Operator's 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
 - Sizes of the pre-processing and post-processing scripts for tape-copy
 - 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
 - Size of the backup volume locking specification file for backups for tape copy
 - 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
 - The required size for the policy file for package backup of management information is the size of the policy file backup.
 - The size required for tape management information is 1.2 times the size of backup of the tape management information.
For details on the size of tape management information, refer to "Maintaining Databases" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide for Tape Backup Option".

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 cluster 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 Operator's Guide" for the operating system running on the Storage Management Server.
 - The required size for the authentication management list is the size of the authentication management list backup data.
For the backup data size, refer to "Maintaining a database" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide" for the operating system running on the Storage 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 Operator's 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.
 - The size required for tape management information is 2.2 times the size of the backup of tape management information.
For details on the size of tape management information, refer to "Maintaining Databases" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide for Tape Backup Option".
5. Execute the stgmgrinfo command as described in "Management information package backup server registration/deletion command (stgmgrinfo)", 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: Policy files

Management information package backup execution server		Policy file path
OS	Operation type	
Windows	Non-cluster	environment-setting-directory\etc\stgmgrinfo_policy.ini
	Cluster	<AdvancedCopy-Manager-shared-data>:\etc\opt\swstorage\etc\stgmgrinfo_policy.ini
Solaris OE	Non-cluster	/etc/opt/FJSVswstf/stgmgrinfo_policy.ini
HP-UX Linux AIX	Cluster	etc/opt/FJSVswstf/<logical-node-name>/stgmgrinfo_policy.ini



Note

- Do not specify more than one Storage Server as a management information package backup execution server.
- When this server is used for executing transactions on the Storage Server in a cluster system, create a directory for storing backup data on a shared disk.

- For databases, backup management, authentication management lists, and tape backup management, backup commands are executed on the associated "management information package backup" target servers. For database backups, backup management (ie, backup management lists) backups, and tape backup management information 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: Policy file settings

Section name	Key name	Description	Whether specifiable on		
			Storage Management Server	Storage Server	Tape server
Server name	DBData	Specifies whether to create a backup of the database. "Yes" Create a backup of the database. "No" Do not create a backup of the database. If this key is not specified or if "Yes" is not in the form "Y" (uppercase) + "es" (lowercase), then "No" is set as a default.	A	C	C
	AuthData	Specifies whether to create a backup of the authentication feature. "Yes" Create a backup of the authentication feature. "No" Do not create a backup of the authentication feature. If this key is not specified or if "Yes" is not in the form "Y" (uppercase) + "es" (lowercase), then "No" is set.	A	C	C
	BkupData	Specifies whether to create a backup of the backup management. "Yes" Create a backup of the backup management. "No" Do not make a backup of the backup management. If this key is not specified or if "Yes" is not in the form "Y" (uppercase) + "es" (lowercase), then "No" is set.	B	B	C
	ReplData	Specifies whether to create a backup of the replication management. "Yes" Create a backup of the replication management. "No" Do not create a backup of the replication management. If this key is not specified or if "Yes" is not in the form "Y" (uppercase) + "es" (lowercase), then "No" is set.	B	B	C

Section name	Key name	Description	Whether specifiable on		
			Storage Management Server	Storage Server	Tape server
	TapeData	Specifies whether to create a backup of tape backup management information. "Yes" Create a backup of tape backup management information. "No" Do not create a backup of tape backup management information. Note that failing to specify "Yes" with the correct capitalization ("Y" in uppercase or "es" in lowercase) or failing to specify the key implicitly sets the value to "No."	C	C	B

- A: Required,
- B: Required only when the applicable function is operating on the server,
- C: Not Applicable

Example: Set up a policy file as follows:

- The Storage Management Server named ACMMGR creates backups for the database, the authentication feature, and backup management.
- The Storage Server named ACMAGT1 creates backups for backup management.
- The Storage Server named ACMAGT2 creates backups for replication management.
- The tape server named ACMTAPE makes a tape backup management backup.

```
[ACMMGR]
DBData=Yes
AuthData=Yes
BkupData=Yes
[ACMAGT1]
BkupData=Yes
[ACMAGT2]
Rep1Data=Yes
[ACMTAPE]
TapeData=Yes
```



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 the management information package backup command (stgmgrinfobkup) on the management information package backup execution server stores the backup data in the directory under the following structure:

Table: Directory structure and data types

Directory		Data type	
*1	Server name	DBData	Database
		AuthData	Authentication management list
		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 Preprocessing and post-processing scripts for tape-copy
		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 Backup volume locking specification file for tape copy
	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
	TapeData	Tape management list	

*1 is the backup data storage directory specified by the management information package backup server registration/deletion command (stgmgrinforeset)

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 initial operations are as follows:

1. Execute the management information package backup command (stgmgrinfofbkup) on the management information package backup execution server to create backups of the registries on all of the Storage Servers.
2. Execute the management information package backup command (stgmgrinfofbkup) on the management information package backup execution server to create a backup of the policy file.

Normal use operations

Execute the management information package backup command (stgmgrinfofbkup) 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 Operator's 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 Storage Server in the policy file.

2. Execute the management information package backup command (stgmgrinfobkup) on the management information package backup execution server to make a backup of the registry of the Storage Server that has been added.
3. Execute the management information package backup command (stgmgrinfobkup) 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 Storage Server deleted from the policy file.
2. In the backup data directory on the management information package backup execution server, delete the directory having the name of the Storage Server that has been deleted.
3. Execute the management information package backup command (stgmgrinfobkup) 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 Storage 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 Storage Server.
3. On the management information package backup execution server, set up a policy for the renamed Storage Server in the policy file.
4. Execute the management information package backup command (stgmgrinfobkup) on the management information package backup execution server to create a backup of the registry of the renamed Storage Server.
5. Execute the management information package backup command (stgmgrinfobkup) on the management information package backup execution server to create a backup of the policy file.

8.1.2.3 Restoration operations

This section describes how to use the backup data created by the management information package backup command (stgmgrinfobkup) 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 the management information package backup server registration/deletion command (stgmgrinfofet) 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"
```

3. Transfer the registry backup that was copied in step 2 to the target server in binary mode.
4. Restore the registry on the server. For details on restoring a registry, refer to "Restoring a registry" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide" for the operating system running on the target server.

Restoring a database

1. Execute the management information package backup server registration/deletion command (stgmgrinfo) 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>\<storage-management-server-name>\RegiData"
```

[Solaris, Linux, HP-UX, or AIX version]

```
"<directory-determined-in-step1>/<storage-management-server-name>/RegiData"
```

3. Transfer the backup of the database that was copied in step 2 to the Storage Management Server in binary mode.
4. On the Storage Management Server, expand the database backup transferred in step 3. On the Windows version Storage Management Server, expand the database backup by using the "expand backup data" command. For the backup data expansion command, refer to "Backup data expansion command (stgarcext)" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide (Windows)".
On a Solaris/Linux version Storage Management Server, expand the database backup by using the tar command supported by the operating system.
5. On the Storage Management Server, restore the database. For details on restoring a database, refer to "Restoring a database" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide" for the operating system running on the Storage Management Server. The backup data storage directory specified for database restoration is the directory expanded in step 4.

Restoring an authentication feature

1. Execute the management information package backup server registration/deletion command (stgmgrinfo) on the management information package backup execution server to determine which directory contains the backup data.
2. In the management information package backup execution server, copy the backup of the authentication management list to the following directory:

[Windows version]

```
"<directory-determined-in-step1>\<storage-management-server-name>\AuthData"
```

[Solaris, Linux, HP-UX, or AIX version]

```
"<directory-determined-in-step1>/<storage-management-server-name>/AuthData"
```

3. Transfer the backup of the authentication management list copied in step 2 to the Storage Management Server in binary mode.
4. On a Solaris/Linux version Storage Management Server, change the access permission assigned to the backup of the authentication management list transferred in step 3 by using the chmod command supported by the operating system. Set execution permission for the root user, which is the owner.
5. On a Windows version Storage Management Server, copy the backup of the authentication management list transferred in step 3 to "`<AdvancedCopy_Manager_program_directory>\bin`".
6. On the Storage Management Server, restore the authentication management list. For details on restoring the authentication management list, refer to "Restoring an authentication management list" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide" for the operating system running on the Storage Management Server. A definition output batch file and a definition output shell script are the backup data of the authentication management list handled in steps 4 and 5.

Restoring a backup management list and/or recovery control file

1. Execute the management information package backup server registration/deletion command (stgmgrinfo) 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 by using the expand backup data command. For the backup data expansion command, refer to "Backup data expansion command (stgarcext)" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide (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 Operator's 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 backup management and tape-copy pre-processing/post-processing scripts

1. Execute the management information package backup server registration/deletion command (stgmgrinfo) 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 and tape-copy pre-processing/post-processing scripts 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 and tape-copy pre-processing/post-processing scripts copied in step 2 to the target server in binary mode.
4. Transfer the backups of the backup management and tape-copy 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 Operator's Guide" for the operating system running on the target server.
5. On a Solaris/Linux/HP-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 the management information package backup server registration/deletion command (stgmgrinfo) 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 “Pre-processing and Post-processing of Backup and Restoration”, in the “ETERNUS SF AdvancedCopy Manager Operator’s Guide (Windows)” .

Restoring pre-processing and post-processing scripts for replication management

1. Execute the management information package backup server registration/deletion command (stgmgrinfo) 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 files to be restored>\ReplData\Scripts

[Solaris, Linux, HP-UX, or AIX version]

“<directory determined in step1>\<server whose files 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 Operator’s Guide” for the operating system running on the target server.
5. On a Solaris/Linux/HP-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 the management information package backup server registration/deletion command (stgmgrinfo) 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 Operator's Guide (Windows)".

Restoring a policy file for management information package backup

1. Execute the management information package backup server registration/deletion command (stgmgrinfo) 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 step 1>\<server whose files to be restored>\MgrInfoData

[Solaris, Linux, HP-UX, or AIX version]

"<directory determined in step 1>/<server whose files 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 section "Preliminaries".

Restoring the tape management information

1. On the package-backup-of-the-management-information execution server, execute the management information package backup server registration/deletion command (stgmgrinfo) **to check the directory that contains the backup data.**
2. On this server, copy the backup data of the tape management information from the directory indicated below:

[For the Windows version]

"<directory checked at step 1>\<name of the target server of restoration>\TapeData

[For the Solaris OE/Linux/HP-UX/AIX version]

"<directory checked at step 1>/<name of the target server of restoration>/TapeData

3. Transfer the backup data of the tape management information copied in Step 2 to the target server in binary mode.
4. Expand the tape management information transferred in Step 3 on the target server. If the target server is the Solaris OE or Linux version, expand the tape management information by using the OS tar command.
5. On the target server, restore the tape management information. For details on how to restore it, refer to "Restoring the tape management information" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide for Tape Backup Option".

8.2 Troubleshooting

8.2.1 Insufficient free space in a repository

For information about the actions to take if there is insufficient free space in a repository used by a Storage Management Server, refer to the version of the following manual that is applicable to the OS of the Storage Management Server:

"Troubleshooting for insufficient free space in a repository" in "ETERNUS SF AdvancedCopy Manager Operator's Guide"

8.2.2 Collecting maintenance information

8.2.2.1 Collecting information using a Web GUI client window

Collect information as follows:

1. Start the Java Plug-in control panel, as follows:
Select [Settings] -> [Control Panel] -> [Java Plug-in].
2. Create the startup settings of the Java console.
On the Java control panel, access the [Basics] tag, mark the [Display Java Console] check box, click the [Apply] button, then exit the Java Plug-in control panel.
3. Start the Web GUI window of AdvancedCopy Manager.
4. The Java console will be started and a message will be displayed. Save the message to a file, and then attach the file to your report.

8.2.2.2 Collecting error survey information

Log into the Storage Management Server and the Storage Server. Collect the following information. (When the Storage Management Server is also the Storage Server, collect resources from both servers).

1. When an error occurs during backup management

Table: Error analysis resources to collect when a backup management error occurs

Type	Storage Management Server	Storage Server
System log	Must be collected.	Must be collected.
Trace log	Acquire the files under the following directory: /var/opt/FJSVswtrc/trc/ Acquire the file when the following file exists. /tmp/f3cvtrc.log	Acquire the files under the following directory: /var/opt/FJSVswtrc/trc/ Acquire the file when the following file exists. /tmp/f3cvtrc.log
Backup management list	Need not be collected.	Acquire the files under the following directory: /etc/opt/FJSVswsts (Note)
Common information	Acquire the files under the following directory: /var/opt/FJSVswstf (Note)	Acquire the files under the following directory: /var/opt/FJSVswstf (Note) /etc/opt/swstorage (Note)

Note: When cluster operation is used, the backup management list and the common information are located on the shared disk, where a symbolic link is setup. Collect information from the shared disk as follows.

```
# cd /etc/opt
# tar -cvfh /tmp/FJSVswsts_nodel.tar FJSVswsts swstorage
```

2. When an error occurs during replication management

Table: Error analysis resources to collect when a replication management error occurs

Type	Storage Management Server	Storage Server
System log	Must be acquired.	Must be acquired.
Trace log	Get the files under the following directory: /var/opt/FJSVswtrc/trc/ Acquire the file when the following file exists. /tmp/f3cvtrc.log	Get the files under the following directory: /var/opt/FJSVswtrc/trc/ Acquire the file when the following file exists. /tmp/f3cvtrc.log
Replication management list	Replication management list	Get the files under the following directory: /etc/opt/FJSVswsrp (Note)
Common information	Get the files under the following directory: /var/opt/FJSVswstf (Note)	Get the files under the following directory: /var/opt/FJSVswtf (Note) /etc/opt/swstorage (Note)

Note: When cluster operation is in use, the replication management list and the common information are located on the shared disk, where a symbolic link is set. Collect information from the shared disk as follows.

```
# cd /etc/opt
# tar -cvfh /tmp/FJSVswsrp_nodel.tar FJSVswsrp swstorage
```

8.2.3 Failure in accessing a repository database

For information about how to find the cause of a failure in accessing a repository database and the actions to perform when this occurs, refer to the version of the following manual that is applicable to the OS of the Storage Management Server:

“Troubleshooting for a failure in accessing a repository database” in “ETERNUS SF AdvancedCopy Manager Operator’s Guide”

8.3 Troubleshooting a 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:

- The backup execution status display command (swstbackstat) displays "failed" in the Status column.
- The synchronous backup status display command (swstsyncstat) displays "failed" in the Status column.
- The restoration execution status display command (swstreststat) displays "failed" in the Status column.
- The history information display command (swsthistdisp) displays "failed" in the Status column.
- The tracking status display command (swstrkstat) displays "failed" in the Status column.
- Code 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 code 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 AdvancedCopy to fail.

If a hardware error occurs during remote copying, the following symptoms may be displayed:

- The backup execution status display command (swstbackstat) displays "halt" in the Status column.
- The synchronous backup status display command (swstsyncstat) displays "halt" in the Status column.
- The restoration execution status display command (swstreststat) displays "halt" in the Status column.
- The history information display command (swsthistdisp) displays "halt" in the Status column.


If the execution status display command and history information display command display "halt" in the Status column, a hardware error has occurred during physical copying, and the remote copy processing has failed.

If these symptoms are 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, resolve 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: Troubleshooting flow for hardware and similar faults



Note

- Refer to "Overview" for details of the "status column" and "fault location".
- Use ETERNUSmgr to check the copy process error codes. Use the following method to check the error codes:
 - On the [Display status] menu, click [AdvancedCopy 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: 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 disk array device automatically changes the copy processing to the error suspend state. This prevents 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 disk array units:

- ETERNUS4000/ETERNUS8000 (firmware version V11L30-0000 or later)
- ETERNUS6000 (firmware version V31L40-0000 or later)

For disk array units 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 Troubleshooting a hardware error on a transaction volume

When a hardware fault occurs in transaction volume, perform the following procedures to repair the fault:

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: History information deletion command (swsthistdel)
 - If an error occurred during synchronous processing (EC): Cancel backup of synchronous processing command (swstcancelsync)
 - If an error occurred during restore (OPC) physical copy: Cancel restore command (swstcancelrest)
 - If an error occurred during tracking processing (QuickOPC): Cancel tracking command (swstcanceltrk).
2. Execute the execution status display commands (swstbackstat, swstsyncstat, swstreststat, swsttrkstat), and history information display command (swsthistdisp) to verify that no other errors have occurred.
3. Fix the hardware error in the transaction volume.
4. Execute the restoration execution command (swstrestore) to restore the transaction volume.



Note

- The backup history information for which copy failed cannot be used to restore the volume.
- If there is no normal (status = "succeeded") backup history information, the volume cannot be restored.

8.3.2.2 Troubleshooting a hardware error on a backup volume

When a hardware fault occurs in backup volume, perform the following to repair the fault.

1. Cancel the processing in which the error occurred, with one of the following commands:
 - If an error occurred during backup (OPC) physical copy: History information deletion command (swsthistdel)
 - If an error occurred during synchronous processing (EC): Cancel backup of synchronous processing command (swstcancelsync)
 - If an error occurred during restoration (OPC) physical copy: Cancel restore command (swstcancelrest)
 - If an error occurred during tracking processing (QuickOPC): Cancel tracking command (swstcanceltrk).
2. If the processing cannot be canceled by using one of the above commands, use ETERNUSmgr/GRmgr to cancel it.
3. Execute the resource match command (swstsrsemch)
4. Cancel the processing in which the error occurred. At this point, the execution status display commands (swstbackstat, swstreststat) and history information display command (swsthistdisp) 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 canceled, so no response is required.
 - If an error occurred during backup (OPC) physical copy: History information deletion command (swsthistdel)
 - If an error occurred during restoration (OPC) physical copy: Cancel restore command (swstcancelrest).

5. Execute the execution status display commands (swstbackstat, swstsyncstat, swstreststat, swsttrkstat) and history information display command (swsthistdisp) to verify that no other errors are found.
6. Use the device information setting command (swstdevinforeset) to delete the backup volume in which the error occurred.
7. Use the device information setting command (swstdevinforeset) 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 GUI client and store the information again in the backup volume.
8. 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): Cancel backup of synchronous processing command (swstcancelsync). If the bad sector occurred during QuickOPC tracking: Cancel tracking command (swstcanceltrk).
2. Execute the execution status display commands (swstbackstat, swstsyncstat, swstreststat, swsttrkstat) and the history information display command (swsthistdisp) 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 restoration methods below and perform restoration.
 - 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 the restore command (swstrestore) 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 if there is no backup history information.

8.3.2.4 Troubleshooting errors (halt) on a remote copy processing

1. Fix the hardware error.
2. Re-execute the suspended copy processing.

8.4 Troubleshooting a Hardware Error or Other Error Occurring 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 a replication, the following symptoms are displayed:

- The operation status display command (swsrpstat) displays "failed" or "?????" in the Status column.
- The error Code swsrp2606 is output during command execution.
- The error Code 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 either 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 AdvancedCopy to fail.

If a hardware error occurs during remote copying, the following symptoms may be displayed:

- The operation status display command (swsrpstat) displays "halt" in the Status column.
- The error Code swsrp2648 is displayed during command execution.

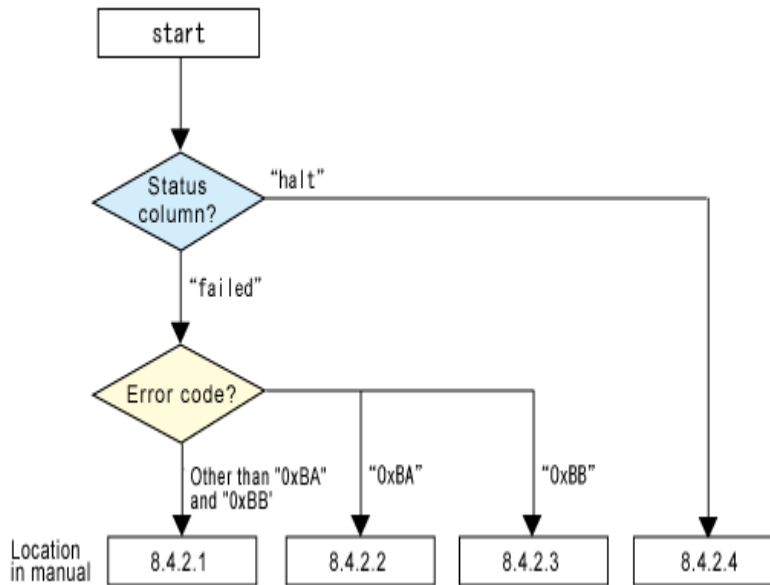
If the operation status display command displays "halt" in the Status column, and swsrp2648 is displayed during command execution, a hardware error has occurred during physical copying, and the remote copy processing has failed.

If any of these symptoms are displayed, a hardware error may have occurred in the source or replica volume. Check the hardware status, and if a hardware error is found, resolve 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: Troubleshooting flow when a hardware or similar fault occurs



Note

4. Refer to "Overview" for details of the "status column".

If the status column is "????", check if ETERNUSmgr/GRmgr 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, another cause (Storage Server, switch, and so on) is likely, so contact your Fujitsu SE.

- The error code can be checked with either of the following methods:
 - Checking method using Operation status display command (swsrpstat):
In the options for Operation status display command (swsrpstat) specify "-0", then execute.
 - Checking method using ETERNUSmgr:
 - 1. On the [Display status] menu, click [AdvancedCopy 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: Meanings of error codes]

Error code	Meaning
0xBA	If a) or b) below applies, a bad sector was created in the copy source 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 copy source volume when a) or b) applies, the disk array unit automatically changes the copy processing to the error suspend state. This prevents restart of QuickOPC or EC/REC resume and prevents the copy destination volume from being overwritten with invalid copy source volume data.
0xBB	An insufficiency 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 disk array units:

- ETERNUS4000/ETERNUS8000 (firmware version V11L30-0000 or later)
- ETERNUS6000 (firmware version V31L40-0000 or later)

For disk array units other than the above, the events indicated by error code 0xBA and 0xBB are identified by the following methods:

[Table: 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 ETERNUSmgr to check the capacity already used on the Snap Data Volume in order to determine whether or not an insufficiency of free space has occurred: <ul style="list-style-type: none"> ● 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 " <u>Troubleshooting when an insufficiency of free space has occurred in the Snap Data Volume or Snap Data Pool</u> ".

8.4.2.1 Troubleshooting a hardware error on a replication volume

When a hardware error occurs in a duplicate volume, perform the following to repair the fault.

1. Use the replication cancel command (swsrpcancel) to cancel the processing in which the error occurred. If the processing cannot be canceled from the operation server when interserver replication is performed, cancel it from a non-operation server.
2. If the processing cannot be canceled by using the command, use ETERNUSmgr/GRmgr to cancel it.
3. Execute the resource adjustment command (swsrprecoverres)
4. Execute the operation status display command (swsrpstat) to verify that no other errors have occurred.
5. Use the replication volume information deletion command (swsrpdelvol) to delete the replication volume in which the error occurred.
6. Use the replication volume information configuration command (swsrpsetvol) 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 GUI client 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 the replication cancellation command (swsrpcancel) 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 ETERNUSmgr to cancel it.
2. Execute the operation status display command (swsrpstat) 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 restoration methods below and perform restoration.
 - 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 the replica creation command (swsrpmake) 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 an insufficiency of free space has occurred in the Snap Data Volume or Snap Data Pool

A Snap Data Volume insufficiency of free space occurs when the Snap Data Pool is not being used, whereas a Snap Data Pool insufficiency of free space occurs when the Snap Data Pool is being used.

The following section provides information on recovery when an insufficiency 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 an insufficiency 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 the Cancel replication command (swsrpcancel) 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 ETERNUSmgr to cancel it.

The likely causes of an insufficiency 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 appropriate.
- b. The estimate of the physical size of the Snap Data Volume is appropriate 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 size of the Snap Data Volume is being unnecessarily consumed.

The use status of the Snap Data Volume can be checked by specifying "stat" in the Snap Data Volume operation/reference command (swstsdv) subcommand.

If "a." applies, re-estimate the physical size of the Snap Data Volume, and recreate the Snap Data Volume.

If "b." applies, use ETERNUSmgr or, specify "init" in the Snap Data Volume operation/reference command (swstsdv) 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 an insufficiency 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 the Cancel replication command (swsrpcancel).
- 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 ETERNUSmgr to cancel it.

The following are likely causes of an insufficiency of free space in the Snap Data Pool:

- c. The estimate of the size of the Snap Data Pool is not appropriate.
- d. The estimate of the size of the Snap Data Pool is appropriate but, as a result of a large volume being updated in the Snap Data Volume when a SnapOPC/SnapOPC+ session does not exist, the size of the Snap Data Pool is being wastefully consumed.

The use status of the Snap Data Pool can be checked by specifying "poolstat" in the Snap Data Volume operation/reference command (swstsdv) 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 ETERNUSmgr or, specify "init" in the Snap Data Volume operation/reference command (swstsdv) 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 Troubleshooting error “halt” on a remote copy processing

The REC restart (Resume) method differs for different REC Recovery modes.

[For Automatic Recovery mode]

1. Remove the cause that made all paths close (halt).
2. The ETERNUS disk array automatically restarts (Resumes) REC.

[For Manual Recovery mode]

1. Remove the cause that made all paths close (halt).
2. Use the replication creation command (swsrpmake) 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 the replication start command (swsrpstartsync) to restart (Resume) REC. (To restart REC after a forcible suspend, the -t option must be specified.)

[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

The “Server information package” changes operational server information as a batch while maintaining consistency between the server information in all systems managed by AdvancedCopy Manager. Therefore, the daemons and transactions of all Storage Servers in the systems that include Storage Management Server transactions must be running.



Note If the server information change command (`stgxfwcmmodsrv`) is used to change the IP address, etc.

Either before or during execution of the server information change command (`stgxfwcmmodsrv`), do not execute backup management commands or replication management commands in any Storage Servers that are managed by a Storage Management Server that executes the server information change command (`stgxfwcmmodsrv`) or by Storage Management Server transactions that execute the server information change command (`stgxfwcmmodsrv`).

If these commands are executed, commands may terminate abnormally.

If the server information change command (`stgxfwcmmodsrv`) terminates abnormally, take the action indicated in the error message to perform recovery.

If a backup management or replication management command is accidentally executed either before or during execution of the server information change command (`stgxfwcmmodsrv`) and subsequently terminates abnormally, wait until the server information change command (`stgxfwcmmodsrv`) terminates normally, then re-execute the backup management or replication management command.

8.5.1 Changing the IP address of a Storage Management Server

For information about changing the IP address of a Storage Management Server, refer to “Changing the IP address of a Storage Management Server” in the “ETERNUS SF AdvancedCopy Manager Operator’s Guide” relevant to the OS of the Storage Management Server:

8.5.2 Changing the IP address of a Storage Server

Change the IP address of a Storage Server as follows:

1. If replication management functions were used to define it as a copy source or copy destination volume, delete the copy source or copy destination volume settings. Refer to “Deleting a copy source or copy destination volume” for the method for deleting the copy source or copy destination volume settings.
2. On a target Storage Server, stop the daemons of AdvancedCopy Manager. For more information on stopping daemons, refer to “Starting and Stopping Daemons”.
3. On the target Storage Server, change the system IP address.
4. On the target Storage Server, start the daemons. For more information on starting daemons, refer to “Starting and Stopping Daemons”.
5. Create a server information change instruction file on the Storage Management Server, then execute the server information change command with the “-f” option specified to change the IP address. Refer to “server information change command (`stgxfwcmmodsrv`)” in the “ETERNUS SF AdvancedCopy Manager Operator’s Guide” for the OS under which the Storage Management Server is running for the details.
6. If the replication management functions are used, reconstruct the replication management environment (perform the “Setting copy source volume and copy destination volume” task and subsequent tasks).



Note

For information on change when in a cluster operation, refer to “Changing the IP address of a Storage Server transaction”.

8.5.3 Changing the port number of a Storage Management Server or Storage Server

For information about changing the port number of a Storage Management Server, refer to “Changing the port number of a Storage Management Server or Storage Server” in the “ETERNUS SF AdvancedCopy Manager Operator’s Guide” version that is applicable to the OS of the Storage Management Server.

Change the port number of a Storage Server as follows:

1. If replication management functions were used at the target Storage Management Server or Storage Server to define it as a copy source or copy destination volume, delete the copy source or copy destination volume settings.
Refer to “Deleting a copy source or copy destination volume” for the method for deleting the copy source or copy destination volume settings.
2. On the target Storage Server, change the port number of the AdvancedCopy Manager communications daemon (stgxfws).
3. On the target Storage Server, restart the daemons of AdvancedCopy Manager. For more information on restarting a daemon, refer to “Starting and Stopping Daemons”.
4. Create a server information change instruction file on the Storage Management Server, then execute the server information change command with the “f” option specified to change the IP address. Refer to “server information change command (stgxfwcmmodsrv)” in the “ETERNUS SR AdvancedCopy Manager Operator’s Guide” of the OS running the Storage Management Server for the details.
5. If the replication management functions are used, reconstruct the replication management environment (perform the “Setting copy source volume and copy destination volume” task and subsequent tasks).



Note

For information about performing the above change in cluster operation, refer to “Changing the port number of a Storage Server transaction”.

8.5.4 Changing the server name of a Storage Management Server or Storage Server

For information about changing the server name of a Storage Management Server, refer to the “Changing the server name of a Storage Management Server or Storage Server” in the “ETERNUS SF AdvancedCopy Manager Operator’s Guide” relevant to the OS of the Storage Management Server:

Change the server name of a Storage Server as follows:

1. If replication management functions were used at the target Storage Management Server or Storage Server to define it as a copy source or copy destination volume, delete the copy source or copy destination volume settings.
Refer to “Deleting a copy source or copy destination volume” for the method for deleting the copy source or copy destination volume settings.
2. Create a server information change instruction file on the Storage Management Server.

3. Execute the server information change command with the “-f” option specified (specify the created file) to change the server name. Refer to “Server information change command (stgxfwcmmodsrv)” in the “ETERNUS SF AdvancedCopy Manager Operator’s Guide” relevant to the OS running the Storage Management Server for details of the server information change command.
4. If the replication management functions are used, reconstruct the replication management environment (perform the “Setting copy source volume and copy destination volume” task and subsequent tasks).



Note

For information about performing this change when in cluster operation, refer to “Changing the server name of a Storage Server transaction”.

8.5.5 Changing the device configuration

Changing the device configuration refers to any of the following tasks:

- Mount name change
- Change from File System operation to RawDevice operation
- Capacity change of partitions (mainly for the recreation of LUN mapping)
- Start position for change of partitions (mainly for the recreation of LUN mapping)

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 “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 “Deleting a backup volume”.
3. Change the device configuration. If you have deleted a logical drive in the primary partition or extended partition, be sure to 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 “Fetching device information on a Storage Server” for details on changing the device information.
5. Reset the deleted transaction volume or backup volume in the backup. Refer to “Configuring the operation type for a device,” for details on adding the transaction volume or backup volume.



Point

If the device configuration is changed (in step 3 above) before the transaction volume and backup volume 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 operation mode, and delete the history information or cancel the synchronized backup processing.

For details on the emergency operation mode, refer to “history information deletion command (swsthistdel)” and “backup synchronous processing cancel command (swstcancelsync)” .

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 and delete them. For information about this, refer to "Deleting a source volume or replica volume."
2. Change the device configuration. If you have deleted a logical drive in the primary partition or extended partition, be sure to 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 "Fetching device information from a Storage Server," for details about changing the device information.
4. Reconfigure a deleted replica volume in the replication. Refer to "Configuring the source volume and replica volume," for details about adding the replica volume.



If the device configuration is changed (in step 2 above) before the replication volumes are deleted (step 1), replication processing cannot be cancelled. This could possibly prevent the deletion of the replication volumes.

In this case, cancel replication processing in emergency operation mode. For details of the emergency operation mode, refer to "replication volume information deletion command (swsrpdelvol)".

8.5.6 Changing the IP address for GUI client connection of a Storage Management Server

For information about changing the server name of a Storage Management Server, refer to the "Changing the IP address for GUI connection of a Storage Management Server" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide" relevant to the OS of the Storage Management Server:

8.5.7 Change of password or deletion of the repository access user

For information about Change of password or deletion of the user who was specified as the repository access user, refer to "Change of deletion/password of the user who specified it as the repository access user" in "ETERNUS SF AdvancedCopy Manager Operator's Guide" relevant to the OS of the Storage Management Server:

8.5.8 Changing the operational configuration in cluster operation

This section describes how to make a change to the operational configuration in cluster operation.

8.5.8.1 Changing the IP address of a Storage Server transaction

Change the IP address of a Storage Server transaction by performing the following steps:

1. Check if the target Storage Management Server transaction is also used for Storage Server transactions.
If it is also used for Storage Server transactions and replication management functions were used to define it as a copy source or copy destination volume, delete the copy source or copy destination volume settings.
Refer to "Deleting a copy source or copy destination volume" for the method for deleting the copy source or copy destination volume settings.
2. Stop the Storage 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 Storage 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 on the Storage Management Server that manages the Storage Server transactions, then execute the server information change command with the "-f" option specified to change the IP address. Refer to "Server information change command (stgxfwcmmodsrv)" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide" relevant to the OS running the Storage Management Server for details of the server information change command.
6. If the replication management functions are used, reconstruct the replication management environment (perform the "Setting copy source volume and copy destination volume" task and subsequent tasks).



Note

The IP address displayed by the cluster setup command executed to delete the cluster environment of the Storage Server transaction is not the IP address newly specified, but the IP address specified by the cluster setup command.

8.5.8.2 Changing the port number of a Storage Server transaction

Change the port number of a Storage Server transaction by performing the following steps:

1. If replication management functions were used to define it as a copy source or copy destination volume, delete the copy source or copy destination volume settings.
Refer to "Deleting a copy source or copy destination volume" for the method for deleting the copy source or copy destination volume settings.
2. Stop the target Storage Server transaction.
For details on how to stop it, refer to the relevant cluster software manual.
3. On the primary node and secondary node, change the port number of the transaction-intended communication daemon (stgxfws_<logical node name>) specified in /etc/services.
4. Start the target Storage Server transaction.
For details on how to start it, refer to the relevant cluster software manual.
5. Create a server information change instruction file on the Storage Management Server that manages the Storage Server transactions, then execute the server information change command with the "-f" option specified to change the port number. Refer to "Server information change command (stgxfwcmmodsrv)" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide" relevant to the OS running the Storage Management Server for details of the server information change command.
6. If the replication management functions are used, reconstruct the replication management environment (perform the "Setting copy source volume and copy destination volume" task and subsequent tasks).

8.5.8.3 Changing the server name of a Storage Server transaction

Change the server name of a Storage Server transaction by performing the steps below.

1. If replication management functions were used by the target Storage Management Server or Storage Server to define it as a copy source or copy destination volume, delete the copy source or copy destination volume settings.
Refer to "Deleting a copy source or copy destination volume" for the method for deleting the copy source or copy destination volume settings.
2. Create a server information change instruction file on the Storage Management Server that manages the Storage Server transactions.
3. Execute the server information change command with the "-f" option specified (specify the created file) to change the server name. Refer to "Server information change command (stgxfwcmmodsrv)" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide" relevant to the OS running the Storage Management Server for details of the server information change command.
4. If the replication management functions are used, reconstruct the replication management environment (perform the "Setting copy source volume and copy destination volume" task and subsequent tasks).

8.5.8.4 Changing the logical node name

Change the logical node name of a Storage Management Server transaction or a Storage Server transaction by performing the following steps:

1. Check whether the target Storage Management Server transaction or Storage Server transaction is operating.
If the transaction is not operating, start the Storage Management Server transaction or Storage Server transaction.
For details of how to start it, refer to the relevant cluster software manual.
2. Delete the configuration of the Storage Server transaction.
When using the backup management function, delete the transaction volume and backup volume. For details on how to perform this, refer to "Deleting a device".
When it is defined as a source or replica volume by using the replication control function, delete the source/replica volume configuration. For details of how to do this, refer to "Deleting a source volume or replica volume".
3. Save the AdvancedCopy Manager data.
In the Storage Server transaction or the Storage Management Server transaction that also functions as a Storage Server transaction, save the following data:
 - Pre-processing/post-processing scripts
4. Save the pre-processing and post-processing scripts for backup management and replication control.
For the server that executes the package backup of the management information of AdvancedCopy Manager, save the following data:
 - Policy file for package backup of the management information
For details on the policy file path, refer to "Preliminaries".
5. For a Storage Server transaction (except a Storage Management Server transaction that also functions as a Storage Server transaction), delete /etc/opt/swstorage/<logical node name>/swnode.ini on the node operating the transaction.
6. Delete the cluster environment of the target transaction.
For details about how to do this, refer to "Deleting a Storage Management Server transaction or Storage Server transaction" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide for cluster environment".



- Do not perform the following operation in the "ETERNUS SF AdvancedCopy Manager Operator's Guide for cluster environment":
 - Step 1 (deleting the server) described in "Detailed steps for deletion" in "Deletion on MC/ServiceGuard for HP-UX"

- Step 1 (deleting the server) described in “Detailed steps for deletion” in “Deletion on VERITAS Cluster Server for HP-UX”
7. Create the cluster environment of the target transaction.
For details of how to create a cluster environment, refer to “Customizing a Storage Management Server transaction or Storage Server transaction” in the “ETERNUS SF AdvancedCopy Manager Operator’s Guide for cluster environment” .



Note

- For the “-n” option of the cluster setup command, specify the name of the new logical node.
 - Do not perform the steps described in “Creating a Storage Management Server transaction or Storage Server transaction environment” in “Detailed steps for customization” on creating a cluster environment:
 - Step 18 in “Details of the steps for customization” on MC/ServiceGuard for HP-UX
 - Step 18 in “Details of the steps for customization” on VERITAS Cluster Server for HP-UX
8. Start the target transaction.
For details on how to start it, refer to the applicable cluster software manual.
 9. Restore the AdvancedCopy Manager data.
In the Storage Server transaction or Storage Management Server transaction that also functions as a Storage Server transaction, restore the following data:
 - Pre-processing/post-processing scripts
Restore the pre-processing and post-processing scripts for backup management and replication control.
 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 “Preliminaries”.
 10. Change the server name.
For a Storage Server transaction, perform “Changing the server name of a Storage Server transaction”.
 11. To use the backup management or replication control function, rebuild the backup management or replication control environment. (Perform the “[Environment settings for the back up server](#)” task and subsequent tasks.)
 12. For the “package backup of the management information” execution server of AdvancedCopy Manager, execute the “package backup of the management information” server registration/deletion command to register the server.

8.5.8.5 Changing the service name

Change the name of an AdvancedCopy Manager transaction by taking the following steps:

1. Check whether the target Storage Management Server transaction or Storage Server transaction is operating.
If the transaction is not operating, start the Storage Management Server transaction or Storage Server transaction.
For details about how to start it, refer to the applicable cluster software manual.
2. Save the AdvancedCopy Manager data.
In the Storage Server transaction or Storage Management Server transaction that also functions as a Storage Server transaction, save the following 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 make a backup of a backup management list, refer to “Making a backup of backup management list”.
 - Pre-processing/post-processing scripts

Save the pre-processing and post-processing scripts for backup management and replication control.

For the server that executes the package backup of the management information, save the following data:

- Policy file for package backup of the management information
For details about the policy file path, refer to "Preliminaries".

3. For a Storage Server transaction (except a Storage Management Server transaction that also functions as a Storage Server transaction), 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 "Deleting a Storage Management Server transaction or Storage Server transaction" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide for cluster environment".



Note

- Do not perform the following operation in the "ETERNUS SF AdvancedCopy Manager Operator's Guide for cluster environment":
 - Step 1 (deleting the server) described in "Detailed steps for deletion" in "Deletion on MC/ServiceGuard for HP-UX"
 - Step 1 (deleting the server) described in "Detailed steps for deletion" in "Deletion on VERITAS Cluster Server for HP-UX"
5. On the cluster system, change the name of the userApplication.
For details on how to do this, refer to the applicable cluster software manual.
 6. Create the cluster environment of the target transaction.
For details on how to do this, refer to "Customizing a Storage Management Server transaction or Storage Server transaction" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide for the cluster environment".



Note

- Do not perform the steps described in "Creating a Storage Management Server transaction or Storage Server transaction environment" in "Detailed steps for customization" about creating a cluster environment.
 - Step 18 in "Details of the steps for customization" on MC/ServiceGuard for HP-UX
 - Step 18 in "Details of the steps for customization" on VERITAS Cluster Server for HP-UX
 - Do not change any information other than the name of a transaction by using the cluster setup command.
7. Start the target transaction.
For details on how to do this, refer to the applicable cluster software manual.

8. Restore the AdvancedCopy Manager data.
 - a) In the Storage Server transaction or Storage Management Server transaction that also functions as a Storage Server transaction, restore the following 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 on how to do this, refer to "Restoring the backup management list".
 - Pre-processing/post-processing scripts
Restore the pre-processing and post-processing scripts for backup management and replication control.
 - When using the replication control function, restore the environment by using the resource adjustment command (swsrprecoverres).
Specify the "-r" option for the resource adjustment command. For details about this, refer to "Resource adjustment command (swsrprecoverres)".
 - b) 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 on the policy file path, refer to "Preliminaries".
9. For the "package backup of the management information" execution server of AdvancedCopy Manager, execute the "package backup of the management information" server registration/deletion command to register the server.

8.5.8.6 Changing a shared disk for shared data

Change a shared disk that is used for AdvancedCopy Manager shared data (e.g., the name of the mount point of a physical disk or shared disk) by performing the following steps:

1. Check whether the target Storage Management Server transaction or Storage Server transaction is operating.
If the transaction is not operating, start the Storage Management Server transaction or Storage Server transaction.
For details on how to start it, refer to the applicable cluster software manual.
2. Save the AdvancedCopy Manager data.
In the Storage Server transaction or Storage Management Server transaction that also functions as a Storage Server transaction, save the following 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 "Making a backup of backup management list".
 - Pre-processing/post-processing scripts
Save the pre-processing and post-processing scripts for backup management and replication control.
 For the execution server of the package backup of the management information 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 "Preliminaries".
3. For a Storage Server transaction (except a Storage Management Server transaction that also functions as a Storage Server transaction), 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 "Deleting a Storage Management Server transaction or Storage Server transaction" in the "ETERNUS SF AdvancedCopy Manager Operator's Guide for a cluster environment".



Note

- Do not perform the following operation in the “ETERNUS SF AdvancedCopy Manager Operator’s Guide for a cluster environment” :
 - Step 1 (deleting the server) described in “Detailed steps for deletion” in “Deletion on MC/ServiceGuard for HP-UX”
 - Step 1 (deleting the server) described in “Detailed steps for deletion” in “Deletion on VERITAS Cluster Server for HP-UX”
5. Create the cluster environment of the target transaction.
For details about how to do this, refer to “Customizing a Storage Management Server transaction or Storage Server transaction” in the “ETERNUS SF AdvancedCopy Manager Operator’s Guide for the cluster environment” .



Note

- To change the name of a mount point, specify the new mount point name in the `-m` option of the cluster setup command.
 - To change a physical disk, specify the new physical disk when you mount it.
 - Change the configuration of a shared disk in the cluster system.
 - Do not perform the steps described in “Creating a Storage Management Server transaction or Storage Server transaction environment” in “Detailed steps for customization” on creating a cluster environment.
 - Step 18 in “Details of the steps for customization” on MC/ServiceGuard for HP-UX
 - Step 18 in “Details of the steps for customization” on VERITAS Cluster Server for HP-UX
6. Start the target transaction.
For details about how to start it, refer to the applicable cluster software manual.
7. Restore the AdvancedCopy Manager data.
- a) In the Storage Server transaction or Storage Management Server transaction that also functions as a Storage Server transaction, restore the following 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 “Restoring the backup management list”.
 - Pre-processing/post-processing scripts
Restore the pre-processing and post-processing scripts for backup management and replication control.
 - When using the replication control function, restore the environment by using the resource adjustment command (`swsrprecoverres`).
Specify the `-r` option for the resource adjustment command. For details about this, refer to “Resource adjustment command (`swsrprecoverres`)”.
 - b) 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 “Preliminaries”.
8. For the “package backup of the management information” execution server of AdvancedCopy Manager, execute the “package backup of the management information” server registration/deletion command to register the server.

Chapter 9 Commands

This chapter provides information on 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: List of environment definition commands for backup management

Function	Command name	Explanation
Storage server configuration information configuration command	swstsvrset	Sets the configuration information of a Storage Server.
Storage server configuration information display command	swstsvrdisp	Displays the configuration information of a Storage Server.
Device information configuration command	swstdevinfo	Sets the usage type of a device.
Device use status display command	swstdevdisp	Displays the information of a device.
Backup policy configuration 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: 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.
Restore execution status display command	swstreststat	Displays the execution status of restoration.
History information display command	swsthistdisp	Displays the history information that has been backed up.
History information deletion command	swsthistdel	Deletes the history information that has been backed up.

Function	Command name	Explanation
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: List of maintenance commands for backup management

Function	Command name	Explanation
Resource adjustment command	swstsrsemth	Recovers the consistency of information in a backup management file.
Resource backup command	swstresback	Backs up a backup management files
Resource restore command	swstresrst	Restores a backup management files

Configuration management commands

Table: List of configuration management commands

Function	Command name	Explanation
Management server information change command	stgcmmodnode	Changes the server information of the Storage Management server
Repository access user change command	stguserst	This command is executed from the Storage Management Server.
GUI CLIENT connection information configuration command	stgguiipset	This command is executed from the Storage Management Server.
Server information addition command	stgxfwcmaddsrv	This command is executed from the Storage Management Server.
Server information change command	stgxfwcmmodsrv	This command is executed from the Storage Management Server.
Server information deletion command	stgxfwcmdelsrv	This command is executed from the Storage Management Server.

Function	Command name	Explanation
Device information fetch/reflect command	stgxfwcmsetdev	This command is executed from the Storage Management Server.
Device information deletion command	stgxfwcmdev	This command is executed from the Storage Management Server.
Server information display command	stgxfwcmdispsrv	This command is executed from the Storage Management Server.
Device information display command	stgxfwcmdispdev	This command is executed from the Storage Management Server.
Partition information display command	stgxfwcmdisppat	This command is executed from the Storage Management Server.
Environment information display command	stgenvdisp	Displays the environment values set for AdvancedCopy Manager
Management information batch backup server registration/deletion command	stgmgrinfofet	Registers servers for execution of batch backup of management information
Management information batch backup command	stgmgrinfobkup	Backs up Storage server and Storage Management server management information
Management information batch backup status display command	stgmgrinfofdisp	Displays the execution status of the management information batch backup command

Replication management commands

Environment definition commands

Table: List of environment definition commands for replication management

Function	Command name	Explanation
Replication volume information configuration command	swsrpsetvol	Sets the information about a replication volume.
Replication volume information display command	swsrpvoinfo	Displays the information about a replication volume.
Replication volume information deletion command	swsrpdelvol	Deletes replication volume information that has been set.
REC transfer buffer information display command	swsrprebuffstat	Displays information for the REC transfer buffer.
REC transfer buffer configuration change command	swsrprebuffset	Changes the configuration for the REC transfer buffer.

Operation commands

Table: 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

Function	Command name	Explanation
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	swsrpchsnc	Changes the operation mode of inter-box synchronization.
Synchronous processing reverse command	swsrprevsnc	Reverses the copying direction of synchronization in suspended status.

Maintenance commands

Table: 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/stop commands

Table: List of daemon start/stop commands

Function	Command name	Explanation
Communication daemon start and stop	stgfwcom	Starts and stops communication daemons
AdvancedCopy Manager daemon start and stop	startacm stopacm	Starts and stops AdvancedCopy Manager

SnapOPC/SnapOPC+ support commands

Table: List of SnapOPC/SnapOPC+ support commands

Function	Command name	Explanation
Update size measurement command	swstestupdate	Uses the pseudo-SnapOPC session configuration function to measure the physical size of a Snap Data Volume disk
<u>Snap Data Volume operation/reference command</u>	swstsdv	Subcommand to initialize a Snap Data Volume, display the status of a Snap Data Volume, or display the status of a Snap Data Pool.

9.2 Backup Management Commands

This section describes operations using backup management commands.



Note

Only a root or superuser can execute all backup management commands. When commands are executed by other users, the message "swst0601 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 operating environment requirements for different commands.

Table: Operating environment requirements for executing commands

Command	Operating environment of AdvancedCopy Manager
swstbackstat	The daemons and services of the Storage Management Server and target Storage Server must be operating in order to execute this command, if the Storage Management Server has either of the following parameter designations: <ul style="list-style-type: none"> • -h option specified, or • No device name specified
swstbackup	To execute this command on a Storage Management Server with the -h option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating. When the Storage Server runs on Windows and is in EC operation, the daemons and services of the Storage Server must be operating.
swstbkpodel	To execute this command on a Storage Server, the daemons and services of the Storage Management Server must be operating. To execute this command on a Storage Management Server with the -h option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating.
swstbkpoldisp	The daemons and services of the Storage Management Server and target Storage Server must be operating in order to execute this command, if the Storage Management Server has either of the following parameter designations: <ul style="list-style-type: none"> • -h option specified, or • No device name specified
swstbkpolset	To execute this command on a Storage Server, the daemons and services of the Storage Management Server must be operating. To execute this command on a Storage Management Server with the -h option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating.
swstcancelrest	To execute this command on a Storage Management Server with the -h option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating.
swstcancelsync	To execute this command on a Storage Management Server with the -h option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating. When the Storage Server runs on Windows and is in EC operation, the daemons and services of the Storage Server must be operating.
swstcanceltrk	When this command is executed with the -h option specified on the Storage Management Server, the daemons/services of the Storage Management Server and target Storage Server must be active.
swstdevdisp	The daemons and services of the Storage Management Server and target Storage Server must be operating in order to execute this command, if the Storage Management Server has either of the following parameter designations: <ul style="list-style-type: none"> • -h option specified, or • No device name specified

Command	Operating environment of AdvancedCopy Manager
swstdevinfo	To execute this command on a Storage Server, the daemons and services of the Storage Management Server must be operating. To execute this command on a Storage Management Server with the -h option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating.
swstexecstat	The daemons and services of the Storage Management Server and target Storage Server must be operating in order to execute this command on a Storage Management Server with either of the following parameter designations: <ul style="list-style-type: none"> • -h option specified, or • No device name specified
swsthistdel	To execute this command on a Storage Management Server with the -h option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating.
swsthistdisp	The daemons and services of the Storage Management Server and target Storage Server must be operating in order to execute this command on a Storage Management Server with either of the following parameter designations: <ul style="list-style-type: none"> • -h option specified, or • No device name specified
swstresback	-
swstresrst	To execute this command on a Storage Management Server with the -h option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating. When the Storage Server runs on Windows, the daemons and services of the Storage Server must be operating. To execute this command with the -x option specified, the daemons and services of the Storage Management Server must be operating.
swstrestore	To execute this command on a Storage Management Server with the -h option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating.
swstreststat	The daemons and services of the Storage Management Server and target Storage Server must be operating in order to execute this command on a Storage Management Server with either of the following parameter designations: <ul style="list-style-type: none"> • -h option specified, or • No device name specified
swstsrsemch	To execute this command on a Storage Management Server with the -h option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating. When the Storage Server runs on Windows, the daemons and services of the Storage Server must be operating. To execute this command with the -x option specified, the daemons and services of the Storage Management Server must be operating.
swststartsync	To execute this command on a Storage Management Server with the -h option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating. When the Storage Server runs on Windows, the daemons and services of the Storage Server must be operating.
swstsvrdisp	To execute this command on a Storage Management Server with the -h option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating.
swstsvrset	To execute this command on a Storage Server, the daemons and services of the Storage Management Server must be operating. To execute this command on a Storage Management Server with the -h option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating.
swstsyncstat	The daemons and services of the Storage Management Server and target Storage Server must be operating in order to execute this command on a Storage Management Server with either of the following parameter designations: <ul style="list-style-type: none"> • -h option specified, or • No device name specified

Command	Operating environment of AdvancedCopy Manager
swstrkstat	When this command is executed in either of the following formats on the Storage Management Server, the daemons/services of the Storage Management Server and target Storage Server must be active: <ul style="list-style-type: none"> • -h option is specified, or • No device name is specified.



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. For "bsh"shell, 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: Cluster operation commands

Command	Environment variable	Command execution node	Remarks
swstsvrset	Required	Active cluster server	
swstsvrdisp	Required	Active cluster server	
swstdevinfofet	Required	Active cluster server	
swstdevdisp	Required	Active cluster server	
swstbkpolset	Required	Active cluster server	
swstbkpoldisp	Required	Active cluster server	
swstbkpoldel	Required	Active cluster server	
swstbackup	Required	Active cluster server	
swstbackstat	Required	Active cluster server	
swstrestore	Required	Active cluster server	
swstcancelrest	Required	Active cluster server	
swstreststat	Required	Active cluster server	
swsthistdisp	Required	Active cluster server	
swsthistdel	Required	Active cluster server	

Command	Environment variable	Command execution node	Remarks
swststartsync	Required	Active cluster server	
swstcancelsync	Required	Active cluster server	
swstsyncstat	Required	Active cluster server	
swstexecstat	Required	Active cluster server	
swstcanceltrk	Required	Active cluster server	
swsttrkstat	Required	Active cluster server	
swstsrsemch	Required	Active cluster server	
swstresback	Required	Active cluster server	
swstresrst	Required	Transaction operation node	



Note

When the display commands that have 'disp' or 'stat' at the end of their name are executed at the Storage Management Server, the content displayed depends on whether or not the -h option and the device name are specified. See the following table for details.

[Table: Contents displayed by the Storage Management Server display commands]

-h option	Device name specification	Displayed content
No	No	All information of all Storage Servers registered to the backup management function is displayed.
	Yes	The information of only the specified device in the Storage Management Server is displayed.
Yes	No	All information of the Storage Server specified in the -h option is specified.
	Yes	The information of only the specified device in the Storage Server specified in the -h option is displayed.

9.2.1 Environment definition commands

This chapter describes environment definition commands for backup management.

9.2.1.1 Storage Server configuration information configuration command (swstsvrset)

This command sets the configuration information of a Storage Server. Execute this command to declare that the Storage Server that executed the command is subject to the backup.

a. Format

```
[To be executed on a storage server]
/opt/FJSVswsts/bin/swstsvrset

[To be executed on a storage management server]
/opt/FJSVswsts/bin/swstsvrset [-h Server]
```

b. Description of options

Option	Description
-h	Specifies the name of the Storage Server to be subject to the backup. This option is valid only on a Storage Management Server. If this option is omitted, the Storage Server on which the command is entered will be subject to the processing.

c. End status

```
= 0: Normal end
> 0: Abnormal end
```

d. Example

Set the configuration information for a Storage Server.

```
# /opt/FJSVswsts/bin/swstsvrset
swstsvrset completed
#
```

Set the configuration information for a Storage Server (job1) from a Storage Management Server.

```
# /opt/FJSVswsts/bin/swstsvrset -h job1
swstsvrset completed
#
```



Note

This command can only be executed if none of the following commands are running:

- Storage Server configuration information Configuration command (swstsvrset)
- Backup policy Configuration command (swstbkpolset)
- Backup policy deletion command (swstbkpoldel)
- Backup execution command (swstbackup)
- Restore execution command (swstrestore)
- Restore cancellation command (swstcancelrest)
- Backup synchronous processing start command (swststartsync)
- Backup synchronous processing cancel command (swstcancelsync)
- Device information Configuration command (swstdevinfo)
- History information deletion command (swsthistdel)
- Tracking cancel command (swstcanceltrk)
- Resource backup command (swstresback)
- Resource restore command (swstresrst)
- Resource match command (swstsrsemth)

9.2.1.2 Storage Server configuration information display command (swstsvrdisp)

This command displays configuration information for a Storage Server that has been set using the Storage Server configuration information Configuration command (swstsvrset).

a. Format

```
[To be executed on a storage server]
/opt/FJSVswsts/bin/swstsvrdisp

[To be executed on a storage management server]
/opt/FJSVswsts/bin/swstsvrdisp [-h Server]
```

b. Description of options

Option	Description
-h	Specifies a Storage Server name. This option is valid only on a Storage Management Server.

c. End status

```
= 0: Normal end
> 0: Abnormal end
```

d. Example

Display configuration information of a Storage Server (Work1).

```
# /opt/FJSVswsts/bin/swstsvrdisp
Storage-Server = job1
#
```

The following information will be displayed:

Title	Description
Storage-Server	Displays a Storage Server name.

9.2.1.3 Device information Configuration command (swstdevinfofet)

Set a volume that has been allocated to the disk array unit Fujitsu ETERNUS storage system, as a transaction or backup volume.

A transaction volume refers to a volume that stores working data to be backed up.

A backup volume refers to a volume used to store a backup of this transaction data.

a. Format

```
[To be executed on a storage server]
/opt/FJSVswsts/bin/swstdevinfofet -t | -b | -o Device-Name

[To be executed on a storage management server]
/opt/FJSVswsts/bin/swstdevinfofet [-h Server] -t | -b | -o Device-Name
```

b. Description of options

Option	Description
-h	Specifies a Storage Server name. This option is valid only on a Storage 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.

c. Description of operand

Operand	Description
Device-Name	Specify a device name.

d. End status

- = 0: Normal end
- > 0: Abnormal end

e. Example

Register a device (/dev/dsk/clt0d0) as a transaction volume.

```
# /opt/FJSVswsts/bin/swstdevinfofet -t /dev/dsk/clt0d0
swstdevinfofet completed
#
```

Register a device (/dev/dsk/clt0d2) as a backup volume.

```
# /opt/FJSVswsts/bin/swstdevinfofet -b /dev/dsk/clt0d2
swstdevinfofet completed
#
```

Deregister a device (/dev/dsk/clt0d0) that has been registered as a transaction volume.

```
# /opt/FJSVswsts/bin/swstdevinfofet -o /dev/dsk/clt0d0
swstdevinfofet completed
#
```



Note

A device cannot be registered as a transaction volume if:

- Another device with the same logical volume information (box identifier, OLU number, EXTENT start 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 transaction or backup volume.



Note

A device cannot be registered as a backup volume if:

- Another device with the same logical volume information (box identifier, OLU number, EXTENT start 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 transaction or backup volume.
- The device to be set up is already mounted. However, when configuring a volume group as backup volume, the check of whether the logical volume in a volume group is mounted is not carried out. If you set a volume group as a backup volume, unmount the logical volume beforehand.
- The device to be set up is registered as a replication destination volume of 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.



Note

If using a volume group as a transaction volume:

- Do not register two or more physical disks into one volume group.
- Create a logical disk so that it does not consist of two or more physical disks.
- Please note that a volume group that is not active cannot be set as a transaction volume.



Note

If using a volume group as a backup volume:

- Do not register two or more physical disks into one volume group.
- Create a logical disk so that it does not consist of two or more physical disks.
- The backup volume must be the same physical size as the transaction volume. In addition, the backup volume's logical disk must be in a different volume group to the transaction volume, but it must have the same configuration as the transaction volume.
- Please note: a volume group that is not active cannot be set as a backup volume.



Note

You must delete related backup history information before:

- Changing a backup volume to a transaction volume
- Canceling the Configuration of a backup volume



Note

You must delete a backup policy specified on a transaction volume before:

- Changing a transaction volume to a backup volume
- Canceling the Configuration of a transaction volume



Note

In order to make a configuration change to a transaction volume or a backup volume, delete the Configuration, and re-register.



Note

The following changes are not possible unless the related synchronous processing (EC session) is cancelled:

- Changing a Configuration of a suspended transaction volume
- Changing a Configuration of a suspended backup volume



Note

To perform the following changes, the related tracking processing (OPC session) must be canceled in advance:

- Changing the Configuration of a transaction volume in tracking processing
- Changing the Configuration of a backup volume in tracking processing



- This command can only be executed if none of the following commands is running:
- Storage Server configuration information Configuration command (swstsvrset)
 - Device information Configuration command (swstdevinfo)
 - Backup policy Configuration command (swstbkpolset)
 - Backup policy deletion command (swstbkpoldel)
 - Backup execution command (swstbackup)
 - Restore execution command (swstrestore)
 - Restore cancellation command (swstcancelrest)
 - Backup synchronous processing start command (swststartsync)
 - Backup synchronous processing cancel command (swstcancelsync)
 - History information deletion command (swsthistdel)
 - Tracking cancel command (swstcanceltrk)
 - Resource match command (swstsrsemch)
 - Resource backup command (swstresback)
 - Resource restore command (swstresrst)

9.2.1.4 Device usage status display command (swstdevdisp)

This command displays information about a transaction or backup volume or any other definable device, which has been set using the device information Configuration command (swstdevinfo).

a. Format

```
[To be executed on a storage server]
/opt/FJSVswsts/bin/swstdevdisp [-t] [-b [-u]] [-o] [Device-Name]

[To be executed on a storage management server]
/opt/FJSVswsts/bin/swstdevdisp [-h Server] [-t] [-b [-u]] [-o] [Device-Name]
```

b. Description of options

Option	Description
-h	Specifies a Storage Server name. This option can be specified only on a Storage 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.

c. Description of operand

Operand	Description
Device-Name	Specifies a device name. Displays device information on all the transaction volumes if this operand and the associated options are omitted.

- d. End status
 - = 0: Normal end
 - > 0: Abnormal end

- e. Example

Display the usage statuses of the registered devices.

```
# /opt/FJVSwsts/bin/swstdevdisp -t - b
Server Device          Size          Device-Mode    Mount-Point (Method)
Backup-Engine
job2 /dev/dsk/clt0d0    1.0 Gbyte Transaction /mnt/tran1 (vxfs)    AdvancedCopy
job2 /dev/vg01        2.0 Gbyte Transaction ---- (----)          AdvancedCopy
job2 /dev/dsk/clt0d2  1.0 Gbyte Backup (used) ---- (----)          AdvancedCopy
job2 /dev/dsk/clt0d3  1.0 Gbyte Backup (free) ---- (----)          AdvancedCopy
job2 /dev/vg02        2.0 Gbyte Backup (free) ---- (----)          AdvancedCopy
job2 /dev/dsk/clt0d7  1.0 Gbyte Backup (free) ---- (----)          AdvancedCopy
:
#
```

The following information will be displayed:

Title	Description
Server	Displays a Storage Server name.
Device	Displays an AdvancedCopy Manager device name.
Size	Displays the size of a partition allocated to a device. In the case of a volume group, the display also shows the size of the physical disk on which the logical disk is defined.
Device-Mode	Displays text by which a device can be identified: <ul style="list-style-type: none"> • "Transaction": Represents a transaction volume. • "Backup": Represents a backup volume. • "Other": Represents a device that can be registered as a transaction 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 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.
Backup-Engine	Displays the operation type. <ul style="list-style-type: none"> • "AdvancedCopy": Backup using AdvancedCopy Manager



Note

This command cannot display volume sizes to the byte. Refer to the notes under "Configuring the device usage type" for details of how to display and check sizes in bytes.

9.2.1.5 Backup policy Configuration command (swstbkpolset)

This command sets a backup policy (the numbers of preservation generations and interval days). The number of preservation generations means how many generations of backup data should be kept. The number of interval days means the number of days after the execution of a backup that the next backup should be performed.



Note

Even if the number of interval days is specified, AdvancedCopy Manager does not automatically create a backup.



Note

When you specify a backup policy, there must be as many backup volumes in operation as specified in the policy. To view information on the number of backup volumes required to perform the backup, refer to "Preparing a backup volume".

a. Format

```
[To be executed on a storage server]
/opt/FJSVswsts/bin/swstbkpolset [-i Interval-Days] [-s Save-Number] [Device-Name]

[To be executed on a storage management server]
/opt/FJSVswsts/bin/swstbkpolset [-h Server] [-i Interval-Days] [-s
Save-Number] [Device-Name]
```

b. Description of options

Option	Description
-h	Specifies a Storage Server name. This option can be specified only on a Storage Management Server.
-i	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	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.

c. Description of operand

Operand	Description
Device-Name	Specify an AdvancedCopy Manager device name. If the device name is omitted, the backup policy will be set for all the transaction volumes. When a device name is omitted, a backup policy is set up to all transaction volumes.

d. End status

- = 0: Normal end
- > 0: Abnormal end

e. Example

To set 3 as the number of preservation generations for a transaction volume (/dev/dsk/c1t0d0):

```
# /opt/FJSVswsts/bin/swstbkpolset -s 3 /dev/dsk/c1t0d0
/dev/dsk/c1t0d0 swstbkpolset completed
#
```



Note

If this command is executed without a transaction volume name, backup policies are set for all transaction volumes.



Note

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.



Note

No backup policy can be set if:

- There are not as many backup volumes as the specified number of preservation generations.
- The specified number of preservation generations is greater than the maximum allowed.
- The specified number of interval days is greater than the maximum allowed.
- The replication control function or tape backup function uses the target transaction volume
- The specified device is a LU (disk) unit



Note

This command can be executed only while none of the following commands is running:

- Resource match command (swstsrsemth)
- Device information Configuration command (swstdevinfo)
- Storage Server configuration information Configuration command (swstsvrset)
- Resource backup command (swstresback)
- Resource restore command (swstresrst)

9.2.1.6 Backup policy display command (swstbkpoldisp)

This command displays a backup policy that has been set for a transaction volume.

a. Format

```
[To be executed on a storage server]
/opt/FJVSswsts/bin/swstbkpoldisp [Device-Name]

[To be executed on a storage management server]
/opt/FJVSswsts/bin/swstbkpoldisp [-h Server] [Device-Name]
```

b. Description of options

Option	Description
-h	Specifies a Storage Server name. This option can be specified only on a Storage Management Server.

c. Description of operand

Operand	Description
Device-Name	Specify an AdvancedCopy Manager device name. If the device name is omitted, a backup policy for a transaction volume with a registered backup policy will be displayed.

- d. End status
 - = 0: Normal end
 - > 0: Abnormal end

- e. Example

To display the backup policies of a transaction volume (/dev/dsk/clt0d0):

```
# /opt/FJJSVswsts/bin/swstbkpoldisp /dev/dsk/clt0d0
Server Device          Interval-Days Save-Number Mount-Point (Method)
job2 /dev/dsk/clt0d0 30              3             /mnt/tran1 (vxfs)
#
```

The following information will be displayed:

Title	Description
Server	Displays the Storage Server name.
Device	Displays the device name.
Interval-Days	Displays the number of interval days.
Save-Number	Displays the number of preservation generations.
Mount-Point (Method)	Displays the mount point of a device. The file system type of the mount point will be displayed inside the parentheses.

9.2.1.7 Backup policy deletion command (swstbkpoldel)

This command deletes an existing backup policy.

If the transaction volume to be deleted has backup history information, delete the backup history information prior to the execution of this command.

- a. Format

```
[To be executed on a storage server]
/opt/FJJSVswsts/bin/swstbkpoldel Device-Name

[To be executed on a storage management server]
/opt/FJJSVswsts/bin/swstbkpoldel [-h Server] Device-Name
```

- b. Description of options

Option	Description
-h	Specifies a Storage Server name. This option can be specified only on a Storage Management Server.

- c. Description of operand

Operand	Description
Device-Name	Specifies an AdvancedCopy Manager device name.

- d. End status
 - = 0: Normal end
 - > 0: Abnormal end

e. Example

To delete the backup policies of a transaction volume (/dev/dsk/clt0d0):

```
# /opt/FJSVswsts/bin/swstbkpodel /dev/dsk/clt0d0
/dev/dsk/clt0d0 swstbkpodel completed
#
```



Note

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 control function or tape backup function uses the target transaction volume



Note

This command can be executed only while none of the following commands are running:

- Storage Server configuration information Configuration command (swstsvrset)
- Device information Configuration command (swstdevinfo)
- Resource match command (swstsrsemch)
- Resource backup command (swstresback)
- Resource restore command (swstresrst)

9.2.2 Operation commands

This section describes operation commands for backup management.

9.2.2.1 Backup execution command (swstbackup)

When using this command, the operation changes according to the 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 information, such as the date and backup volume of a copy target, is registered into backup history information. This is called snapshot type high-speed backup.
- When synchronous processing (EC) is performed
The state of synchronous processing is checked and, in the case of an equivalency 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. This is called synchronous high-speed backup. Note that it aborts when an equivalent maintenance state is not reached.
To use the Suspend/Resume function for synchronous high-speed backup, execute the backup command with "-suspend" specified, and suspend backup synchronous processing.

In both cases, a backup is established, and access to the backup volume is available immediately after the command execution.

AdvancedCopy Manager executes a pre-processing Shell script (OpcBackup.pre) before backup and then a post-processing Shell script (OpcBackup.post) after backup. For information on customizing these Shell scripts, refer to "Pre-processing and Post-processing of Backup and Restoration".

a. Format

```
[To be executed on a storage server]
/opt/FJVSswsts/bin/swstbackup [-suspend|-T] Device-Name [-Xdevmap Device-Map-File]

[To be executed on a storage management server]
/opt/FJVSswsts/bin/swstbackup [-h Server] [-suspend|-T] Device-Name [-Xdevmap
Device-Map-File]
```

b. Description of options

Option	Description
-h	Specifies a Storage Server name. This option can be specified only on a Storage Management Server.
-Xdevmap	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 is 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 Storage Server) using an absolute pathname. A file name specified in Device-Map-File must not include any national characters.
-suspend	This results in synchronous high-speed backup using the 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 cabinet when the disk array supports the QuickOPC function. If this option is not specified, ordinary snapshot processing (OPC without using the QuickOPC function) is started. This option must always be specified to perform a backup using differential snapshots.

c. Description of operand

Operand	Description
Device-Name	Specifies a device name.

d. End status

```
= 0: Normal end
> 0: Abnormal end
```

e. Example

To back up a business volume:

```
# /opt/FJVSswsts/bin/swstbackup /dev/dsk/c1t0d0
/dev/dsk/c1t0d0 swstbackup completed
#
```



Note

To guarantee the data, this command prevents other processes from accessing the transaction and backup volumes before commencing the backup. Thus, the drive letters of devices are not assigned. This command terminates with an error if a drive letter of a transaction volume is assigned or if:

- The mount point has directories to which other volumes are to be mounted.
- The volume has a file in use.



Note

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 the backup command.
3. Execute the fsck command for the backup volume, and check the file system.



Note

When configuring a volume group as a transaction volume, and backing up or restoring it, modify the pre-processing and post-processing scripts for backup such that unmount/mount of all logical devices is carried out.



Note

You cannot perform backup in the following cases. In these cases, perform the action described in the [System Administrator's Response] message that will be output.

- No backup policy has been specified for the specified transaction volume.
- The backup that was executed one generation earlier was unsuccessful.
- Restore is currently being executed for the transaction volume that is targeted for processing.
- Synchronous backup processing is being performed but the transaction volume has not yet reached equivalency maintenance status.
- The mount point set at registration of a transaction volume has been changed.
- Restoration is being executed using backup history information that will cause a generation overflow.
- No unused backup volume could be acquired to use as the copying destination.
- Any of the cabinet information (box identifier, OLU number, EXTENT start position, and EXTENT size) of a transaction volume has been changed.
- A backup volume defined as the copy destination is mounted.
- Changes have been made to the logical disk configuration which cannot be supported by AdvancedCopy Manager.
- A volume group is inactive in the volume group specification.
- The replication control function or tape backup function uses the target transaction volume
- The replication control function or tape backup function uses the target backup volume
- The replication control function or tape backup function uses backup records exceeding the number of preservation generations



Note

You cannot perform backup if:

- The backup volume specified by the device map file (in the "to" column) cannot be used in a pair with the transaction volume. This occurs if:
 1. The specified backup volume is in use.
 2. The partition size of the transaction volume and the backup volume are different.
- The device map file contains an error (refer to "Preparing a device map file").
- The device map file name includes a national character.



Point

At this time, if existing backup history information satisfies the condition of the number of saved generations specified by the backup policy Configuration command (swstbkpolset), 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 is specified for the backup history information, then a backup volume with the Suspend or tracking state is not cleared and remains in use.



Note

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.



Note

For notes on executing backup, refer to "General notes".



Note

This command can be executed only while none of the following commands are running:

- Storage Server configuration information Configuration command (swstsvrset)
- Device information Configuration command (swstdevinfoset)
- Resource match command (swstsrsemetch)
- Resource backup command (swstresback)
- Resource restore command (swstresrst)

9.2.2.2 Backup execution status display command (swstbackstat)

This command displays the completion status of a backup that is being performed using the backup execution command (swstbackup).



Note

Do not use the OPC stop function of ETERNUSmgr/GRmgr while executing the snapshot high speed backup. If the ETERNUSmgr/GRmgr stop function is used, it will display "succeeded", regardless of whether or not the backup was actually successful.

a. Format

```
[To be executed on a storage server]
/opt/FJSVswsts/bin/swstbackstat [Device-Name]

[To be executed on a storage management server]
/opt/FJSVswsts/bin/swstbackstat [-h Server] [Device-Name]
```

b. Description of options

Option	Description
-h	Specifies a Storage Server name. This option can be specified only on a Storage Management Server.

c. Description of operand

Operand	Description
Device-Name	Specifies a device name corresponding to a transaction volume. If this operand is omitted, the backup execution statuses of all the transaction volumes will be displayed.

d. End status

- = 0: Normal end
- > 0: Abnormal end

e. Example

Display the actual copying execution status.

```
# /opt/FJSVswsts/bin/swstbackstat
Server Transaction-Disk Backup-Disk      Status      Execute
job1 /dev/dsk/clt0d0 /dev/dsk/clt0d3 succeeded ----
job1 /dev/vg01      ----          ----          ----
:
#
```

The following information will be displayed:

Keyword	Description
Server	Displays a Storage Server name.
Transaction-Disk	Displays the 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. <ul style="list-style-type: none"> • "----": 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. • "halt": Copying using OPC has been halted.
Execute	Displays how much has been copied in percentage terms if the status is "executing", otherwise displays "----".



Note

You cannot display the backup processing status if:

- Changes have been made to the logical disk configuration that cannot be supported by AdvancedCopy Manager.
- A volume group is inactive in the volume group specification.



Note

If "failed" or "halt" is displayed in the Status field, a hardware error may have occurred. For more information refer to "Troubleshooting for a hardware error, etc. occurring during backup".

9.2.2.3 Restore execution command (swstrestore)

The data that exists in backup history information is restored using OPC.



Note

Restoration of each individual file cannot be performed because copying between devices will be performed.

When the restore execution command is executed, AdvancedCopy Manager executes the restore pre-processing Shell script (OpcRestore.pre), copies data using OPC, then executes the restore post-processing Shell script (OpcRestore.post). For information on customizing these scripts, refer to "Pre-processing and Post-processing for Backup and Restoration".



Point

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 must have the same amount of space available as the specified transaction volume.

a. Format

```
[To be executed on a storage server]
/opt/FJSVswsts/bin/swstrestore [-g Generation-No | -v Version-No] [-r
Restore-Device-Name] Device-Name

[To be executed on a storage management server]
/opt/FJSVswsts/bin/swstrestore [-h Server] [-g Generation-No | -v Version-No]
[-r Restore-Device-Name] Device-Name
```

b. Description of options

Option	Description
-h	Specifies a Storage Server name. This option can be specified only on a Storage Management Server.
-g	Specifies the relative generation number of data to be restored. Check the relative generation number using the history information display command (swsthistdisp).
-v	Specifies the absolute generation number of data to be restored. Check the absolute generation number using the history information display command (swsthistdisp).
-r	Specifies a restore destination device name to restore data to any device other than a transaction volume.

Note: If neither the -g or -v option is specified, the latest history management information will be restored.

c. Description of operand

Operand	Description
Device-Name	Specifies a device name corresponding to a transaction volume.

d. End status

- = 0: Normal end
- > 0: Abnormal end

e. Example

Restore data with relative generation number 2 (For information on the relative generation number, refer to the data displayed using the history information display command (swsthistdisp)).

```
# /opt/FJSVswsts/bin/swstrestore -g 2 /dev/dsk/clt0d0
/dev/dsk/clt0d0 swstrestore completed
#
```

If the Suspend/Resume function is used for the backup 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 the backup synchronous processing execution status display command (swstsyncstat). Execute the backup synchronous processing cancel command (swstcancelsync) to cancel the synchronous processing.

```
# /opt/FJSVswsts/bin/swstsyncstat /dev/dsk/clt0d0s6
Server Transaction-Disk Backup-Disk      Status      Execute
job2 /dev/dsk/clt0d0 /dev/dsk/clt0d3 suspend      ----
job2 /dev/dsk/clt0d0 /dev/dsk/clt0d2 executing 75%
# /opt/FJSVswsts/bin/swstcancelsync -all /dev/dsk/clt0d0
/dev/dsk/clt0d0 swstcancelsync completed
# /opt/FJSVswsts/bin/swstsyncstat /dev/dsk/clt0d0
Server Transaction-Disk Backup-Disk Status Execute
job2 /dev/dsk/clt0d0 ----      ----      ----
# /opt/FJSVswsts/bin/swstrestore /dev/dsk/clt0d0
/dev/dsk/clt0d0 swstrestore completed
#
```



Point

On a transaction volume on which a file system has been constructed physically, restore a specific file as follows:

1. Mount a backup volume. Check the backup volume by executing the Backup History List view/History information display command (swsthistdisp).
2. Copy the file to be restored by executing, for example, the cp command.
3. Unmount the backup volume.



Point

In restoring a specific file in the transaction volume that is having the file system of a volume group built, the following procedures are performed:

1. Mount a backup volume. Check the backup volume by executing the Backup History List view/History information display command (swsthistdisp).
2. Copy the file to be restored, for example by executing the cp command.
3. Unmount the backup volume.



Note

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.



Note

You cannot perform restoration in the following cases. In these cases, perform the action described in the [System Administrator's Response] message that will be output.

- Both relative and absolute generation numbers have been specified.
- A device registered as a transaction volume has been specified in a parameter with the "-r" option.
- 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. You are recommended never to specify a backup volume as a restore destination device.
- You perform restoration immediately after starting backup (ie, while 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). In this case, perform restoration 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 above (when restoring a specific file at a transaction volume on which a file system has been constructed).
- A transaction volume performing synchronous backup processing is specified.
- Any of the cabinet information (box identifier, OLU number, EXTENT start position, and EXTENT size) of a transaction volume to be restored has been changed.
- The mount point that was set during the transaction volume registration has been changed.
- Changes have been made to the logical disk configuration that cannot be supported by AdvancedCopy Manager.
- A volume group is inactive in the volume group specification.
- The replication control function or tape backup function uses the target transaction volume
- The replication control function or tape backup function uses the device specified in the - r option



Note

The suspended or tracking backup volume cannot be used as the restore destination volume.



Note

The restore command cannot be executed on a transaction volume during synchronous processing (during copying by EC, in the equivalency maintenance state, or in the Suspend state).

If either of the following errors occur during restore processing, execute the backup synchronous processing cancel command (swstcancelsync) (with the -all option specified) to cancel all of the synchronous processing of the transaction volume, then execute the restore command.

swst0639 This Command cannot be run. This is because a copy is currently running.

or

swst0634 OPC cannot be run. Error code=26 (SCSI command busy)



Note

Refer to "General notes", for notes on executing restore.



Note

This command can only be executed while none of the following commands are running:

- Storage Server configuration information Configuration command (swstsvrset)
- Device information Configuration command (swstdevinfofet)
- Resource match command (swstsrsemctch)
- Resource backup command (swstresback)
- Resource restore command (swstresrst)

9.2.2.4 Restore cancellation command (swstcancelrest)

This command cancels OPC-based restoration. If a hardware error is detected while data is being copied using OPC, remove the cause of the hardware error and perform restoration again using the restore execution command (swstrestore).

Check whether an abnormality has occurred using the restore execution status display command (swstreststat).



Note

If you cancel OPC-based restoration using this command, the transaction volume will be incomplete and can no longer be used because part of the data has already been copied. Perform restoration again using the restore execution command (swstrestore).

a. Format

```
[To be executed on a storage server]
/opt/FJSVswsts/bin/swstcancelrest [-g Generation-No | -v Version-No] [-r
Restore-Device-Name] [-emergency] Device-Name

[To be executed on a storage management server]
/opt/FJSVswsts/bin/swstcancelrest [-h Server] [-g Generation-No | -v Version-No] [-r
Restore-Device-Name] [-emergency] Device-Name
```

b. Description of options

Option	Description
-h	Specifies a Storage Server name. This option can be specified only on a Storage Management Server.
-g	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 the restore execution status display command (swstreststat).
-v	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 the restore execution status display command (swstreststat).
-r	Specifies the device name if data is restored to any device other than a transaction volume. Check the device name using the restore execution status display command (swstreststat).
-emergency	Emergency operation mode is available for operations. In such cases, OPC sessions are neither confirmed nor cancelled.

Note: The restoration of the latest history information will be canceled if neither the -g or -v option is specified.

c. Description of operand

Operand	Description
Device-Name	Specifies a device name.

- d. End status
= 0: Normal end
> 0: Abnormal end

e. Example

Cancel restore processing of a transaction volume (/dev/dsk/clt0d0).

```
# /opt/FJJSVswsts/bin/swstcancelrest -g 1 /dev/dsk/clt0d0
/dev/dsk/clt0d0 swstcancelrest completed.
#
```

Restoration of the latest history information on the transaction volume (/dev/dsk/clt0d0) is cancelled in emergency operation mode.

```
# /opt/FJJSVswsts/bin/swstcancelrest - emergency /dev/dsk/clt0d0
/dev/dsk/clt0d0 swstcancelrest completed.
#
```



Note

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.)
- Changes have been made to the logical disk configuration that cannot be supported by AdvancedCopy Manager.
- A volume group is inactive in the volume group specification.
- The replication control function or tape backup function uses the target transaction volume



Note

OPC sessions are neither confirmed nor cancelled by execution in emergency operation mode. If an OPC session is in the error suspended state or halt state, then unmanaged sessions remain. In this event, GRmgr or ETERNUSmgr must be used to cancel these sessions.



Note

This command can be executed only while none of the following commands are running:

- Storage Server configuration information Configuration command (swstsvrset)
- Device information Configuration command (swstdevinfoset)
- Resource match command (swstsrsemch)
- Resource backup command (swstresback)
- Resource restore command (swstresrst)

9.2.2.5 Restore execution status display command (swstreststat)

This command displays the execution status of a restoration.



Note

Do not use the OPC stop function of ETERNUSmgr/GRmgr while the snapshot fast backup is in progress. If the ETERNUSmgr/GRmgr stop function is used, it will display "succeeded", regardless of whether or not the restoration was actually successful.

a. Format

```
[To be executed on a storage server]
/opt/FJSVswsts/bin/swstreststat [-g Generation-No | -v Version-No] [Device-Name]

[To be executed on a storage management server]
/opt/FJSVswsts/bin/swstreststat [-h Server] [-g Generation-No | -v Version-No]
[Device-Name]
```

b. Description of options

Option	Description
-h	Specifies a Storage Server name. This option can be specified only on a Storage Management Server.
-g	Specifies the relative generation number of data whose restore execution status is to be displayed.
-v	Specifies the absolute generation number of data whose restore execution status is to be displayed.

Note: The restore execution statuses of all the histories will be displayed if both the -g and -v options are omitted.

c. Description of operand

Operand	Description
Device-Name	Specifies a device name. If this operand is omitted, the restore execution statuses of all the transaction volumes will be displayed.

d. End status

= 0: Normal end
> 0: Abnormal end

e. Example

Display the restore execution status of transaction volumes.

```
# /opt/FJSVswsts/bin/swstreststat
Server Transaction-Disk  Generation  Version  Backup-Disk      Restore-Device
Status  Execute
job2  /dev/dsk/c1t0d0  1          10      /dev/dsk/c1t0d8  /dev/dsk/c1t0d11
executing 75%
job2  /dev/dsk/c1t0d2  ----      ----      ----      ----      ----
----
job2  /dev/dsk/c1t0d5  2          12      /dev/dsk/c1t0d9  ----
succeeded ---
job2  /dev/dsk/c1t0d6  3          13      /dev/dsk/c1t0d10 /dev/dsk/c1t0d12
executing 75%
:
#
```

The following information will be displayed.

Title	Description
Server	Displays a Storage Server name.
Transaction-Disk	Displays the 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 a restore is not in progress.
Restore-Device	Displays a restore target device name unless data is copied and restored to a transaction volume. "----" will be displayed if a restore is not in progress or data is restored to a transaction volume.
Status	Displays an execution status. <ul style="list-style-type: none"> • "----": No copying using OPC is in progress. • "executing": Copying using OPC is in progress. • "succeeded": Copying using OPC has been completed. • "failed": Copying using OPC has been interrupted due to an error. • "halt": Copying using OPC has been halted.
Execute	If the status is "executing" then displays how much has been copied in percentage terms, otherwise displays "----".



If "failed" or "halt" is displayed in the "Status" field, a hardware error may have occurred. For more information refer to "Troubleshooting for a hardware error, etc. occurring during backup".



If no generation number has been specified in the -v or -g option, the restore execution statuses of all the history will be displayed. If, at this time, restoration is not in progress from any of the history, "----" will be displayed in all the information fields other than "Server" and "Transaction-Disk".



- You cannot display the restore processing status in the following cases:
- Changes have been made to the logical disk configuration which cannot be supported by AdvancedCopy Manager.
 - A volume group is in an inactive state in volume group specification.

9.2.2.6 History information display command (swsthistdisp)

This command displays history information that has been backed up.

a. Format

```
[To be executed on a storage server]
/opt/FJSVswsts/bin/swsthistdisp [Device-Name]

[To be executed on a storage management server]
/opt/FJSVswsts/bin/swsthistdisp [-h Server] [Device-Name]
```

b. Description of options

Option	Description
-h	Specifies a Storage Server name. This option can be specified only on a Storage Management Server.

c. Description of operand

Operand	Description
Device-Name	Specifies a device name corresponding to a transaction volume. If this operand is omitted, the backup history information of all the transaction volumes will be displayed.

d. End status

= 0: Normal end
> 0: Abnormal end

e. Example

Display the backup history information of all the transaction volumes.

```
# /opt/FJSVswsts/bin/swsthistdisp
Server=StrgSV01 Device=/dev/dsk/c1t0d0 Mount-Point=/mnt/tran1 (vxfs)
Generation Version Backup-Date Backup-Device Status Execute
1 10 2000/11/12 22:00 /dev/dsk/c1t0d2 succeeded ----
2 9 2000/11/11 22:00 /dev/dsk/c1t0d4 succeeded ----
:
#
```

The following information will be displayed:

Keyword	Description
Server	Displays a Storage Server name.
Device	Displays a device name.
Mount-Point	Displays the mount point name of the device and the file system type 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 date at which backup has been completed. This item will not be displayed if no backup history information exists.
Backup-Device	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.

Keyword	Description
Status	<p>Displays an actual copying execution status using one of the following character strings:</p> <ul style="list-style-type: none"> • "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. • "halt": Copying using OPC has been halt status. <p>This item will not be displayed if no backup history information exists.</p>
Execute	<p>If the status is "executing" then displays how much has been copied in percentage terms, otherwise displays "----".</p> <p>This item will not be displayed if no backup history information exists.</p>



Note

- You cannot perform processing of displaying history information in the following cases:
- Changes have been made to the logical disk configuration which cannot be supported by AdvancedCopy Manager.
 - A volume group is in an inactive state in volume group specification.



Note

If "failed" or "halt" is displayed in the "Status" field, a hardware error may have occurred. For more information refer to "Troubleshooting for a hardware error, etc. occurring during backup".

9.2.2.7 History information deletion command (swsthistdel)

This command deletes backup history information.

The backup volume of the deleted history is released and it turns into empty backup volume.



Note

If the history information display command (swsthistdisp) displays "failed" in the "Status" field, a hardware error may have occurred. For more information refer to "Troubleshooting for a hardware error, etc. occurring during backup".

a. Format

```
[To be executed on a storage server]
/opt/FJSVswsts/bin/swsthistdel -g Generation-No | -v Version-No | -z [-emergency]
Device-Name

[To be executed on a storage management server]
/opt/FJSVswsts/bin/swsthistdel [-h Server] -g Generation-No | -v Version-No | -z
[-emergency] Device-Name
```

b. Description of options

Option	Description
-h	<p>Specifies a Storage Server name.</p> <p>This option can be specified only on a Storage Management Server.</p>
-g	<p>Specifies the relative generation number of backup history information to be deleted.</p> <p>This option cannot be specified at the same time as the -v or -z.</p> <p>Check the relative generation number using the history information display command (swsthistdisp).</p>

Option	Description
-v	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. Check the absolute generation number using the history information display command (swsthistdisp).
-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.
-emergency	An emergency operation mode is available. In such cases, OPC sessions are neither confirmed nor cancelled.

c. Description of operand

Operand	Description
Device-Name	Specifies a device name corresponding to a transaction volume.

d. End status

= 0: Normal end
> 0: Abnormal end

e. Example

Delete the backup history of absolute generation number 10 in a transaction volume (/dev/dsk/clt0d0).

```
# /opt/FJSVswsts/bin/swsthistdel -v 10 /dev/dsk/clt0d0
/dev/dsk/clt0d0 swsthistdel completed
#
```

The history information of relative generation number 1 of the transaction volume (/dev/dsk/clt0d0) is deleted in emergency operation mode.

```
# /opt/FJSVswsts/bin/swsthistdel -g 1 -emergency /dev/dsk/clt0d0
/dev/dsk/clt0d0 swsthistdel completed
#
```



Note

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 being used in the restore process.
- Changes have been made to the logical disk configuration which cannot be supported by AdvancedCopy Manager.
- A volume group is inactive state in the volume group specification.
- The replication control function or tape backup function uses the target transaction volume
- The replication control function or tape backup function uses the target backup volume



Note

OPC sessions are neither confirmed nor cancelled by execution in emergency operation mode. If an OPC session is in the error suspended state or halt state, then unmanaged sessions remain. In this event, ETERNUSmgr or GRmgr must be used to cancel these sessions.



Point

Deleting the suspended history information does not cancel the synchronous processing (EC session).

Deleting the tracking history information does not cancel the tracking processing (OPC session).



Point

In synchronous high-speed backup with "-suspend" specified, the backup volume remains in use even after the history information is deleted.

In differential snapshot high-speed backup with - T specified, the backup volume remains in use even after the history information is deleted.



Note

If the history information is deleted while the OPC physical copy has not been completed in a backup 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 of the 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 the 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.



Note

This command can be executed only while none of the following commands are running:

- Storage Server configuration information Configuration command (swstsvrset)
- Device information Configuration command (swstdevinfo)
- Resource match command (swstsrsem)
- Resource backup command (swstresback)
- Resource restore command (swstresrst)

9.2.2.8 Backup synchronous processing start command (swststartsync)

This command starts backup synchronous processing (starts disk-to-disk copy using EC).

a. Format

```
[To be executed on a storage server]
/opt/FJSVswsts/bin/swststartsync Device-Name [-Xdevmap Device-Map-File]

[To be executed on a storage management server]
/opt/FJSVswsts/bin/swststartsync [-h Server] Device-Name [-Xdevmap Device-Map-File]
```

b. Description of options

Option	Description
-h	Specify the name of a Storage Server. This option can be specified only on the Storage Management Server.
-Xdevmap	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, the 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 Storage Server) using an absolute pathname. A file name specified in Device-Map-File must not include any national characters.

c. Description of operand

Operand	Description
Device-Name	Specifies a device name.

d. End status

- = 0: Normally terminated
- > 0: Abnormally terminated

e. Example

Start synchronous backup processing of a transaction volume (/dev/dsk/c1t0d0).

```
# /opt/FJSVswsts/bin/swststartsync /dev/dsk/c1t0d0
/dev/dsk/c1t0d0 swststartsync completed.
#
```



Note

When you execute the backup synchronous processing start command for a transaction volume to which the backup synchronous processing is already being executed, the message, "swst0301 Backup synchronous processing is already in progress." is output and then the command is terminated normally.



Note

You cannot start backup synchronous processing in the following cases. In any of the cases below, perform the action specified in the "System administrator response" message.

- When specifying a device that does not support the EC function.
- When the backup volumes required to start 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.
- Changes have been made to the logical disk configuration which cannot be supported by AdvancedCopy Manager.
- A volume group is inactive in the volume group specification
- When the replication control function or tape backup function uses the target transaction volume



Note

If the backup volume is a volume group, do not execute the `vgchange (1M)` command to deactivate the volume group under synchronous processing (EC/REC) execution. If this occurs and the backup volume is inactive then it cannot be reactivated. Additionally, it will not be possible to perform commands including those to stop synchronous processing, such as the backup execution command (`swstbackup`), and the backup synchronous processing cancel command (`swstcancelsync`). If it has been inactivated accidentally, make the backup volume active after stopping the synchronous processing using `ETERNUSmgr/GRmgr`, etc.



Point

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.



Point

Even if already in use, the transaction volume where the synchronous processing started and the suspended backup volume can be used as backup destinations.



Note

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.



Note

This command cannot be executed if any of the following commands are running:

- Resource match command (`swstsrsemch`)
- Storage Server configuration information Configuration command (`swstsvrset`)
- Device information Configuration command (`swstdevinfo`)
- Resource backup command (`swstresback`)
- Resource restore command (`swstresrst`)

9.2.2.9 Backup synchronous processing cancel command (`swstcancelsync`)

The backup synchronous processing can be canceled during copying by EC or under the equivalency maintenance status.

The synchronous processing can be canceled with the backup volume during copying by EC, under the equivalency maintenance 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.

When the system detects an error while backup synchronous processing is being executed, the copying being performed at that time will terminate abnormally. If this occurs, use this command to cancel the backup synchronous processing. Use the backup synchronous processing progress display command (`swstsyncstat`) to check for the error.

a. Format

```
[To be executed on a storage server]
/opt/FJSVswsts/bin/swstcancelsync [-emergency] [-bd Backup-Device-Name | -all]
Device-Name

[To be executed on a storage management server]
/opt/FJSVswsts/bin/swstcancelsync [-h Server] [-emergency] [-bd Backup-Device-Name |
-all] Device-Name
```

b. Description of options

Option	Description
-h	Specify the name of a Storage Server. This option can be specified only with a Storage Management Server.
-emergency	Emergency operation mode is available for operations. In such cases, EC sessions are neither confirmed nor cancelled.
-bd	The synchronous processing on a specific backup volume can be cancelled. To do so, specify the backup volume name in an operand. Note that the backup volume name and the -all option cannot be used together.
-all	All of the synchronous processing of the transaction volume can be cancelled. Note that this option and -bd cannot be used together.

c. Description of operand

Operand	Description
Device-Name	Specifies a device name.
Backup-Device-Name	Specifies the device name for the copy destination backup volume .

d. End status

```
= 0: Normally terminated
> 0: Abnormally terminated
```

e. Example

Cancel synchronous backup processing of the transaction volume (/dev/dsk/clt0d0).

```
# /opt/FJSVswsts/bin/swstcancelsync /dev/dsk/clt0d0
/dev/dsk/clt0d0 swstcancelsync completed.
#
```

Backup synchronous processing of the transaction volume (/dev/dsk/clt0d0) is cancelled in emergency operation mode.

```
# /opt/FJSVswsts/bin/swstcancelsync -emergency /dev/dsk/clt0d0
/dev/dsk/clt0d0 swstcancelsync completed.
#
```

Cancel the backup synchronous processing of the backup volume (/dev/dsk/clt0d2) from the transaction volume (/dev/dsk/clt0d0).

```
# /opt/FJVSswsts/bin/swstsyncstat /dev/dsk/clt0d0
Server Transaction-Disk Backup-Disk      Status   Execute
job2 /dev/dsk/clt0d0 /dev/dsk/clt0d3 suspend ----
job2 /dev/dsk/clt0d0 /dev/dsk/clt0d2 executing 75%
# /opt/FJVSswsts/bin/swstcancelsync /dev/dsk/clt0d0 - bd /dev/dsk/clt0d2
/dev/dsk/clt0d0 swstcancelsync completed.
# /opt/FJVSswsts/bin/swstsyncstat /dev/dsk/clt0d0
Server Transaction-Disk Backup-Disk      Status   Execute
job2 /dev/dsk/clt0d0 /dev/dsk/clt0d3 suspend ----
#
```

Cancel all the backup synchronous processing of the transaction volume (/dev/dsk/clt0d0).

```
# /opt/FJVSswsts/bin/swstsyncstat /dev/dsk/clt0d0
Server Transaction-Disk Backup-Disk      Status   Execute
job2 /dev/dsk/clt0d0 /dev/dsk/clt0d3 suspend ----
job2 /dev/dsk/clt0d0 /dev/dsk/clt0d2 executing 75%
# /opt/FJVSswsts/bin/swstcancelsync /dev/dsk/clt0d0 - all
/dev/dsk/clt0d0 swstcancelsync completed.
# /opt/FJVSswsts/bin/swstsyncstat /dev/dsk/clt0d0
Server Transaction-Disk Backup-Disk Status Execute
job2 /dev/dsk/clt0d0 ----      ----      ----
#
```



Note

You cannot cancel synchronous backup processing in the following cases:

- Changes have been made to the logical disk configuration which cannot be supported by AdvancedCopy Manager.
- When a volume group is inactive in the volume group specification.
- When the replication control function or tape backup function uses the target transaction volume



Note

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 halted state, then unmanaged sessions remain. In this event, ETERNUSmgr or GRMgr must be used to cancel these sessions.



Note

This command cannot be executed when any of the following commands is running:

- Resource match command (swstrsrsemrch)
- Device information Configuration command (swstdevinfoet)
- Storage Server configuration information Configuration command (swstsvrset)
- Resource backup command (swstresback)
- Resource restore command (swstresrst)

9.2.2.10 Backup synchronous processing progress display command (swstsyncstat)

The execution status of the backup synchronous processing is displayed (during copying, under the equivalency maintenance state or in the suspended state).

If one transaction volume has multiple EC sessions, then the statuses of all of these EC sessions are displayed.

a. Format

```
[To be executed on a storage server]
/opt/FJSVswsts/bin/swstsyncstat [Device-Name]

[To be executed on a storage management server]
/opt/FJSVswsts/bin/swstsyncstat [-h Server] [Device-Name]
```

b. Description of options

Option	Description
-h	Specifies the name of a Storage Server. This option can be specified only with the Storage Management Server.

c. Description of operand

Operand	Description
Device-Name	Specifies a device name. If you omit the device name, the system displays the status of backup synchronous processing for all of the transaction volumes.

d. End status

```
= 0: Normally terminated
> 0: Abnormally terminated
```

e. Example

Display the synchronous backup processing status of a transaction volume (/dev/dsk/c1t0d0).

```
# /opt/FJSVswsts/bin/swstsyncstat /dev/dsk/c1t0d0
Server Transaction-Disk Backup-Disk      Status   Execute
job1 /dev/dsk/c1t0d0 /dev/dsk/c1t0d3 suspend    ----
job1 /dev/dsk/c1t0d0 /dev/dsk/c1t0d2 executing 75%
#
```

Display the synchronous backup processing status of all business volumes.

```
# /opt/FJSVswsts/bin/swstsyncstat
Server Transaction-Disk Backup-Disk      Status   Execute
job1 /dev/dsk/c1t0d0 /dev/dsk/c1t0d1 suspend    ----
job1 /dev/dsk/c1t0d0 /dev/dsk/c1t0d2 executing 75%
job1 /dev/dsk/c1t0d0 /dev/dsk/c1t0d4 suspend    ----
job1 /dev/dsk/c1t0d1 /dev/dsk/c1t0d3 equivalent 100%
job1 /dev/dsk/c1t0d5 /dev/dsk/c1t0d7 failed     ----
job1 /dev/dsk/c1t0d6 ----          ----          ----
#
```

The table below lists the meaning of each title.

Title	Description
Server	The name of the Storage Server.
Transaction-Disk	The name of the transaction volume.
Backup-Disk	The name of the backup volume. When backup synchronous processing is not being executed, the system displays "----".

Title	Description
Status	The current backup synchronous processing status: <ul style="list-style-type: none"> • "----": Backup synchronous processing is not being executed. • "executing": Backup synchronous processing is being executed but equivalency maintenance status has not been established. • "equivalent": Equivalency maintenance 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). • "halt": Backup synchronous processing has the halt status.
Execute	The progress of the backup synchronous processing as a percentage (0 to 100%). "----" is displayed when the status is Suspend and when the backup synchronous processing is not executing.



Note

- You cannot display the status of the synchronous processing in the following cases:
- Changes have been made to the logical disk configuration which cannot be supported by AdvancedCopy Manager.
 - A volume group is inactive in the volume group specification.



Note

If the status is "failed" or "halt", a hardware failure may have occurred. Refer to "Troubleshooting for a hardware error, etc. occurring during backup" for more information.



Note

If the status is "nosession", there may be resource information inconsistency. Use the resource match command (swstsrsemctch) to remove the inconsistency from the resource information.

9.2.2.11 Execution status display command (swstexecstat)

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.

a. Format

```
[To be executed on a storage server]
/opt/FJSVswsts/bin/swstexecstat [-j] [Device-Name]

[To be executed on a storage management server]
/opt/FJSVswsts/bin/swstexecstat [-h Server] [-j] [Device-Name]
```

b. Description of options

Option	Description
-h	Specifies a Storage Server name. This option can be specified only on a Storage Management Server.
-j	Specifies to display information in another format.

c. Description of operand

Operand	Description
Device-Name	Specifies a device name corresponding to a transaction volume. If this operand is omitted, the execution statuses of all the transaction volumes will be displayed.

d. End status

- = 0: Normal end
- > 0: Abnormal end

e. Example

Display the execution statuses of transaction volumes.

```
# /opt/FJSVswsts/bin/swstexecstat
Server Device          Last-Backup-Date Interval Status Mount-Point (Method) Execute
job1 /dev/dsk/c1t0d0 2000/12/11 12:20 OK      IDLE /mnt/trans1 (vxfs)  ----
job1 /dev/dsk/c1t0d5 2000/12/10 12:20 DELAY=1 IDLE /mnt/usr1 (hfs)     sync (22%)
job1 /dev/dsk/c1t0d6 2000/12/09 12:20 DELAY=2 IDLE /mnt/usr2 (ufs)
snapshot (45%)
:
#
```

The following information will be displayed:

Keyword	Description
Server	Displays a Storage Server name.
Device	Displays an AdvancedCopy Manager device name.
Last-Backup-Date	Displays the date and time when the last backup execution command was accepted.
Interval	If more than the specified number of days have passed since the last backup then the number of days overdue will be displayed as "DELAY=number-of-days-passing". "OK" will be displayed if the number of interval days has not been exceeded.
Status	Indicates whether any of the following commands is being executed on a transaction volume. If so, "EXEC" will be displayed. Otherwise, "IDLE" will be displayed. <ul style="list-style-type: none"> • <u>Backup policy setting command (swstbkpolset)</u> • <u>Backup policy deletion command (swstbkpoldel)</u> • <u>Backup execution command (swstbackup)</u> • <u>Restore execution command (swstrestore)</u> • <u>History information deletion command (swsthistdel)</u> • <u>Backup synchronous processing start command (swststartsync)</u> • <u>Backup synchronous processing cancel command (swstcancelsync)</u> • <u>Restore cancellation command (swstcancelrest)</u>
Mount-Point (Method)	Displays the mount point name of a transaction volume, with the file system type in parentheses.

Keyword	Description
Execute	<p>Displays the name of a command being executed or the copying progress percentage.</p> <ul style="list-style-type: none"> • "----": Neither a command or OPC/EC is being executed. • "swststartsync": Backup synchronous processing start command is being executed. • "swstcancelsync": Backup synchronous processing cancel command is being executed. • "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 configuration 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%)": How far the EC operation has progressed. • "sync(failed)": EC has been abnormally terminated. • "sync(halt)": EC has been halted. • "snapshot(xxx%)": Progress ratio of backup using OPC • "snapshot(failed)": Backup using OPC has abnormally terminated. • "snapshot(halt)": Backup using OPC has been has been halted. • "restore(xxx%)": Progress ratio of restoration using OPC • "restore(failed)": Restoration using OPC has abnormally terminated. • "restore(halt)": Restoration using OPC has been halted.

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 snap shot 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 execution 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 cannot be performed during disk copy of the snapshot backup or backup synchronous processing due to a restriction of the Fujitsu ETERNUS storage system.)

The "-j" option can be used as in the following example:

```
# /opt/FJVSswsts/bin/swstexecstat - j
Server Device          Execute-Date      Status Mount-Point (Method) Execute
job1 /dev/dsk/c1t0d0 2000/12/11 12:20 IDLE /usr1 (ufs)      ----
job1 /dev/dsk/c1t0d2 2000/12/10 12:20 IDLE /usr2 (ufs)      sync(22%)
job1 /dev/dsk/c1t0d5 2000/12/09 12:20 IDLE /usr3 (ufs)      snapshot(45%)
:
#
```

The following information will be displayed:

Keyword	Description
Server	A Storage Server name.
Device	An AdvancedCopy Manager device name.
Execute-Date	The command execution time.

Keyword	Description
Status	<p>Indicates whether any of the following commands are being executed on a transaction volume. If so, "EXEC" will be displayed. Otherwise, "IDLE" will be displayed.</p> <ul style="list-style-type: none"> • Backup policy setting command (swstbkpolset) • Backup policy deletion command (swstbkpoldel) • Backup execution command (swstbackup) • Restore execution command (swstrestore) • History information deletion command (swsthistdel) • Backup synchronous processing start command (swststartsync) • Backup synchronous processing cancel command (swstcancelsync) • Restore cancellation command (swstcancelrest)
Mount-Point(Method)	<p>Displays the mount point name of a transaction volume, with the file system type in parentheses.</p>
Execute	<p>Displays the name of a command being executed or the copying progress percentage.</p> <ul style="list-style-type: none"> • "----": Neither a command or OPC/EC is being executed. • "swststartsync": Backup synchronous processing start command is being executed. • "swstcancelsync": Backup synchronous processing cancel command is being executed. • "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 configuration 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 the EC operation. • "sync(failed)": EC has been abnormally terminated. • "sync(halt)": EC has been halted. • "snapshot(xxx%)": Progress ratio of backup using OPC • "snapshot(failed)": Backup using OPC has abnormally terminated. • "snapshot(halt)": Backup using OPC has been halted." • restore(xxx%): Progress ratio of restoration using OPC • "restore(failed)": Restoration using OPC has abnormally terminated. • "restore(halt)": Restoration using OPC has been halted.



Point

During synchronous high-speed backup with "-suspend" specified, no operation ("----") is indicated in the Execute field.



Note

You cannot display the execution status in the following cases:

- Changes have been made to the logical disk configuration which cannot be supported by AdvancedCopy Manager.
- A volume group is inactive in the volume group specification.



Note

If "sync (failed)", "snapshot(failed)", or "restore(failed)", "sync (halt)", "snapshot(halt)" or "restore(halt)" is displayed in the "Execute" field, a hardware error may have occurred. After fixing the cause of the hardware error:

- For "sync (failed)"
Execute the backup synchronous processing cancel command (swstcancelsync) to cancel the backup synchronous processing.
- If "snapshot(failed)" is displayed
Delete backup history information using the history information deletion command (swsthistdel).
- If "restore(failed)" is displayed
Cancel restoration using the restore cancellation command (swstcancelrest).

9.2.2.12 Tracking cancel command (swstcanceltrk)

This command stops tracking processing.

a. Format

```
[To execute on the storage server]
/opt/FJSVswsts/bin/swstcanceltrk [-copy | -emergency] Device-Name (-bd
Backup-Device-Name | -all)
[To execute on the storage management server]
/opt/FJSVswsts/bin/swstcanceltrk -h Server-Name [-copy | -emergency] Device-Name (-bd
Backup-Device-Name | -all)
```

b. Options

Option	Description
-h	Specifies a Storage Server name. This option can only be specified on the Storage Management Server.
-copy	Specifies that physical copying, if active, should also be stopped (ie, the OPC session should be canceled). 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 the history information deletion command (swsthistdel).
-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	Stops tracking processing for the specified backup volume. Specify the target backup volume name as an operand. This option cannot be used with the - all option.
-all	Stops the entire tracking processing specified for the transaction volumes. This option cannot be used with the - bd option.

c. Operand

Operand	Description
Device-Name	Specifies a device name.

d. End status

- = 0: Normal end
- > 0: Abnormal end

e. Example

The following example stops the tracking that was specified for the backup volume (/dev/dsk/clt0d8) from the transaction volume (/dev/dsk/clt0d3).

```
# /opt/FJSVswsts/bin/swstcanceltrk /dev/clt0d3 -bd /dev/clt0d8
/dev/clt0d3 swstcanceltrk completed.
#
```

The following example stops the tracking and physical copying that were specified for the transaction volume (/dev/dsk/clt0d3).

```
# /opt/FJSVswsts/bin/swsttrkstat /dev/clt0d3
Server Transaction-Disk Backup-Disk Status Execute Update
job2 /dev/clt0d3 /dev/clt0d3 executing 33% -----
# /opt/FJSVswsts/bin/swstcanceltrk - copy /dev/clt0d3 -bd /dev/clt0d8
/dev/clt0d3 swstcanceltrk completed.
#
```



Note

If the command is executed in emergency operation mode, no OPC session is canceled. A session that is not managed remains if the command is entered during tracking. The session must be canceled using ETERNUSmgr, GRmgr, or GDS.



Note

Tracking processing cannot be stopped if tracking processing is not performed for the specified copy source and destination volumes. In this case, take action according to "System administrator response" provided for the message output.



Note

This command can only be executed while if none of the following commands are being processed:

- Storage Server configuration information Configuration command (swstsvrset)
- Device information Configuration command (swstdevinfofet)
- Resource matching command (swstsrsemctch)
- Resource backup command (swstresback)
- Resource restore command (swstresrst)

9.2.2.13 Tracking status display command (swsttrkstat)

This command displays the status of tracking processing.

When multiple types of tracking processing are running on one transaction volume, the command displays the statuses of all types of tracking processing.

a. Format

```
[To execute on the storage server]
/opt/FJSVswsts/bin/swsttrkstat [Device-Name]
[To execute on the storage management server]
/opt/FJSVswsts/bin/swsttrkstat [-h Server] [Device-Name]
```

b. Options

Option	Description
-h	Specifies a Storage Server name. This option can be specified only on the Storage Management Server.

c. Operands

Operand	Description
Device-Name	Specifies a device name. If this operand is omitted, the command displays the statuses of tracking processing for all transaction volumes.

d. End status

- = 0: Normal end
- > 0: Abnormal end

e. Example

The following example displays the status of tracking processing of a transaction volume (/dev/dsk/clt0d3):

```
# /opt/FJSVswsts/bin/swsttrkstat /dev/clt0d3
Server Transaction-Disk Backup-Disk Status Execute Update
job1 /dev/clt0d3 /dev/clt0d8 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/clt0d3 /dev/clt0d8 nosession ---- ----
job1 /dev/clt0d4 /dev/clt0d9 executing 75% ----
job1 /dev/clt0d5 /dev/clt0d10 failed ---- ----
job1 /dev/clt0d6 /dev/clt0d11 tracking ---- 12%
job1 /dev/clt0d7 ---- ---- ---- ----
:
#
```

The displayed data is explained below.

Title	Description
Server	A Storage Server name.
Transaction-Disk	A transaction volume name.
Backup-Disk	A backup volume name. Displays "----" when the volume is not in the tracking status.
Status	The execution status. "----": No tracking is in progress "executing": Physical copying and tracking is in progress "tracking": Tracking is in progress "failed": Physical copying or tracking terminated abnormally. "nosession": No session exists (ie, the resource information is inconsistent).
Execute	The copy progress as a percentage when the status is "executing". "----" is displayed when the status is "tracking".
Update	The amount of updated data after logical copying as a percentage when the status is "tracking". "----" is displayed when the status is "executing".



Note

When the execution status (displayed in the "Status" column) is "failed," a hardware problem has occurred. After fixing the problem, re-execute the differential snapshot backup.



Note

When the execution status (displayed in the "Status" column) is "nosession," the resource information is probably inconsistent. Use the resource matching command (swstrsemth) to remove any inconsistencies.



Note

The execution status of tracking processing cannot be displayed under the following conditions:

- - When the transaction volume or backup volume is a logical volume, and the volume configuration of the logical volume has been changed to one that is not supported by AdvancedCopy Manager.
- - When the transaction volume or backup volume is a volume group, and the volume group is not active.

9.2.3 Maintenance commands

This section describes maintenance commands for backup management.

9.2.3.1 Resource match command (swstrsemth)

This command recovers resource information if a system failure, etc. has caused inconsistencies.

a. Format

```
[To be executed on a storage server]

/opt/FJSVswsts/bin/swstrsemth [-x]

[To be executed on a storage management server]

/opt/FJSVswsts/bin/swstrsemth [-x] [-h Server]
```

b. Description of operands

Operand	Description
-h	Specifies a Storage Server name. This option can be specified only on a Storage Management Server.
-x	Achieves consistency with a repository. Specify this option if a backup management file and a repository become inconsistent.

c. Example

```
# /opt/FJSVswsts/bin/swstrsemth
swstrsemth completed
#
```



Note

This command can only be executed if none of the following commands are running:

- Storage Server configuration information Configuration command (swstsvrset)
- Device information Configuration command (swstdevinfo)
- Backup policy Configuration command (swstbkpolset)
- Backup policy deletion command (swstbkpoldel)
- Backup execution command (swstbackup)
- Restore execution command (swstrestore)
- Restore cancellation command (swstcancelrest)
- Backup synchronous processing start command (swststartsync)
- Backup synchronous processing cancel command (swstcancelsync)
- History information deletion command (swsthistdel)
- Tracking cancel command (swstcanceltrk)
- Resource match command (swstsrsemch)
- Resource backup command (swstresback)
- Resource restore command (swstresrst)

9.2.3.2 Resource backup command (swstresback)

This command copies a backup management file to a specified directory.

a. Format

```
/opt/FJSVswsts/bin/swstresback Backup-Directory
```

b. Description of operand

Operand	Description
Backup-Directory	Specifies the directory to which a backup management file should be saved.

c. Example

Create a backup copy of the backup management files in /home/backup.

```
# /opt/FJSVswsts/bin/swstresback /home/backup
swstresback completed
#
```



Note

You cannot perform backup if:

- The specified directory does not exist.
- Copying of data to the specified directory fails, for example because of insufficient free space



Note

This command can only be executed while none of the following commands are running:

- Storage Server configuration information Configuration command (swstsvrset)
- Device information Configuration command (swstdevinfo)
- Backup policy Configuration command (swstbkpolset)
- Backup policy deletion command (swstbkpodel)
- Backup execution command (swstbackup)
- Restore execution command (swstrestore)
- Restore cancellation command (swstcancelrest)
- Backup synchronous processing start command (swststartsync)
- Backup synchronous processing cancel command (swstcancelsync)
- History information deletion command (swsthistdel)
- Tracking cancel command (swstcanceltrk)
- Resource match command (swstsrsemch)
- Resource backup command (swstresback)
- Resource restore command (swstresrst)

9.2.3.3 Resource restore command (swstresrst)

The backup management files are restored from the specified directory.

a. Format

```
/opt/FJSVswsts/bin/swstresrst Backup-Directory
```

b. Description of operand

Operand	Description
Backup-Directory	Specifies the directory in which to save the backup management files.

c. Example

Restore the backup copy of the backup management files from /home/backup

```
# /opt/FJSVswsts/bin/swstresrst /home/backup
swstresrst completed
#
```



Note

Before the backup management files can be restored by this command, the Storage Server configuration information Configuration command (swstsvrset) must be executed.



Note

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, for example because of insufficient free space.



This command can be executed only while none of the following commands are running:

- Storage Server configuration information Configuration command (swstsvrset)
- Device information Configuration command (swstdevinfo)
- Backup policy Configuration command (swstbkpolset)
- Backup policy deletion command (swstbkpoldel)
- Backup execution command (swstbackup)
- Restore execution command (swstrestore)
- Restore cancellation command (swstcancelrest)
- Backup synchronous processing start command (swststartsync)
- Backup synchronous processing cancel command (swstcancelsync)
- History information deletion command (swsthistdel)
- Tracking cancel command (swstcanceltrk)
- Resource match command (swstsrsemch)
- Resource backup command (swstresback)
- Resource restore command (swstresrst)

9.3 Configuration Management Commands

This section explains the configuration management commands. These commands manage the information on a server device which performs backup and replication, or the node for operation in a cluster environment.



Note Operating environment for command execution

The table below lists the operating environment requirements for each command:

Table: Operating environment required for executing commands

Command	Operating environment of AdvancedCopy Manager
Stgcmmmodnode	To execute this command on a Storage Management Server (Storage Management Server transaction), the Storage Management Server (Storage Management Server transaction) must be operating.
Stgenvdisp	This command is executable when both Storage Management Server (storage management server transaction) and Storage Server (storage server transaction) are not operating. Note that when the cluster service (which is named differently in respective cluster systems) is not in cluster operation, this command cannot display certain information.
Stgmgrinfoset	The package_backup_of_the_management_information execution server must be operating.
stgmgrinfobkup	The Storage Management Server (Storage Management Server transaction), package_backup_of_the_management_information execution server, and package_backup_of_the_management_information target servers must be operating.
stgmgrinfodisp	The package_backup_of_the_management_information execution server must be operating.



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. For "bsh"shell, 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 specifies the nodes on which commands are executable. For notes and cautions on command execution, see the Remarks column.

Table: Cluster operation commands

Command	Environment variable	Command execution node	Remarks
stgcmmmodnode	Required	Cluster active server	
stgenvdisp	Not required	Cluster active server	
stgmgrinfobkup	Required	Cluster active server	

Command	Environment variable	Command execution node	Remarks
stgmgrinfodisp	Required	Cluster active server	
stgmgrinfofet	Required	Cluster active server	

9.3.1 Server information change command (stgcmmodnode)

This command can be used to change the following information on a Storage Management Server:

- To change the IP address of the Storage Management Server or the PORT number (stgxfws) used by the communication daemon, follow the procedure below:
 1. Change the server information by executing this command on the Storage Management Server.
 2. Change the server information of the Storage Management Server by executing this command on every Storage Server. This command must be executed on all Storage Servers managed by the Storage Management Server.
- To change the server name of the Storage Management Server, follow the procedure below:
 1. Change the server name from the GUI client of the Storage Management Server.
 2. Change the server name of the Storage Management Server by executing this command on every Storage Server managed by the Storage Management Server.

This command has the following two functions:

- Displays Storage Management Server information specified on the Storage Server on which this command is currently executed
- Changes Storage Management Server information.

a. Format

```
/opt/FJSVswstf/bin/stgcmmodnode [-i IP-address] [-p PORT number] [-n server name]
```

b. Description of options

Option	Description
None	Displays the server information currently managed.
-i	Specifies the IP address of a Storage Management Server to change the IP address of the Storage Management Server.
-p	Specifies the port number of a Storage Management Server to change the port number of the Storage Management Server.
-n	Specifies the server name of a Storage Management Server to change the server name of the Storage Management Server. This option cannot be specified if this command is executed on a Storage Management Server.

c. End status

- = 0: Normal end
- > 0: Abnormal end

d. Example

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 Storage Management Server to 10.12.13.14.

```
# /opt/FJSVswstf/bin/stgcmmodnode -i 10.12.13.14
swstf2401 Storage Management Server registry updated.
```



Note

Additional information about this command:

- This command can only be executed by root users.
- The PORT number of a Storage Management Server can be between 1024 and 65535.
- This command must specify the same value as the one specified on the Storage Management Server. An error will occur if an invalid value is specified.
- If this command is executed, restart the AdvancedCopy Manager daemons.

9.3.2 Repository access user change command (stguserset)

The repository access user change command (stguserset) is executed by the Storage Management Server.

For more information refer to the "ETERNUS SF AdvancedCopy Manager Operator's Guide" relevant to the OS of the Storage Management Server.

9.3.3 GUI client connection information Configuration command (stgguiipset)

The GUI CLIENT connection information Configuration command (stgguiipset) is executed by the Storage Management Server.

For more information refer to the "ETERNUS SF AdvancedCopy Manager Operator's Guide" relevant to the OS of the Storage Management Server.

9.3.4 Server information addition command (stgxfwcmaddsrv)

The server information addition command (stgxfwcmaddsrv) is executed by the Storage Management Server.

For more information refer to the "ETERNUS SF AdvancedCopy Manager Operator's Guide" relevant to the OS of the Storage Management Server.

9.3.5 Server information change command (stgxfwcmmodsrv)

The server information change command (stgxfwcmmodsrv) is executed by the Storage Management Server.

For more information refer to the "ETERNUS SF AdvancedCopy Manager Operator's Guide" relevant to the OS of the Storage Management Server.

9.3.6 Server information deletion command (stgxfwcmdelsrv)

The server information deletion command (stgxfwcmdelsrv) is executed by the Storage Management Server.

For more information refer to the "ETERNUS SF AdvancedCopy Manager Operator's Guide" relevant to the OS of the Storage Management Server.

9.3.7 Device information collection/reflection command (stgxfwcmsetdev)

The Device information collection/reflection command (stgxfwcmsetdev) is executed by the Storage Management Server.

For more information refer to the "ETERNUS SF AdvancedCopy Manager Operator's Guide" relevant to the OS of the Storage Management Server.

9.3.8 Device information deletion command (stgxfwcmdeldev)

The device information deletion command (stgxfwcmdeldev) is executed by the Storage Management Server.

For more information refer to the "ETERNUS SF AdvancedCopy Manager Operator's Guide" relevant to the OS of the Storage Management Server.

9.3.9 Server information display command (stgxfwcmdisprv)

The server information display command (stgxfwcmdisprv) is executed by the Storage Management Server.

For more information refer to the "ETERNUS SF AdvancedCopy Manager Operator's Guide" relevant to the OS of the Storage Management Server.

9.3.10 Device information display command (stgxfwcmdispdev)

The device information display command (stgxfwcmdispdev) is executed by the Storage Management Server.

For more information refer to the "ETERNUS SF AdvancedCopy Manager Operator's Guide" relevant to the OS of the Storage Management Server.

9.3.11 Partition information display command (stgxfwcmdisppat)

The partition information display command (stgxfwcmdisppat) is executed by the Storage Management Server.

For more information refer to the "ETERNUS SF AdvancedCopy Manager Operator's Guide" relevant to the OS of the Storage Management Server.

9.3.12 Environment information display command (stgenvdisp)

This command displays the environment configuration in AdvancedCopy Manager.



Note

- 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 cluster environment, the environment variable (SWSTGNODE) need not be set.
- Only users with root authority can execute this command.

a. Format

```
/opt/FJSVswstf/bin/stgenvdisp [-n transaction-logical-node-name]
```

b. Description of options

Option	Description
-n	Displays the logical node name (as specified in the cluster setup command) of the transaction. <ul style="list-style-type: none">• In a cluster environment If this option is specified, only information on the selected transactions is displayed. Otherwise, information on all transactions is displayed. <ul style="list-style-type: none">• In a non-cluster environment This option cannot be specified, otherwise message swstf6001 is displayed and processing ends.

- c. End status
 - = 0: Normal end
 - > 0: Abnormal end

d. Example

Information in a non-cluster environment is displayed.

```
# /opt/FJSVswstf/bin/stgenvdisp
Server type      : Server
Manager IP address : 10.10.10.10
Communication    : 1226
-----
Version          : 13.3
Language         : JAPANESE
Program install directory : /opt
Fixed configuration directory : /etc/opt
Modifying configuration directory : /var/opt
#
```

Information on all transactions in a cluster environment (MC/ServiceGuard) is displayed.

```
# /opt/FJSVswstf/bin/stgenvdisp
Server type      : Server
Cluster system   : MC/ServiceGuard
Node type        : Primary
Node name        : nodeagt1
Service name     : agent_service1
Mount point      : /acmagt1disk
Volume group     : vg01
Subnet mask      : 255.255.255.0
IP by cluster setup : 10.10.10.10
Manager IP address : 10.10.10.10
Communication    : 11000
-----
Server type      : Server
Cluster system   : MC/ServiceGuard
Node type        : Secondary
Node name        : nodeagt2
Service name     : agent_service2
Mount point      : /acmagt2disk
Volume group     : vg02
Subnet mask      : 255.255.255.0
IP by cluster setup : 10.10.10.10
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          : 13.3
Language         : JAPANESE
Program install directory : /opt
Fixed configuration directory : /etc/opt
Modifying configuration directory : /var/opt
#
```

Information on all transactions in a cluster environment (VERITAS Cluster Server) is displayed.

```
# /opt/FJVSwtstf/bin/stgenvdisp
Server type          : Server
Cluster system      : VERITAS Cluster Server
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      : VERITAS Cluster Server
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             : 13.3
Language            : JAPANESE
Program install directory : /opt
Fixed configuration directory : /etc/opt
Modifying configuration directory : /var/opt
#
```

The displayed information is described below:

Title	Description
Server type	The server type is displayed as follows: Management server: Storage Management Server or Storage Management Server transaction Server: Storage Server of Storage Server package
Cluster system	The types of cluster system are as follows: SynfinityCluster or PRIMECLUSTER: SynfinityCluster or PRIMECLUSTER VERITAS Cluster Server: VERITAS Cluster Server ynfinityCluster: SynfinityCluster PRIMECLUSTER: PRIMECLUSTER MC/ServiceGuard: HP MC/ServiceGuard MSCS: Microsoft Cluster Server or WSFC HACMP: IBM High Availability Cluster Multi-Processing
Node type	The node type is displayed as follows: 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.
Volume group	The volume group name is displayed.

Title	Description
Network name	The name of the network 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 Configuration command is displayed. However, note that the displayed IP is not the active logical IP address if the IP address of either the Storage Management Server transaction or Storage Server transaction is changed during operation.
Manager IP address	The IP address of the Storage 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 Edition only] The name of the DB log file storage destination folder is displayed.
DB dictionary file	[Windows Edition only] The name of the DB dictionary file storage destination folder is displayed.
Repository database	[Windows Edition only] The name of the repository DB storage destination folder is displayed.
DB file	[Solaris Edition] The name of the DB file system device or directory is displayed. [Linux Edition] The DB file directory name is displayed.
RDB dictionary	[Solaris Edition] The RDB dictionary raw device name is displayed. [Linux Edition] The RDB dictionary directory name is displayed.
RDB log file	[Solaris Edition] The name of the RDB log file raw device or directory is displayed. [Linux OE Edition] The RDB log file directory name is displayed.
Database space	[Solaris Edition] The name of the raw device containing the repository data storage DB is displayed. [Linux Edition] The name of the directory containing the repository data storage DB is displayed.
Database size	[Windows Edition and Linux Edition only] The database size is displayed.
IP address for GUI	The GUI CLIENT connection IP address is displayed.
Communication	The communication service/daemon port number (the service name in the services file: stgxfws) is displayed. In a cluster environment, the transaction communication service/daemon port number (the service name in the services file: stgxfws_<logical node name>) is displayed.
Remote access	The remote access service/daemon port number (the service name in the services file: rmiswstf) is displayed.
Exclusive control	The exclusive control service/daemon port number (the service name in the services file: semswstf) is displayed.
Authentication feature	The authentication feature service/daemon port number (the service name in the services file: smaclmgr) is displayed.
Repository	The repository service/daemon port number (service name in the services file: rdaswstf) is displayed.
Resource Type	[For Solaris and Linux] Displays the resource type Procedure: Procedure resource CmdLine: Command line resource
Network Resource	[Solaris only] Displays the name of the logical host name resource or the name of the

Title	Description
	shared address resource
Program install directory	Displays the program directory specified during installation
Fixed configuration directory	Displays the environment configuration directory specified during installation
Modifying configuration directory	Displays the transactions directory specified during installation
Version	Displays the version and level of AdvancedCopy Manager
Language	Displays the AdvancedCopy Manager language type JAPANESE: Domestic edition ENGLISH: International edition
Code	[Solaris and Linux only] Displays the encoding specified during installation



Note

- "???" is displayed whenever information cannot be displayed because of an error.
- Some items may not be output, depending on the cluster type.

9.3.13 Management information package backup server registration/deletion command (stgmgrinfo set)

This command can be used to register, delete or display information about the management information package backup execution server (the server on which the management information package backup command (stgmgrinfo backup) is executed).

a. Format

```
program-directory%bin%stgmgrinfo set [-del | -d directory-name -c
storage-management-server-encoding]
```

b. Description of options

Option	Description
None	Displays information about the management information package backup execution server.
-del	Deletes the management information package backup execution server.
-d	Specify the full 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	Specify the encoding designated when the Storage Management Server was installed. Specify ASCII.

c. End status

- =0: Normal end
- >0: Abnormal end

d. Example

Register the management information package backup execution server (assuming that the directory where backup data is to be stored is D:/¥stgbkupdata, and the encoding used on the Storage Management Server is ASCII):

```
# /opt/FJSVswstf/bin/stgmgrinfo set -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/stgmgrinfo set -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/stgmgrinfo set
Directory      : /stgbkupdata
Manager code   : ASCII
#
```



Note

Additional information about this command:

- Only a root or superuser or an account which belongs to the Administrators group can execute this command.
- 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 Storage Server as the management information package backup execution server.
- Set the name of the logical node in the SWSTGNODE environment variable when the Storage Server registered as the management information package backup execution server is in a cluster system.
- Create the directory specified in the “-d” option in advance.

9.3.14 Management information package backup command (stgmgrinfobkup)

This command makes backups of the management information stored on the Storage Servers and Storage Management Server.

a. Format

```
program-directory¥bin¥stgmgrinfobkup [-n server-name] [-t backup-type]
```

b. Description of options

Option	Description
-n	Creates backups of the management information stored on the specified server. If this option is not specified, the management information on all Storage Servers (including the Storage Management Server) is backed up.
-t	Specifies the types of data to be backed up. The available options are 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 AuthData : Authentication management list BkupData : Backup management ReplData : Replication management TapeData : Tape backup management backup If this option is not specified, backups are created as defined in the policy file. For details on the policy file, refer to "Preliminaries".

c. End status

- =0: Normal end
- >0: Abnormal end

d. Example

Backup of the management information on all the Storage Servers, 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.
#
```

Backup of the management information on Storage Server acm01, 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.
#
```

Backup of the registries on all of the Storage Servers.

```
# /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.
#
```

Backup of backup management located on Storage Server acm01:

```
# /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.
#
```



Note

Additional information about this command:

- Only a root or superuser or an account which belongs to the Administrators group can execute this command.
- Be sure to check the information (ie, the message indicating completion or an error) that is output by this command.
- This command can only be executed on the management information package backup execution server.
- Set the name of the logical node in the SWSTGNODE environment variable when the management information package backup execution server is in a cluster system.
- Do not execute other AdvancedCopy Manager commands or GUI CLIENT commands while this command is running. Conversely, do not execute this command while another AdvancedCopy Manager command or a GUI CLIENT command is running.
- Do not execute this command again if it is already running.
- 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 Storage Servers with Softek AdvancedCopy Manager 10.3 (V10.0L40) or earlier.

9.3.15 Management information package backup status display command (stgmgrinfodisp)

This command displays the execution status of the management information package backup command (stgmgrinfobkup).

a. Format

```
program-directory\bin\stgmgrinfodisp [-n server-name]
```

b. Description of options

Option	Description
-n	Specifies the name of the server for which the backup status is to be displayed. If this option is not specified, the status of backups on all Storage Servers managed by the Storage Management Server are displayed.

c. End status

- =0: Normal end
- >0: Abnormal end

d. Output information

The types of information output are listed below. 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	The server name
RegiData	The 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) yy/mm/dd hh:mm:ss
DBData	The 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
AuthData	The date and time at which the management information package backup command created a backup of the authentication management list (ie, the date and time that the management information package backup command was executed) yy/mm/dd hh:mm:ss
BkupData	The date and time at which the management information package backup command created a backup of the backup management (date and time that the management information package backup command was executed) yy/mm/dd hh:mm:ss
ReplData	The 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	The 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
TapeData	The date when a backup of tape backup management was made by package backup using the management information command (ie, the date and time that the management information package backup command was executed) yy/mm/dd hh:mm:ss

e. Example

Output information about all Storage Servers:

```
# /opt/FJSVswstf/bin/stgmgrinfodisp
Server:      acm01
  DBData:    04/08/02 16:11:26
  AuthData:  04/08/02 16:32:00
  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
Server:      tape
  RegiData:  05/08/01 09:00:00
  TapeData:  05/08/03 16:00:00
#
```

Output information about Storage Server acm01:

```
# /opt/FJSVswstf/bin/stgmgrinfodisp -n acm01
Server:      acm01
  DBData:    04/08/02 16:11:26
  AuthData:  04/08/02 16:32:00
  RegiData:  04/08/01 09:00:00
  ReplData:  04/08/02 16:32:22
#
```



Note

Additional information about this command:

- Only a root or superuser or an account which belongs to the Administrators group can execute this command.
- This command can only be executed on the management information package backup execution server.
- When the management information package backup execution server is in a cluster system, set the name of the logical node in the SWSTGNODE environment variable.
-

9.4 Replication Management Commands

This section describes operations using Replication Management Commands.



Note

Only a root or superuser can execute all the commands of replication management. If commands are executed 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 environments in which to execute the respective commands.

Table: Operating environment required for executing commands

Command	Operating environment of AdvancedCopy Manager
swsrpcancel	To execute this command on a Storage Management Server with the <code>-h</code> option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating. When the Storage Server runs on Windows and is in EC operation, the daemons and services of the copy destination Storage Server must be operating (except when the <code>-m</code> option or <code>-t</code> option is specified). For inter-server replication in EC equivalency maintenance status, the associated server must be operating (except when the <code>-m</code> option is specified).
swsrpchsnc	To execute this command on a Storage Management Server with the <code>-h</code> option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating.
swsrpdelvol	To execute this command on a Storage Server, the daemons and services of the Storage Management Server must be operating. To execute this command on a Storage Management Server with the <code>-h</code> option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating. For inter-server replication, the daemons and services of the communicating Storage Servers must be operating (except when the <code>-d</code> option is specified).
swsrprebuffstat	To execute this command on a Storage Management Server with the <code>-h</code> option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating.
swsrprebuffset	To execute this command on a Storage Server, the daemons and services of the Storage Management Server must be operating. To execute this command on a Storage Management Server with the <code>-h</code> option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating.
swsrpmake	To execute this command on a Storage Management Server with the <code>-h</code> option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating. When the Storage Server runs on Windows and is engaged in intra-server replication and EC operation, the daemons and services of the Storage Server must be operating. For inter-server replication, the daemons and services of the communicating Storage Servers must be operating (except when the <code>-m</code> option or <code>-t</code> option is specified).
swsrprecoverres	To execute this command on a Storage Management Server with the <code>-h</code> option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating. When the Storage Server runs on Windows, the daemons and services of the Storage Server must be operating. To specify the <code>-r</code> option, the daemons and services of the Storage Management Server must be operating.

Command	Operating environment of AdvancedCopy Manager
swsrprevsync	To execute this command on a Storage Management Server with the -h option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating.
swsrpsetvol	To execute this command on a Storage Server, the daemons and services of the Storage Management Server must be operating. To execute this command on a Storage Management Server with the -h option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating. For inter-server replication, the daemons and services of the communicating Storage Servers must be operating.
swsrpstartsync	To execute this command on a Storage Management Server with the -h option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating. When the Storage Server runs on Windows, the daemons and services of the Storage Server must be operating. For inter-server replication, the associated server must be operating (except a case where the -m option is specified).
swsrpstat	To execute this command on a Storage Management Server with the -h option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating.
swsrpvoinfo	To execute this command on a Storage Management Server with the -h option specified, the daemons and services of the Storage Management Server and target Storage Server must be operating.



Note In cluster operation

In cluster operation, execute a command by taking the following steps:

- 3) If necessary, specify the logical node name to the transaction in environment variable SWSTGNODE. For "bsh"shell, specify the environment variable as follows:

```
# SWSTGNODE=logical-node-name
# export SWSTGNODE
```

- 4) Execute the command.

The table below lists the commands available in cluster operation. It indicates whether environment variables must be configured, and specifies the nodes on which commands are executable. For notes and cautions on command execution, see the Remarks column.

Table: Cluster operation commands

Command	Environment variable	Command execution node	Remarks
swsrpsetvol	Required	Active cluster server	
swsrpvoinfo	Required	Active cluster server	
swsrpdelvol	Required	Active cluster server	
swsrprebuffstat	Required	Active cluster server	
swsrprebuffset	Required	Active cluster server	
swsrpstartsync	Required	Active cluster server	

Command	Environment variable	Command execution node	Remarks
swsrpmake	Required	Active cluster server	
swsrpstat	Required	Active cluster server	
swsrpcancel	Required	Active cluster server	
swsrpchsync	Required	Active cluster server	
swsrprevsync	Required	Active cluster server	
swsrprecoverres	Required	Active cluster server	

9.4.1 Environment definition commands

This section describes environment definition commands for replication management.

9.4.1.1 Replication volume information Configuration command (swsrpsetvol)

This command sets the source and destination volumes that will be used in a replication.

a. Format

<p>[To be executed on a Storage Server] For replication on a server /opt/FJSVswsrp/bin/swsrpsetvol [-n] [-u] [-Xgroup <i>Group-Name</i>] <i>Original-Volume-Name</i> <i>Replica-Volume-Name</i> For replication between servers /opt/FJSVswsrp/bin/swsrpsetvol [-n] [-u] -o (ORG REP BOTH) [-Xgroup <i>Group-Name</i>] <i>Original-Volume-Name</i> <i>Replica-Volume-Name</i> [To be executed on a Storage Management Server] For replication on a server /opt/FJSVswsrp/bin/swsrpsetvol -h <i>Server-Name</i> [-n] [-u] [-Xgroup <i>Group-Name</i>] <i>Original-Volume-Name</i> <i>Replica-Volume-Name</i> For replication between servers /opt/FJSVswsrp/bin/swsrpsetvol -h <i>Server-Name</i> [-n] [-u] -o (ORG REP BOTH) [-Xgroup <i>Group-Name</i>] <i>Original-Volume-Name</i> <i>Replica-Volume-Name</i></p>
--

b. Description of options

Option	Description
-h	Specifies the name of a Storage Server subject to replication. This option is valid only on a Storage Management Server. If this option is omitted, the Storage Server on which the command is entered will be the target of the command.
-u	Specifies the direction of copying. If this option is specified, copying can only be performed from a 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.

Option	Description
-o	<p>Specifies the operation server.</p> <p>Specify one of the following for this operand:</p> <p>ORG: Sets the replication source server as the operation server.</p> <p>REP: Sets the replication destination server as the operation server.</p> <p>BOTH: Sets both the copy-source server and copy-destination server as operation servers.</p> <p>On the server that is not specified as the operation server, only the following commands can be executed for a replication volume:</p> <ul style="list-style-type: none"> • Replication volume information display command (swsrpvoinfo) • Operation status display command (swsrpstat) • Replication cancellation command (swsrpcancel) (The forcible stop option is required.) <p>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 11.0/V11.0L10 or later.</p> <p>This option is valid only for replication between servers.</p>
-n	<p>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.</p> <p>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).</p> <p>When the sizes of the source and destination volumes are different, copying is performed as follows:</p> <p>The address at which copying starts is the first sector of the replication destination volume.</p> <p>The copy data size is the size of the smaller volume.</p> <p>Therefore, make the copy destination volume size the same or larger than the copy source volume size.</p>
-Xgroup	<p>Specifies that the copy source and copy destination volumes be registered in a group.</p> <p>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.</p> <p>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.</p>

c. Description of operands

Operand	Description
Original-Volume-Name	Specifies a replication source volume name. For the volume names of other servers, use the format "volume-name@storage-server-name."
Replica-Volume-Name	Specifies a replication destination volume name. For the volume names of other servers, use the format "volume-name@storage-server-name."

d. End status

- = 0: Normal end
- > 0: Abnormal end

e. Example

Set replication volume information by specifying volume /dev/dsk/c1t2d0 of Storage Server SRC-SV as the replication source volume, and volume /dev/dsk/c1t2d0 of Storage 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 following example shows the command executed on SRC-SV:

```
# /opt/FJSVswsrp/bin/swsrpsetvol -u -o ORG /dev/dsk/c1t2d0 /dev/dsk/c1t2d0@TARG-SV
swsrpsetvol completed
#
```

When the same configuration is made from TARG-SV, the example is as follows (be careful with the volume name):

```
# /opt/FJSVswsrp/bin/swsrpsetvol -u -o ORG /dev/dsk/c1t2d0@SRC-SV /dev/dsk/c1t2d0
swsrpsetvol completed
#
```



Note

Replication volume information cannot be set if:

- Replication volume information has already been set.
- The specified replication source volume is already set as the replication destination volume, or the replication destination volume is already set as the replication source volume.
- 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 cabinet information (ie, box identifier, OLU number, EXTENT start position, and EXTENT size) as the cabinet 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 bidirectional 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).
- The replication source volume and the replication destination volume are located on another disk array device and the remote copy function of the disk array device is not correctly set.
- Communication with the Storage Management Server fails.
- Communication with a destination server fails for replication between servers.
- Changes have been made to the logical disk configuration that cannot be supported by AdvancedCopy Manager.
- A volume group is inactive in the volume group specification.



Note

When using a volume group as a source volume or replica volume, the following applies:

- Be sure to specify the pairs of a volume group as a source volume and replica volume. Note that this does not work between platforms
- The same volume group cannot be registered as a source volume and replica volume.
- When registering a volume group as a source volume and replica volume, make the logical disk configuration the same.



Note

When you change physical information on volumes (ie, partition size, drive letter, and cabinet information on a disk array device) using such functions as an OS management tool (for example, Disk Administrator) or GRmgr/ETERNUSmgr provided by the disk array device, you need to reset the registered replication volume information associated with the changed volumes, by following these steps:

1. Delete the replication volume information using the replication volume information deletion command (swsrpdelvol).
2. In the GUI client, "information acquisition / Display" of all devices is performed, otherwise the device information acquisition / Display command (stgxfwcmsetdev) is executed in a Storage Management Server.
3. Reset the replication volume information using the replication volume information Configuration command (swsrpsetvol).



Note

When you use a replication destination volume by mounting or by allocating a drive letter after creating a replication, the size of the replication destination volume must exactly match the size of the replication destination volume to the byte. If you do not use a replication destination volume (for example, when you need copy data, you restore it to the replication source volume and then use it), the size of the replication destination volume can be larger than the size of the replication source volume.



Note

Group Configurations cannot be made in the following cases:

- When the group configuration conditions are not met
- When both servers involved in inter-server replication are not running version 13.0 or later



Note

This command cannot be executed while any of the following commands are running:

- Replication volume information Configuration command (swsrpsetvol)
- Replication volume information deletion command (swsrpdelvol)
- Synchronous processing start command (swsrpstartsync)
- Replication creation command (swsrpmake)
- Replication cancellation command (swsrpcancel)
- Synchronization mode change command (swsrpchsnc)
- Synchronization reverse command (swsrprevsync)
- Resource adjustment command (swsrprecoverres)

9.4.1.2 Replication volume information display command (swsrpvoinfo)

This command displays the replication volume information registered by the replication volume Configuration command (swsrpsetvol).

- a. Format

<p>[To be executed on a storage server] /opt/FJSVswsrp/bin/swsrpvoinfo [-L] [<i>Original-Volume-Name</i>] [To be executed on a storage management server] /opt/FJSVswsrp/bin/swsrpvoinfo -h <i>Server-Name</i> [-L] [<i>Original-Volume-Name</i>]</p>
--

b. Description of options

Option	Description
-h	Specifies a Storage Server name. This option is valid only on a Storage Management Server.
-L	Specifies that information will be displayed using the extended format. Refer to the table in the example following for information about the items that are displayed when this option is specified. The items that are displayed when this option is specified may be expanded to support new functions provided by future versions.

c. Description of operands

Operand	Description
Original-Volume-Name	Specifies a replication source volume name. For the volume names of other servers, use the format: "volume-name@<storage-server-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 Storage Server on which this command is executed is displayed.

d. End status

- = 0: Normal end
- > 0: Abnormal end

e. Example

Direct that all the replication volume information on a Storage Management Server be displayed on Storage Server SRC-SV.

```
# /opt/FJSVswsrp/bin/swsrpvoinfo -h SRC-SV
Server Original-Volume      Size      Replica-Volume      Size      Copy
Op-Server
SRC-SV /dev/dsk/c1t2d0@SRC-SV  4.0Gbyte /dev/dsk/c1t2d0@TARG-SV  4.0Gbyte
uni-direction original
SRC-SV /dev/dsk/c1t2d1@SRC-SV  4.0Gbyte /dev/dsk/c1t2d1@TARG-SV  5.0Gbyte
bi-direction original
SRC-SV /dev/dsk/c1t2d2@SRC-SV  9.1Gbyte /dev/dsk/c1t2d2@TARG-SV  9.1Gbte
bi-direction replica
SRC-SV /dev/dsk/c1t2d3@SRC-SV  7.5Gbyte /dev/dsk/c1t2d3@TARG-SV  7.5Gbte
bi-direction both
#
```


The following information will be displayed.

Title	Description
Server	A Storage Server name.
Original-Volume	A replication source volume name.
Size	The volume size. Additionally, in the case of a volume group, displays the size of the physical disk on which the logical disk exists.
Replica-Volume	A replication destination volume name.
Copy	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	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 if the - L option has been specified]	The group name. "----" is displayed if the copy source and destination volumes are not registered in a group.



Note

Replication volume information cannot be displayed if the specified volume has not been set as a replication source volume



Note

- This command cannot be executed while the Resource adjustment command (swsrprecoverres) is running.

9.4.1.3 Replication volume information deletion command (swsrpdelvol)

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.

a. Format

```
[To be executed on a storage server]
/opt/FJSVswsrp/bin/swsrpdelvol [-e] [-d] Original-Volume-Name Replica-Volume-Name
[To be executed on a storage management server]
/opt/FJSVswsrp/bin/swsrpdelvol -h Server-Name [-e] [-d] Original-Volume-Name
Replica-Volume-Name
```

b. Description of options

Option	Description
-h	Specifies a Storage Server name. This option is valid only on a Storage Management Server.
-e	Emergency operation mode. This deletes 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.

Option	Description
-d	Specifies that no communication takes 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.

c. Description of Operand

Operand	Description
Original-Volume-Name	Specifies a replication source volume name. For the volume names of other servers, use the format "volume-name@storage-server-name."
Replica-Volume-Name	Specifies a replication destination volume name. For the volume names of other servers, use the format "volume-name@storage-server-name."

d. End status

- = 0: Normal end
- > 0: Abnormal end

e. Example

Delete replication volume information on Storage Server SRC-SV.

```
# /opt/FJSVswsrp/bin/swsrpdelvol /dev/dsk/c1t2d0 /dev/dsk/c1t2d0@TARG-SV
swsrpdelvol completed
#
```



Note

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 Storage Server that executed the command is not the operation server of the replication volume.
- Communication with a Storage Management Server fails.
- For server to server replication, communication with a destination server fails.
- Changes have been made to the logical disk configuration that cannot be supported by AdvancedCopy Manager.
- A volume group is inactive in the volume group specification.
- If a logical volume is deleted, or if changes are made to its configuration that cannot be supported by AdvancedCopy Manager, replica volume information cannot be deleted. In this case, perform processing of (A) or (B) shown below.
 - (A): Specify the "-e" option and execute the command, after checking that an advanced copy is not processing in the replica volume to be deleted (if an advanced copy is running, stop the advanced copy with disk array equipment, or perform (B)).
 - (B): Execute the command after re-creating the logical volume with a configuration that AdvancedCopy Manager can support, or after returning the logical volume configuration to its original state.
- If emergency operation mode is used, OPC sessions are neither confirmed nor cancelled. If an OPC session is in the error suspended state or halt state, then unmanaged sessions remain. In this event, ETERNUSmgr or GRmgr must be used to cancel these sessions.



Note

This command cannot be executed while any of the following commands are running:

- Replication volume information Configuration command (swsrpsetvol)
- Replication volume information deletion command (swsrpdelvol)
- Synchronous processing start command (swsrpstartsync)
- Replication creation command (swsrpmake)
- Replication cancellation command (swsrpcancel)
- Synchronization mode change command (swsrpchsync)
- Synchronization reverse command (swsrprevsync)
- Resource adjustment command (swsrprecoverres)

9.4.1.4 REC transfer buffer information display command (swsrprebuffstat)

This command displays information for the transfer buffer used by ETERNUS when REC is performed in Consistency mode.

a. Format

```
[To be executed on a storage server]
swsrprebuffstat Volume-Name
[To be executed on a storage management server]
swsrprebuffstat -h Server Volume-Name
```

b. Description of options

Option	Description
-h	Specifies the name of a Storage Server. You can use this option only for a Storage Management Server.

c. Description of operands

Operand	Description
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 Storage Server where the command is executed. Specify the AdvancedCopy Manager device name as the volume name.

d. End status

- =0:Normal end
- >0:Abnormal end

e. Example

Display information for the transfer buffer used by the REC, where the copy source volume is /dev/dsk/clt2d10@Server1 and the copy destination volume is /dev/dsk/clt2d11@Server1, and replication is performed within the Storage Server "Server1":

```
# /opt/FJSVswsrp/bin/swsrprebuffstat /dev/dsk/clt2d10@Server1
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use Rate RemoteBoxID
1 SEND 12% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
# /opt/FJSVswsrp/bin/swsrprebuffstat /dev/dsk/clt2d11@Server1
BoxID = BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
ID Use Rate RemoteBoxID
1 RECV 0% AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
#
```

Display information for the transfer buffer used by the REC in an environment where replication is being performed between Storage Server "Server1" and Storage Server "Server2". Additionally, the copy source volume is /dev/dsk/clt2d10@Server1 and the copy destination volume is /dev/dsk/clt2d20@Server2:

```
[Command executed on storage server "Server1"]
# /opt/FJSVswsrp/bin/swsrprebuffstat /dev/dsk/clt2d10@Server1
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use Rate RemoteBoxID
1 SEND 12% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
2 RECV 0% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
#
[Command executed on storage server "Server2"]
# /opt/FJSVswsrp/bin/swsrprebuffstat /dev/dsk/clt2d20@Server2
BoxID = BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
ID Use Rate RemoteBoxID
1 RECV 0% AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
2 SEND 5% AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
#
```

Display information for the transfer buffer used by the REC if a cabinet where the copy source volume /dev/dsk/clt2d10@Server1 resides has been set up to make remote copies to multiple cabinets, and replication is being performed between Storage Server "Server1" and Storage Server "Server2":

```
[Command executed on storage server "Server1"]
# /opt/FJSVswsrp/bin/swsrprebuffstat /dev/dsk/clt2d10@Server1
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use Rate RemoteBoxID
1 SEND 12% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
2 RECV 0% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
3 SEND 4% CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
4 RECV 0% CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
5 SEND 10% DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD
6 RECV 0% EEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEE
7 SEND 3% FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
8 SEND 5% GGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGG
#
```

Display information for the transfer buffer used by the REC, where the copy source volume is /dev/sfdsk/class1/dsk/volume1@Server1 and the copy destination volume is /dev/sfdsk/class2/dsk/volume1@Server2, and replication is being performed between Storage Server "Server1" and Storage Server "Server2":
(The SDX object is in a two-cabinet mirror configuration.)

```
[Command executed on storage server "Server1"]
# /opt/FJSVswsrp/bin/swsrprebuffstat -Xaltpath
/dev/sfdsk/class1/dsk/volume1:mlb2048
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use Rate RemoteBoxID
1 SEND 12% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
2 RECV 0% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
# /opt/FJSVswsrp/bin/swsrprebuffstat -Xaltpath
/dev/sfdsk/class1/dsk/volume1:mlb3048
BoxID = CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
ID Use Rate RemoteBoxID
1 SEND 11% DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD
2 RECV 0% DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD
```

```
#
[Command executed on storage server "Server2"]
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath
/dev/sfdsk/class2/dsk/volume1:mp1b2148
BoxID = BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
ID Use Rate RemoteBoxID
1 RECV 0% AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
2 SEND 4% AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath
/dev/sfdsk/class2/dsk/volume1:mp1b3148
BoxID = DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD
ID Use Rate RemoteBoxID
1 RECV 0% CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
2 SEND 5% CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
#
```

The following information will be displayed:

Title	Description
BoxID	The box identifier for the cabinet connected to the server where the command was executed.
ID	The ID for the transfer buffer.
Use	The attribute (usage) for the transfer buffer. SEND : Transfer buffer for sending data RECV : Transfer buffer for receiving data
Rate	The usage rate of the transfer buffer as a percentage. If the usage is SEND: Displays the usage rate of the transfer buffer as a percentage. This will never be 0%, even if REC is not being performed in asynchronous Consistency mode. If the usage is RECV : "0%" will always be displayed. "----" will be displayed if the transfer buffer cannot be used.
RemoteBoxID	The box identifier for the cabinet performing a remote copy.



Note

Information is displayed only for cabinets that contain replication volumes for the server where the command was executed.
If inter-cabinet 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 transfer buffers in both cabinets.



Note

It will take a short time for changes to REC transfer buffer configuration 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.



Note

This command cannot be executed while the Resource adjustment command (swsrprecoverses) is running.

9.4.1.5 REC transfer buffer configuration change command (swsrprecbuffset)

This command changes the configuration of the transfer buffer used by ETERNUS when REC is performed in Consistency mode.

Refer to "ETERNUSmgr User's Guide Setup / Maintenance" for information about the configuration method of the REC transfer buffer.

a. Format

[To be executed on a Storage Server]

```
swsrprecbuffset -Xbuffid buffer-id -Xbuffuse {SEND|RECV} Volume-Name
```

[To be executed on a Storage Management Server]

```
swsrprecbuffset -h Server -Xbuffid buffer-id -Xbuffuse {SEND|RECV} Volume-Name
```

b. Description of options

Option	Description
-h	Specifies the name of a Storage Server. You can use this option only for a Storage Management Server.
-Xbuffid	Specifies the ID for the transfer buffer.
-Xbuffuse	Specifies the usage of the transfer buffer. SEND : Transfer buffer for sending data RECV : Transfer buffer for receiving data

c. Description of operands

Operand	Description
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 Storage Server where the command is executed.

d. End status

=0:Normal end

>0:Abnormal end

e. Example

Change the attributes of the transfer buffer used by the REC where the copy source volume is /dev/dsk/clt2d10@Server1 and the copy destination volume is /dev/dsk/clt2d11@Server1, and replication is being performed within the Storage Server "Server1":

```
# /opt/FJSVswsrp/bin/swsrprecbufstat /dev/dsk/clt2d10@Server1
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use Rate RemoteBoxID
1 SEND 12% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
# /opt/FJSVswsrp/bin/swsrpbufset -Xbuffid 1 -Xbuffuse RECV /dev/dsk/clt2d10@Server1
swsrpbufset completed
# /opt/FJSVswsrp/bin/swsrprecbufstat /dev/dsk/clt2d10@Server1
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use Rate RemoteBoxID
1 RECV 0% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
# /opt/FJSVswsrp/bin/swsrprecbufstat /dev/dsk/clt2d11@Server1
BoxID = BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
ID Use Rate RemoteBoxID
1 RECV 0% AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
# /opt/FJSVswsrp/bin/swsrpbufset -Xbuffid 1 -Xbuffuse SEND /dev/dsk/clt2d11@Server1
swsrpbufset completed
# /opt/FJSVswsrp/bin/swsrprecbufstat /dev/dsk/clt2d11@Server1
BoxID = BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
ID Use Rate RemoteBoxID
1 SEND 3% AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
#
```

Change the attributes of the transfer buffer where the copy source volume is /dev/dsk/clt2d10@Server1, and replication is being performed between Storage Server "Server1" and Storage Server "Server2".

```
[Command executed on storage server "Server1"]
# /opt/FJSVswsrp/bin/swsrpbufstat /dev/dsk/clt2d10@Server1
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use Rate RemoteBoxID
1 SEND 12% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
2 RECV 0% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
# /opt/FJSVswsrp/bin/swsrpbufset -Xbuffid 1 -Xbuffuse SEND /dev/dsk/clt2d11@Server1
swsrpbufset completed
# /opt/FJSVswsrp/bin/swsrpbufset -Xbuffid 2 -Xbuffuse RECV /dev/dsk/clt2d11@Server1
swsrpbufset completed
# /opt/FJSVswsrp/bin/swsrprecbufstat /dev/dsk/clt2d11@Server1
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use Rate RemoteBoxID
1 SEND 4% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
2 RECV 0% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
#
```



Note

Configurations for dedicated REC buffers must be made in advance. Change the configuration while there are no asynchronous Consistency mode sessions in either the ACTIVE or ERROR SUSPEND state. It will take a short time for changes to buffer configuration to take effect, during which the transfer 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:

- Replication volume information Configuration command (swsrpsetvol)
- Replication volume information deletion command (swsrpdelvol)
- Synchronous processing start command (swsrpstartsync)
- Replication creation command (swsrpmake)
- Operation status display command (swsrpstat)
- Replication cancellation command (swsrpcancel)
- Synchronization mode change command (swsrpchsyc)
- Synchronization reverse command (swsrprevsync)
- Resource adjustment command (swsrprecoverres)

9.4.2 Operational commands

This section describes commands for replication management.

9.4.2.1 Synchronous processing start command (swsrpstartsync)

This command starts the synchronous type 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.

a. Format

```
[To be executed on a Storage Server]
/opt/FJSVswsrp/bin/swsrpstartsync [[-y [-a | -v] | -k | -i] [-g] [-m] [-Xskip | -Xremain]
[-t] From-Volume-Name To-Volume-Name
[To be executed on a Storage Management Server]
/opt/FJSVswsrp/bin/swsrpstartsync -h Server-Name [[-y [-a | -v] | -k | -i] [-g] [-m]
[-Xskip | -Xremain] [-t] From-Volume-Name To-Volume-Name
[if executed on a Storage Server: group specification]
/opt/FJSVswsrp/bin/swsrpstartsync [[-y [-a | -v] | -k | -i] [-g] [-m] [-Xskip | -Xremain]
[-t] [-Xreverse] -Xgroup Group-Name
[if executed on a Storage Management Server: group specification]
/opt/FJSVswsrp/bin/swsrpstartsync - h Server-Name [ -y [ -a | -v ] | -k | -i ] [-g] [-m]
[ -Xskip | -Xremain ] [-t] [-Xreverse] -Xgroup Group-Name
```

b. Description of options

Option	Description
-h	This option is valid only on a Storage Management Server. This option can be specified only on a Storage Management Server.
-y	Specifies synchronization in synchronous write mode. This option is valid only when total copy is started for inter-box replication. Not specifying the -y, -k, and -i options sets asynchronous write mode. Synchronization in intra-box replication is always handled in synchronous write 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. Specifying the -y option without the -a option or -v option sets Automatic Split mode.

Option	Description
-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. Specifying the -y option without the -a option or -v option sets Automatic Split mode. You must specify this option when specifying synchronization mode (-y option) on a box that does not support Automatic Split mode.
-k	Specifies the configuration 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 options sets asynchronous write mode. You cannot use this option when operating SDX objects in units of logical volumes.
-i	Specifies the configuration 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 options sets asynchronous write mode.
-g	Specifies the configuration 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.
-m	For server to server replication, this option specifies that communication processing with non-operation 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-operation 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. Please use this option only if any one of the following conditions apply: <ul style="list-style-type: none"> • If AdvancedCopy Manager does not perform pre-processing and post-processing of the replication destination volume, and independent pre-processing and post-processing is incorporated before and after AdvancedCopy Manager commands • If it can be determined that replication destination pre-processing and post-processing are not necessary (Example: a database constructed on a RAW device) • After a forcible suspend, if restarting inter-cabinet synchronous processing (REC)
-Xgroup	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.
-Xremain	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 by specifying the "-Xskip" option and backup data has been copied from the tape medium to the copy destination volume.

c. Description of operands

Operand	Description
From-Volume-Name	Specifies a replication source volume name. This is the source volume or the destination volume set with the replication volume information configuration command (swsrpsetvol).
To-Volume-Name	Specifies a replication destination volume name. This is the source volume or the destination volume set with the copy volume information configuration command.

- d. End status
= 0: Normal end
> 0: Abnormal end

e. Example

Start replication from /dev/dsk/c1t2d0 of a Storage Management Server to /dev/dsk/c1t2d0@TARG-SV of Storage Server SRC-SV.

```
# /opt/FJSVswsrp/bin/swsrpstartsync -h SRC-SV /dev/dsk/c1t2d0 /dev/dsk/c1t2d0@TARG-SV
FROM=/dev/dsk/c1t2d0@SRC-SV,TO=/dev/dsk/c1t2d0@TARG-SV swsrpstartsync completed
#
```

Start replication for group "GRP1".

```
# /opt/FJSVswsrp/bin/swsrpstartsync -Xgroup GRP1
GROUP=GRP1 swsrpstartsync completed
#
```



Note

Replication cannot be started if:

- The specified replication source volume and replication destination volume have not been set as replication volumes.
- A replication source volume is specified as the copy source volume and a replication destination volume is specified as the cop destination 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 Storage 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.
- Changes have been made to the logical disk configuration that cannot be supported AdvancedCopy Manager.
- A volume group is inactive in the volume group specification.
- Option "-g" is specified, except for a remote copy.
- The backup management function or tape backup function uses a specified source/destination volume
- Either the copy source volume or the copy destination volume is a Snap Data Volume.



Note

If the "-Xgroup" option is specified, replication volumes are processed in order, but processing will be terminated immediately if an error occurs.



Note

When the replication corresponds to any of the following conditions, perform drive letter release or unmount of the copy source 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 the replication source volume and the replication destination volume differ

In the abovementioned cases, if neither drive letter release nor unmount have been performed, post-processing of the copy source volume at the time of duplicate creation will produce an error.



Point

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.



Note

- The processing of this command varies according to the status of the synchronization processing of the specified replication source and destination volumes.

Table: Processing variations in accordance with synchronization processing status

Status of synchronization processing	Processing
Unmount status	Start of total copy
Total or incremental copy status	Informational message (swsrp2401) is output and the command ends normally
Equivalency maintenance status	Informational message (swsrp2401) is output and the command ends normally
Replication established status (halt status)	Start of incremental copy



Note

Before starting the replication, run the replication pre-processing script for the replication destination volume. For information on customizing this script, refer to "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.



Note

To protect processing that accesses a replication destination volume, set the copied material volume before starting replication so that other processes cannot access it. To do this, cancel the drive letter of the replication destination volume or execute unmount. If you cannot cancel the drive letter of the replication destination volume or execute unmount 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 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 configuration, unmounting, or some other method.
- When configuring a volume group which consists of two or more logical disks as a source volume, and replicating it, make all logical devices unmount/mount using the Administrator's authority.

Refer to "General notes" for notes on starting replication.



Note

If the replica volume is a volume group, do not execute the `vgchange (lM)` command to deactivate the replica volume under synchronous processing (EC/REC) execution. If the replica volume is inactive during synchronous processing execution then it cannot be reactivated. Additionally, it will not be possible to perform commands including those to stop synchronous processing, the replication creation command (`swsrpmake`) and the replication cancel command (`swsrpcancel`). If it has been inactivated accidentally, make the replica volume active after stopping the synchronous processing using `ETERNUSmgr/GRmgr`, etc.



Note

Refer to "General notes" for notes on starting replication.



Note

This command cannot be executed while any of the following commands are running:

- Replication volume information configuration command (`swsrpsetvol`)
- Replication volume information deletion command (`swsrpdelvol`)
- Resource adjustment command (`swsrprecoverres`)

9.4.2.2 Replication creation command (`swsrpmake`)

When synchronization processing (EC) is not being performed, this command will start snapshot processing (OPC/QuickOPC/SnapOPC/SnapOPC+) to create a replication volume.

When synchronization processing is being performed, check whether 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 volume (ie, the replication destination volume) can be accessed.

This command executes the replication pre-processing and post-processing scripts before a replica is created. For details of these scripts, refer to "Pre-processing and Post-processing of Replication".

a. Format

```
[To be executed on a Storage Server]
/opt/FJSVswsrp/bin/swsrpmake [-m][-f][-t] ([-j | -T | -Xconcur | -C | -P])
[-Xgds-slice-skipchk] From-Volume-Name To-Volume-Name
[To be executed on a Storage Management Server]
/opt/FJSVswsrp/bin/swsrpmake -h Server-Name [-m][-f][-t] ([-j | -T | -Xconcur | -C |
-P])
[-Xgds-slice-slipchk] From-Volume-Name To-Volume-Name
[if executed on a Storage Server: group specification]
/opt/FJSVswsrp/bin/swsrpmake [-m] [-f] [-t] ([-j | -T | -Xconcur | -C | -P]) [-Xreverse]
-Xgroup Group-Name
[if executed on a Storage Management Server: group specification]
/opt/FJSVswsrp/bin/swsrpmake -h Server-Name [-m] [-f] [-t] ([-j | -T | -Xconcur | -C
| -P]) [-Xreverse] -Xgroup Group-Name
```

b. Description of options

Option	Description
-h	Specifies a Storage Server name. This option can be specified only on a Storage Management Server.
-m	For server to server replication, specifies that communication processing is not performed with the non-operational server. When this option is specified, volume status check processing, pre-processing and post-processing are not performed for the non-operational server volume. If the command is executed from the copy destination server and the Xconcur option is specified, this option cannot be specified. This option is valid only during server to server replication.
-f	Specifies that pre-processing and post-processing are not performed for a replication source volume. Use this option 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. It can also be used 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 pre-processing and post-processing are not performed for a replication destination volume. Use this option 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. It can also be used 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". This option can only be used for inter-cabinet replication; if attempted for intra-cabinet 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 inter-cabinet synchronous processing (REC), execute the replication start command (swsrpstartsync) with the -t option specified.
-T	Specifies that differential snapshot processing will be executed. This option cannot be specified while synchronous processing is being performed. This option is valid only for a replication within a cabinet and when the disk array 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 using differential snapshots.
-Xgroup	Specifies that operation statuses are displayed by group.
-Xreverse	Specifies that a replication will be created for each group in the reverse direction (ie, from replica volumes to original volumes).
-Xconcur	Specifies that a replication will be created using the concurrent suspension function. This option is not supported when GDS Snapshot linkage is being used. If this option is used, the size of the copy destination volume must be the same or larger than that of the copy source volume.
-C	Specifies that a SnapOPC will be started. This option is only valid for intra-cabinet replications and for disk array devices that support the SnapOPC function. In order to execute a SnapOPC, the copy destination volume must be a Snap Data Volume.

Option	Description
-P	Specifies that a SnapOPC+ session will be started. This option is only valid for intra-cabinet replications and for disk array devices that support 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: <ul style="list-style-type: none"> ● If the same SnapOPC+ session exists at the copy source volume and the copy destination volume ● If the total number of SnapOPC+ sessions at the copy source volume exceeds 8 sessions
-Xgds-slice-sk ipchk	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 option and the -t option must also be specified.

c. Description of operands

Operand	Description
From-Volume-Name	Specifies a replication source volume. This is the replication source volume or the replication destination volume that was set with the replication volume information Configuration command (swsrpsetvol).
To-Volume-Name	Specifies a replication destination volume. This is the replication source volume or the replication destination volume that was set with the replication volume information Configuration command (swsrpsetvol).

d. End status

- = 0: Normal end
- > 0: Abnormal end

e. Example

The Storage Management Server instructs Storage Server SRC-SV to create a copy of /dev/dsk/c1t2d0 in /dev/dsk/c1t2d0@TARG-SV.

```
# /opt/FJSVswsrp/bin/swsrpmake -h SRC-SV /dev/dsk/c1t0d2 /dev/dsk/c1t2d0@TARG-SV
FROM=/dev/dsk/c1t2d0@SRC-SV,TO=/dev/dsk/c1t2d0@TARG-SV swsrpmake completed
#
```

Creates a copy for the group "GRP1".

```
# /opt/FJSVswsrp/bin/swsrpmake -Xgroup GRP1
GROUP=GRP1 swsrpmake completed
#
```



Note

Replication cannot be performed 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, the operation for the corresponding volume cannot continue. Stop the replication operation for the volume, and then delete and reset the replication volume information.
- A replication source volume or a replication destination volume has multiple mount points.

- For server to server replication, the Storage 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.
- Changes have been made to the logical disk configuration which cannot be supported by AdvancedCopy Manager.
- A volume group is inactive in the volume group specification.
- The replication destination volume cannot be locked
- The contents of the replication destination volume lock operation specification file are invalid
- The backup management function or tape backup function uses a specified source or destination volume



Note

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

When the replication corresponds to either of the following conditions, perform drive letter release or unmount of the copy source 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 neither drive letter release nor unmount have been performed, post-processing of the copy source volume at the time of duplicate creation will produce an error.



Point

Before a replica is created, the pre-processing and post-processing scripts for the source volume are executed (for synchronized replication, only the pre-processing script is executed). After the replica has been created, the post-processing for the source volume and destination volume is executed. For information on customizing these scripts, refer to "Pre-processing and Post-processing of Replication". If you do not want to implement the pre-processing and post-processing scripts, use the "-f" option and the "-t" option.

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.



Point

If this command is executed during snapshot processing, stop the copy processing in progress and start the snapshot processing again.



Point

To protect the data, set the replication source volume before creating a replica so that other processes cannot access the volume. To do so, cancel the drive letter of the source volume or execute unmount. If you cannot cancel the drive letter of the source volume or execute unmount because of the situations described below, this command will fail and terminate:

- 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 configuration, unmounting, or some other method. If for whatever reason you do not want to remove the share configuration of the replication source volume or execute unmount or some other method, refer to "Pre-processing and Post-processing of Replication" and note the following:
 - When configuring a volume group which consists of two or more logical disks as a source volume, and replicating it, make all logical devices unmount/mount using the Administrator's authority.
 - When the replication source volume is being used by backup management, unmount may not be possible. If so, "swsrp2613 An error occurred on the pre-processing script at replication creation. Error code = 2" will be output.
 - 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.



Note

For Snap Data Volumes, only SnapOPC/SnapOPC+ can be executed. Additionally, when data is restored from a Snap Data Volume, there must be a SnapOPC session between the copy source volume and the copy destination volume (the Snap Data Volume).



Note

Refer to "General notes" for notes on creating replicas.



Point

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 unmounting the replication destination volume. In cases where the replication destination volume cannot be locked or unmounted, 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 canceling share configuration or unmounting for the replication operation, using the Administrator's authority.



- This command cannot be executed while any of the following commands are running:
- Replication volume information configuration command (swsrpsetvol)
 - Replication volume information deletion command (swsrpdelvol)
 - Resource adjustment command (swsrprecoverres)



- If this command is executed during snapshot processing, snapshot processing is restarted. Refer to the following table for details of copy operation:

Table: Copy operation if the command is executed during snapshot processing

Operation status (Display by swsrpstat -L)	-T option specificat ion	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 that is in progress is terminated and ordinary snapshot processing is restarted.
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.

9.4.2.3 Operation status display command (swsrpstat)

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.

a. Format

```
[To be executed on a Storage Server]
/opt/FJSVswsrp/bin/swsrpstat [-L] [-E] [-H] [-O] [Original-Volume-Name]
[To be executed on a Storage Management Server]
/opt/FJSVswsrp/bin/swsrpstat -h Server-Name [-L] [-E] [-H] [-O] [Original-Volume-Name]
[if executed on a Storage Server: group specification]
/opt/FJSVswsrp/bin/swsrpstat [-L] [-E] [-O] -Xgroup Group-Name
[if executed on a Storage Management Server: group specification]
/opt/FJSVswsrp/bin/swsrpstat -h Server-Name [-L] [-E] [-O] -Xgroup Group-Name
```

b. Description of options

Option	Description
-h	Specifies a Storage Server name. This option can be specified only on a Storage Management Server.
-L	Specifies that display will be in extension format. For the items displayed by this option, refer to the table in the following example following. Note that the items displayed by this option may be subject to change in a future release.
-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. 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	Specifies that the operation statuses are displayed by 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. Note that hard-suspend status will still be displayed as "halt" even if this option is specified for ROPC.
-0	Changes the Status field display method used when in error-suspended (failed) status. If this option is specified, the display is as follows: <ul style="list-style-type: none"> ● If there is a bad sector: "failed(badsector)" ● If there is a Snap Data Volume or Snap Data Pool capacity shortage: "failed(overflow)" In other cases: "failed(other)"

c. Description of operands

Operand	Description
Original-Volume-Name	Specifies a replication source volume name. For the volume names of other servers, use the format "volume-name@storage-server-name." Displays all operation statuses for the specified replication source volume. If this operand is omitted, all the operation statuses of the Storage Server on which this command is executed are displayed.

d. End status

- = 0: Normal end
- > 0: Abnormal end

e. Example

Direct that all the replication statuses for Storage Server SRC-SV be displayed on the Storage Management Server.

```
# /opt/FJSVswsrp/bin//opt/FJSVswsrp/bin/swsrpstat -h SRC-SV
Server Original-Volume      Replica-Volume      Direction Status      Execute
SRC-SV /dev/dsk/c1t2d0@SRC-SV /dev/dsk/c1t2d0@TARG-SV regular replicated ----
SRC-SV /dev/dsk/c1t2d1@SRC-SV /dev/dsk/c1t2d1@TARG-SV ----      ----      ----
SRC-SV /dev/dsk/c1t2d2@SRC-SV /dev/dsk/c1t2d2@TARG-SV reverse snap      45%
#
```

From Storage Management Server to Storage Server: Instruct SRC-SV to display the status of all replication operations (with the -L option specified).

```
# /opt/FJSVswsrp/bin/swsrpstat - L
Server Original-Volume      Replica-Volume      Direction  Status
Execute Trk  Update Rcv  Split Xfer  Snap-Gen
SRC-SV /dev/dsk/c1t2d0@SRC-SV /dev/dsk/c1t2d0@TARG-SV regular    replicated
-----
----- auto ----- async -----
SRC-SV /dev/dsk/c1t2d1@SRC-SV /dev/dsk/c1t2d1@TARG-SV -----
----- on 6% -----
SRC-SV /dev/dsk/c1t2d2@SRC-SV /dev/dsk/c1t2d2@TARG-SV reverse    snap
45% off -----
SRC-SV /dev/dsk/c1t1d0@SRC-SV /dev/dsk/c1t1d1@TARG-SV regular
copy-on-write(inactive) ----- off 0% ----- 1
SRC-SV /dev/dsk/c1t1d0@SRC-SV /dev/dsk/c1t1d2@TARG-SV regular
copy-on-write(inactive) ----- off 5% ----- 2
SRC-SV /dev/dsk/c1t1d0@SRC-SV /dev/dsk/c1t1d3@TARG-SV regular
copy-on-write(active) ----- off 6% ----- 3
#
```

The following information will be displayed.

Title	Description
Server	Displays a Storage Server name.
Original-Volume	Displays a replication source volume name.
Replica-Volume	Displays a replication destination volume name.
Direction	Displays the copy direction. <ul style="list-style-type: none"> 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 progress. <ul style="list-style-type: none"> ----- : Copying is not being performed. sync : Total copy or incremental copy is being performed. equivalent : Equivalency maintenance status replicated : Replication established status copy-on-write: SnapOPC processing is being performed. copy-on-write(active): SnapOPC+ processing is being performed (latest Snap generation). copy-on-write(inactive): SnapOPC+ processing is being performed (an earlier Snap generation). failed : Error suspend status failed(badsector): Bad sector at the copy source failed(overflow): Snap Data Volume or Snap Data Pool capacity shortage failed(other): An error suspend status other than failed(badsector) or failed(overflow) halt : Hard suspend status halt(sync): Hard suspend status (ie, the hard suspend status has been entered during a REC total copy or differential copy) halt(equivalent): Hard suspend status (ie, the hard suspend status has been entered during an REC equivalency maintenance state) snap : OPC or QuickOPC processing is being performed. ????? : Copy status cannot be determined.
Execute	Displays the copy progress rate as a percentage. "-----" is displayed during SnapOPC or when no copying is being performed.

Title	Description
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 any of the following cases. For SnapOPC+, displays the percentage of data updated since the previous Snap generation <ul style="list-style-type: none"> • When the Status column is "----" and the Trk column is "on" (ie, tracking is in progress and the physical copy is not executed). • When the Status column is "copy-on-write" (ie, when SnapOPC processing is being performed) • When the Status column is "copy-on-write(active)" (that is, when SnapOPC+ processing (the latest Snap generation) is in progress) • When the Status column is "copy-on-write(inactive)" (that is, when SnapOPC+ processing (an earlier Snap generation) is in progress) In all other cases, "----" is displayed.
Rcv	Indicates recovery mode for inter-box synchronization. auto: Automatic Recovery mode. manual: Manual Recovery mode. ----: Synchronous duplication (REC) is not performed between boxes.
Split	Indicates split mode for inter-box synchronization. auto: Automatic Split mode. manual: Manual Split mode. ----: Synchronous duplication (REC) is not performed between boxes.
Xfer	Indicates transfer mode in inter-box synchronization. sync: Synchronous mode. async: Asynchronous mode. stack: Stack mode. consist: Consistency mode. ----: Synchronous duplication (REC) is not performed between boxes.
Snap-Gen	Displays the Snap generation number (1 to 8) if SnapOPC+ is being performed. If SnapOPC+ is not being performed, "----" is displayed. A Snap generation number of 1 indicates the oldest generation, and each time the number increases it becomes the latest generation.



Note

If the execution status (displayed in the Status field) is "failed", "halt", or "?????", refer to "Troubleshooting for a hardware error, etc. occurring during replication" and take the required action.



Note

- The operation status cannot be displayed if:
- The volumes have not been set as replication volumes.
 - Changes have been made to the logical disk configuration which cannot be supported AdvancedCopy Manager.
 - A volume group is inactive in the volume group specification.



- This command cannot be executed while the Resource adjustment command (swsrprecoverres) is running.

9.4.2.4 Replication cancellation command (swsrpcancel)

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 and when you want to stop replication processing.

a. Format

```
[To be executed on a Storage 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 a Storage 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
[If executed on a Storage Server: group specification]
/opt/FJSVswsrp/bin/swsrpcancel [-c] ([-f] [-t] [-m] [-b])|[-T] |[-Xforce] [-Xreverse]
-Xgroup Group-Name
[If executed on a Storage Management Server: group specification]
/opt/FJSVswsrp/bin/swsrpcancel -h Server-Name [-c] ([-f] [-t] [-m] [-b])|[-T]
|[-Xforce] [-Xreverse] -Xgroup Group-Name
```

b. Description of options

Option	Description
-h	Specifies a Storage Server name. This option can only be specified on a Storage Management Server.
-c	Specifies that copy processing be canceled from servers other than the operation server if the operation server becomes unavailable (eg, due to system failure). If this option is used, copy processing is stopped without executing 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 for the replication destination volume of a non operational server are not performed. This option is valid only during server to server replication.
-f	When stopping the synchronous processing under the equivalency maintenance status, specifies that pre-processing and post-processing for the copy source volume is not to be executed. Use this option 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. It can also be used 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 in the equivalency maintenance status, this option specifies that post-processing for the copy destination volume will not be executed. Use this option 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. It can also be used when you judge that copy pre-processing and post-processing are unnecessary (e.g., for a database constructed on a raw device).

Option	Description
-b	Specifies that an inter-cabinet copy operation in the halt status be stopped. For replication on a server, the copying is cancelled for both cabinets. For replication between servers, the copying is cancelled only for the cabinet connected to the server on which this command was executed. The option is valid only for a copy operation in the halt status, and should only be performed if the operation is not expected to be recoverable. This option can be used only when the copy processing status is halt. The option can be used on a server not specified as the operation server. 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). This option releases the OPC session if snapshot processing (OPC physical copying) is complete. This option can be used only while physical copying or tracking is in progress, and 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	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 option and the "-t" option must also be specified.
-Xforce	This option enables a SnapOPC+ session to be forcibly stopped. Forced stoppage of 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 specific to SnapOPC+. If this option is specified at the same time as the -Xgroup option, the SnapOPC+ sessions within the group are stopped, and the Snap generation and all earlier Snap generations within the group are stopped.

c. Description of operands

Operand	Description
From-Volume-Name	Specifies a copy source volume name. This is the replication source volume or the replication destination volume set with the replication volume information configuration command (swsrpsetvol).
To-Volume-Name	Specifies a copy destination volume name. This is the replication source volume or the replication destination volume set with the replication volume information configuration command (swsrpsetvol).

d. End status

- = 0: Normal end
- > 0: Abnormal end

e. Example

On the Storage Management Server, direct that copying from /dev/dsk/c1t2d0 to /dev/dsk/c1t2d0@TARG-SV is stopped for Storage Server SRC-SV.

```
#/opt/FJSVswsrp/bin/swsrpcancel -h SRC-SV /dev/dsk/c1t2d0 /dev/dsk/c1t2d0@TARG-SV
FROM=/dev/dsk/c1t2d0@SRC-SV, TO=/dev/dsk/c1t2d0@TARG-SV swsrpcancel completed
#
```

Stops copy processing for group "GRP1".

```
# /opt/FJSVswsrp/bin/swsrpcancel -Xgroup GRP1
GROUP=GRP1 swsrpcancel completed
#
```



Note

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 cabinets.



Note

Copy processing cannot be performed if:

- The specified source and destination volumes are not 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 Storage 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.
- Changes have been made to the logical disk configuration that cannot be supported AdvancedCopy Manager.
- A volume group is inactive in the volume group specification.
- The backup management function or tape backup function uses a specified source or destination volume



Note

If the -Xgroup option is specified, processing will be performed for each active replication volume, but processing will be terminated without canceling replication in the following cases:

- If no replication volumes have been copied
- If you execute the "-b" option when no instance of copy processing is currently in the halt state in the group.
- If you execute the "-t" option when no tracking processing of QuickOPC is currently occurring in the group.

Additionally, if the "-Xgroup" option is specified, the target replication volumes are processed in order, but processing will be terminated immediately if an error occurs.



Note

When SnapOPC+ sessions are stopped, they must be stopped in sequence from the oldest Snap generation.

If the second and subsequent generations are to be forcibly stopped, specify the -Xforce option. However, note that the specified Snap generation and all earlier Snap generations will be stopped.

If the -Xforce option is specified at the same time as the -Xgroup option, the SnapOPC+ sessions within the group are stopped, and the Snap generation and all earlier Snap generations within the group are stopped.



Note

The volume information for the specified volume is not checked (ie, there is no check to match the volume information in the management list and the actual physical information). The operation is canceled unconditionally using the volume information in the management list.



Point

- The following processing is performed depending on the operation status of the specified replication source volume and replication destination volume.

Table: Processing variations in accordance with operation 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.
Equivalency maintenance 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.



Point

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.

If replication is canceled before the equivalency maintenance status is achieved, the replica volume is placed in the incomplete state and, in the case of a file system, the replica volume cannot be mounted.

If the processing target is a volume group, the operation may be customized so that the pre-processing and post-processing scripts for the replica volume are used to mount and unmount the volume. If replication is canceled before the equivalency maintenance state is achieved under these conditions, take action as follows:

- To start synchronous processing
Execute the replication start command (swsrpstartsync) with the "-t" option specified.
- To execute replication
Execute the replication creation command (swsrpmake) with the "-t" option specified.
After replication is completed, mount the replica volume.



Note

This command cannot be executed while any of the following commands are running:

- Replication volume information configuration command (swsrpsetvol)
- Replication volume information deletion command (swsrpdelvol)
- Resource adjustment command (swsrprecoverres)



Point

- When the "-t" option is not specified, command processing varies depending on the operation status of the source and destination volumes.

Table: Processing when -T option not specified

<u>Operation status</u> (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 maintenance 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: Processing when the -T option is specified

<u>Operation status</u> (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.
Equivalency maintenance 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.

9.4.2.5 Synchronization mode change command (swsrpchsnc)

This command changes the operational mode of inter-box synchronization (transfer mode, recovery mode, or split mode).

- Format

```
[To be executed on a Storage Server]
/opt/FJSVswsrp/bin/swsrpchsnc [-F (Sync|Async|Stack|Consist)] [-R (Auto|Manual)] [-S
(Auto|Manual)] From-Volume-Name To-Volume-Name

[To be executed on a Storage Management Server]
/opt/FJSVswsrp/bin/swsrpchsnc -h Server-Name [-F (Sync|Async|Stack|Consist)] [-R
```

```
(Auto|Manual)] [-S (Auto|Manual)] From-Volume-Name To-Volume-Name

[if executed on a Storage Server: group specification]
/opt/FJSVswsrp/bin/swsrpchsnc [-F (Sync|Async|Stack|Consist)] [-R (Auto|Manual)] [-S
(Auto|Manual)] [-Xreverse] -Xgroup Group-Name

[if executed on a Storage Management Server: group specification]
/opt/FJSVswsrp/bin/swsrpchsnc -h Server-Name [-F (Sync|Async|Stack|Consist)] [-R
(Auto|Manual)] [-S (Auto|Manual)] [-Xreverse] -Xgroup Group-Name
```

b. Description of options

Option	Description
-h	Specifies the name of a Storage Server. You can use this option only for a Storage Management Server.
-F	Changes the transfer mode. You can specify one of the following parameters in the operand: Sync: Changes to Synchronous mode. Async: Changes to Asynchronous (sequential transfer) mode. Stack: Changes to Stack mode. Consist: Changes to Consistency mode. 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	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.
-S	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. To specify this option, you must also specify -F Sync, or the transfer mode of the current synchronization must be synchronous mode.
-Xgroup	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.

c. Description of operands

Operand	Description
From-Volume-Name	Specifies the name of a source volume. This is the source volume or replica volume configured by the replication volume information configuration command (swsrpsetvol).
To-Volume-Name	Specifies the name of a destination volume. This is the source volume or replica volume configured by the replication volume information configuration command (swsrpsetvol).

d. End status

- = 0: Normal end
- > 0: Abnormal end

e. Example

From Storage Management Server to Storage Server: Instruct SRC-SV to change the transfer mode for inter-box synchronization between /dev/dsk/c1t2d3@SRC-SV and /dev/dsk/c1t2d3@TARG-SV from Consistency mode to Asynchronous mode.

```
#/opt/FJSVswsrp/bin/swsrpstat -h SRC-SV -L /dev/dsk/c1t2d3@SRC-SV
Server Original-Volume      Replica-Volume      Direction Status Execute Trk
Rcv Split Xfer Snap-Gen
SRC-SV /dev/dsk/c1t2d3@SRC-SV /dev/dsk/c1t2d3@TARG-SV regular sync 99% ----
auto ---- consist ----
:
#/opt/FJSVswsrp/bin/swsrpchsnc - h SRC-SV - F Async /dev/dsk/c1t2d3@SRC-SV
/dev/dsk/c1t2d3@TARG-SV
FROM=/dev/dsk/c1t2d3@SRC-SV, TO=/dev/dsk/c1t2d3@TARG-SV swsrpchsnc completed
:
#/opt/FJSVswsrp/bin/swsrpstat -h SRC-SV - L /dev/dsk/c1t2d3@SRC-SV
Server Original-Volume      Replica-Volume      Direction Status Execute
Trk Rcv Split Xfer Snap-Gen
SRC-SV /dev/dsk/c1t2d3@SRC-SV /dev/dsk/c1t2d3@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 asynchronous mode.

```
# /opt/FJSVswsrp/bin/swsrpchsnc -F Async -Xgroup GRP1
GROUP=GRP1 swsrpchsnc completed
#
```



Note

In any of the following cases, you cannot change the operation mode of inter-box synchronization:

- The specified source/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 Storage 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 has error-suspended status (with "failed" indicated by the operating status display command (swsrpstat)) or hard-suspended status (with "halt" indicated by the operating status display command (swsrpstat)).
- The disk array in which the specified source/replica volume is placed does not support the operation mode change function.
- Changes have been made to the logical disk configuration that AdvancedCopy Manager does not support.
- The specified volume group is inactive.



Note

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 one or more replication volumes meet any of 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.



Note

If the REC operating mode is changed while in equivalency maintenance status, after the operating mode is changed, the user must wait until the equivalence maintenance status is attained.



Note

This command cannot be executed while any of the following commands are running:

- Replication volume information configuration command (swsrpsetvol)
- Replication volume information deletion command (swsrpdelvol)
- Resource adjustment command (swsrprecoverres)

9.4.2.6 Synchronization reverse command (swsrprevsync)

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.

a. Format

```
[To be executed on a Storage Server]
/opt/FJSVswsrp/bin/swsrprevsync From-Volume-Name To-Volume-Name

[To be executed on a Storage Management Server]
/opt/FJSVswsrp/bin/swsrprevsync -h Server-Name From-Volume-Name To-Volume-Name

[if executed on a Storage Server: group specification]
/opt/FJSVswsrp/bin/swsrprevsync [-Xreverse] -Xgroup Group-Name

[if executed on a Storage Management Server: group specification]
/opt/FJSVswsrp/bin/swsrprevsync -h Server-Name [-Xreverse] -Xgroup Group-Name
```

b. Description of options

Option	Description
-h	Specifies the name of a Storage Server. You can only use this option for a Storage Management Server.
-Xgroup	Changes the direction of synchronous processing for each group.
-Xreverse	Specifies that the direction of synchronous processing that is being executed in reverse direction (from replica volumes to original volumes) will be inverted.

c. Description of operands

Option	Description
From-Volume-Name	Specifies the name of a source volume (as the destination in the reverse direction). This is the source volume or replica volume configured by the replication volume information configuration command (swsrpsetvol).
To-Volume-Name	Specifies the name of a destination volume (as the source in the reverse direction). This is the source volume or replica volume configured by the replication volume information configuration command (swsrpsetvol).

d. End status

- = 0: Normal end
- > 0: Abnormal end

e. Example

From Storage Management Server to Storage Server: Instruct SRC-SV to reverse the copying direction of synchronization between /dev/dsk/c1t2d3@SRC-SV and /dev/dsk/c1t2d3@TARG-SV.

```
#/opt/FJSVswsrp/bin/swsrpstat -h SRC-SV /dev/dsk/c1t2d3@SRC-SV
Server Original-Volume      Replica-Volume      Direction  Status    Execute
SRC-SV /dev/dsk/c1t2d3@SRC-SV /dev/dsk/c1t2d3@TARG-SV reverse    replicated ----
#/opt/FJSVswsrp/bin/swsrprevsync /dev/dsk/c1t2d3@TARG-SV /dev/dsk/c1t2d3@SRC-SV
FROM=/dev/dsk/c1t2d3@SRC-SV, TO=/dev/dsk/c1t2d3@TARG-SV swsrprevsync completed
#/opt/FJSVswsrp/bin/swsrpstat /dev/dsk/c1t2d3@SRC-SV
Server Original-Volume      Replica-Volume      Direction  Status    Execute
SRC-SV /dev/dsk/c1t2d3@SRC-SV /dev/dsk/c1t2d3@TARG-SV regular    replicated ----
#
```

Change the direction of synchronous processing for group "GRP1".

```
# /opt/FJSVswsrp/bin/swsrprevsync -Xgroup GRP1
GROUP=GRP1 swsrprevsync completed
#
```



Note

- In any of the following cases, you cannot reverse the direction of synchronization:
- The specified source/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 Storage 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 disk array in which the specified source/replica volume is placed does not support the reverse function.
 - Changes have been made to the logical disk configuration that AdvancedCopy Manager does not support.
 - The specified volume group is inactive.



Note

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 one or more replication volumes meet any of 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.



Note

If the -Xgroup option is specified, processing is not performed for those sessions in the group where the direction has already been changed.



Note

This command cannot be executed while any of the following commands are running:

- Replication volume information configuration command (swsrpsetvol)
- Replication volume information deletion command (swsrpdelvol)

- Resource adjustment command (swsrprecoverres)

9.4.3 Maintenance commands

This section describes the maintenance commands for replication management.

9.4.3.1 Resource adjustment command (swsrprecoverres)

This command executes recovery to make replication management information consistent if a system failure or other problem has made the information inconsistent.

d. Format

```
[To be executed on a Storage Server]
/opt/FJSVswsrp/bin/swsrprecoverres [-r] [-p]
[To be executed on a Storage Management Server]
/opt/FJSVswsrp/bin/swsrprecoverres -h Server-Name [-r] [-p]
```

e. Description of options

Option	Description
-h	Specifies a Storage Server name. This option can be used only on a Storage Management Server.
-r	Specify this option when an incorrect status exists in the replication management files (for example, when the replication volume information configuration command (swsrpsetvol) or the replication volume information deletion command (swsrpdelvol) terminates suddenly because of a forced stop or a system failure, or when the replication management files have been deleted by mistake). When this option is specified, the following processing is performed to re-create the replication management files. Replication management information in the Storage Management Server repository is downloaded to the Storage Server, and the replication volume information is restored. The disk array device is queried, and the execution information for the synchronization processing is restored.
-p	Specifies not to perform any post-processing that has not already been performed. This option can be used when the replication creation command (swsrpmake) processing terminates suddenly because of a forced stop or system failure.

f. End status

- = 0: Normal end
- > 0: Abnormal end

g. Example

```
Perform resource adjustment.

#/opt/FJSVswsrp/bin/swsrprecoverres
swsrprecoverres completed
#
```



Note

The resource adjustment processing cannot be performed if:

- Communication with a Storage Management Server fails and the "-r" option is specified.
- The "-r" option is not specified when this command is executed after processing by the replication volume information configuration command (swsrpsetvol) and the replication volume information deletion command (swsrpdelvol) has been forcibly terminated.
- The "-r" option is specified and the command terminates with an error due to one of the following:
 - Changes have been made to the logical disk configuration that cannot be supported by AdvancedCopy Manager.
 - A volume group is in active in the volume group specification.



Note

The consistency of resource information is restored only for the Storage 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 Storage Server on which the processing was performed and on the Storage Server with which communication is being made.



Note

If this command is executed at a time when the copy creation command (swsrpmake) is running and has been interrupted (that is, the copy process has not completed for the source volume or the destination volume), replication post-processing will be performed.



Note

This command cannot be executed while any of the following commands are running:

- Replication volume information configuration command (swsrpsetvol)
- Replication volume information display command (swsrpvoinfo)
- Replication volume information deletion command (swsrpdelvol)
- Synchronous processing start command (swsrpstartsync)
- Replication creation command (swsrpmake)
- Replication status display command (swsrpstat)
- Replication cancellation command (swsrpcancel)
- Synchronization mode change command (swsrpchsyc)
- Synchronization reverse command (swsrprevsync)
- Resource adjustment command (swsrprecoverres)

9.5 Start/Stop Daemon Commands

This section describes the ways in which AdvancedCopy Manager daemons can be started.



Note

Additional information about this command:

- Confirm that the SWSTGNODE environment variable is not defined.
- Execute this command with the same locale as the LANG environment variable used by the system.
- Only a root or superuser 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

Execute the following command to start the communication daemon:

```
# /opt/FJSVswstf/bin/stgfwcom start
```

- Stopping the communication daemon

Execute the following command to stop the communication daemon:

```
# /opt/FJSVswstf/bin/stgfwcom stop
```

9.5.2 Starting and stopping AdvancedCopy Manager daemons

The daemon listed below is started and stopped:

Table: Daemons started and stopped as a package

Server type	Storage Server
Daemon name	Communication daemon

- * For details on how to start a daemon of a Storage Management Server, refer to the "ETERNUS SF AdvancedCopy Manager Operator's Guide" relevant to the OS of the Storage Management Server.

- Starting an AdvancedCopy Manager daemon

Execute the following command to start an AdvancedCopy Manager daemon:

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

- Stopping an AdvancedCopy Manager daemon

Execute the following command to stop an AdvancedCopy Manager daemon:

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


9.6 SnapOPC/SnapOPC+ Support Commands

This section explains the commands that 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: Operating environment for executing SnapOPC/SnapOPC+ support commands

Command	Operating environment of AdvancedCopy Manager
swstestupdate	-
swstsdv	-



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 environment variable SWSTGNODE. For "bsh" shell, 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. For notes and cautions on command execution, see the Remarks column.

Table: Cluster operation commands

Command	Environment variable	Command execution node	Remarks
swstestupdate	Not required	Active cluster server	-
swstsdv	Not required	Active cluster server	-

9.6.1 Update volume measurement command (swstestupdate)

This command measures the physical capacity of Snap Data Volumes using the pseudo SnapOPC session setup function.

- a. Format

```
[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 (canceling the pseudo session)]
/opt/FJSVswsts/bin/swstestupdate stop Volume-Name
```

- b. Description of operands

Operand	Description
Volume-Name	Specifies the volume name. Specify a copy source volume that is used with SnapOPC/SnapOPC+ operations.

- c. End status
 - = 0 : Normal end
 - > 0 : Abnormal end

d. Example

Measure the update volume for /dev/dsk/clt2d10:

```
(Start taking measurements by starting a pseudosession)
#/opt/FJSVswsts/bin/swstestupdate start /dev/dsk/clt2d10
/dev/dsk/clt2d10 swstestupdate completed

(First measurement)
#/opt/FJSVswsts/bin/swstestupdate status /dev/dsk/clt2d10
Volume-Name      Update
/dev/dsk/clt2d10 644333

(Second measurement)
#/opt/FJSVswsts/bin/swstestupdate status /dev/dsk/clt2d10
Volume-Name      Update
/dev/dsk/clt2d10 1045000

(Stop taking measurements by canceling the pseudosession)
# /opt/FJSVswsts/bin/swstestupdate stop /dev/dsk/clt2d10
/dev/dsk/clt2d10 swstestupdate completed
#
```

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 (shown as the number of blocks, where 1 block = 512 bytes). Displays "----" if there is no pseudo session. Displays "failed" if the pseudo session is error-suspended.



Note

The amount of data updated cannot be measured in the following cases:

- When the specified volume does not exist.
- When the disk array device does not support the SnapOPC/SnapOPC+ function.
- When the pseudo session has already been set up (ie, when the "start" option is specified).
- When the pseudo session does not exist (ie, when the "stop" option is specified).

9.6.2 Snap Data Volume operation/reference command (swstdsv)

This command can be used to initialize a Snap Data Volume, display the Snap Data Volume status, or to display the Snap Data Pool status. When executing this command, specify the following subcommands in accordance with the desired operation.

Table: Subcommand list

Subcommand	Function
init	<u>Snap Data Volume initialization</u>
stat	<u>Snap Data Volume status display</u>
poolstat	<u>Snap Data Pool status display</u>



Note

If a capacity shortage has occurred for a Snap Data Volume, the access path and the ETERNUS disk array logical unit number (LU-Number) can be used to enable operations. However, when a Snap Data Volume is in the capacity shortage state, if the "stat" subcommand is specified with the intention of checking the LU-Number of the Snap Data Volume for which the capacity shortage occurred, access to that volume results in an access error. The command terminates abnormally, and the LU-Number cannot be checked. Therefore, if the LU-Number is to be used, the Snap Data Volume LU-Number must be checked in advance and remembered.

Snap Data Volume initialization

Specify "init" in the subcommand to initialize a Snap Data Volume.



Note

Execute Snap Data Volume initialization before SnapOPC/SnapOPC+ is started.



Point

Initialization of a Snap Data Volume discards all data in the Snap Data Volume. Therefore, the following procedures are recommended.

(When operations are performed from the logical volume manager, access paths must be prepared in advance and LU-Numbers must be collected.)

- 1) If the Snap Data Volume contains data that is still needed, use OS commands or similar to save data in advance.
- 2) Initialize the disk management information.
When operations are performed from the logical volume manager, set the targeted Snap Data Volume (physical disk) such that it is not managed by the logical volume manager.
- 3) Use this command to initialize the Snap Data Volume.
When operations are performed from the logical volume manager, use access paths to perform initialization.

a. SFormat

```
/opt/FJSVswsts/bin/swstdsv init Volume-Name
or
/opt/FJSVswsts/bin/swstdsv init -G Access-Path-Name LU-Number
```

b. Description of options

Option	Description
-G	Specifies the access path (Access-Path-Name). Refer to "Creating an access path" for details of access paths.

c. Description of operands

Operands	Description
Volume-Name	Specifies the volume name of the Snap Data Volume (AdvancedCopy Manager device name) that is to be initialized. If partitions (slices) have been specified, Logical Units (disks) in the partitions (slices) are initialized.
LU-Number	Specifies the ETERNUS disk array logical unit number. Specify a decimal number.

d. End status

- = 0 : Normal end
- > 0 : Abnormal end

e. Example

Initialize the Snap Data Volume (/dev/dsk/c39t0d5).

```
# /opt/FJSVswsts/bin/swstdsv init /dev/dsk/c39t0d5
swstdsv init completed
#
```

If a capacity shortage has occurred, use the access path (/dev/dsk/c39t0d1) to initialize the volume of LU-Number (142).

```
# /opt/FJSVswsts/bin/swstdsv init -G /dev/dsk/c39t0d1 142
swstdsv init completed
#
```

Snap Data Volume status display

Specify "stat" in the subcommand to display the Snap Data Volume logical capacity, physical capacity, the size of the area in use and the size of the unused area.

a. Format

```
/opt/FJSVswsts/bin/swstdsv stat Volume-Name
or
/opt/FJSVswsts/bin/swstdsv stat - G Access-Path-Name LU-Number
```

b. Description of option

Option	Description
-G	Specifies the access path (Access-Path-Name). Refer to "Creating an access path" for details of access paths.

c. Description of operands

Operands	Description
Volume-Name	Specifies the volume name of the Snap Data Volume (AdvancedCopy Manager device name). If partitions (slices) have been specified, information of the Logical Units (disks) in the partitions (slices) is displayed.
LU-Number	Specifies the ETERNUS disk array logical unit number. Specify a decimal number.

d. End status

= 0 : Normal end
> 0 : Abnormal end

e. Example

Display the status of Snap Data Volume (/dev/dsk/c39t0d5).

```
# /opt/FJSVswsts/bin/swstdsv stat /dev/dsk/c39t0d5
BoxID = 00E4000M3#####E450S20A####KD4030639004##
LUN = 142 (0x8E)
Rate Logical (sector) Physical (sector) Used (sector) Copy (sector) Host (sector)
Pool (sector)
0% 8388608 1048576 0 0 0 0
```

Use the access path (/dev/dsk/c39t0d1) to display the status of LU-Number (142).

```
# /opt/FJSVswsts/bin/swstdsv stat -G /dev/dsk/c39t0d1 142
BoxID = 00E4000M3#####E450S20A####KD4030639004##
LUN = 142 (0x8E)
Rate Logical (sector) Physical (sector) Used (sector) Copy (sector) Host (sector)
Pool (sector)
0% 8388608 1048576 0 0 0 0
```

The following table shows the content displayed.

Title	Description
BoxID	Displays the Box ID.
LUN	Displays the logical unit number used at the ETERNUS disk array. The format of the display is "<decimal number>(<hexadecimal number>)".
Rate	The Snap Data Volume usage rate is displayed as a percentage.

Title	Description
Logical	The Snap Data Volume logical capacity is displayed. The display unit is the number of sectors.
Physical	The Snap Data Volume physical capacity is displayed. The display unit is the number of sectors.
Used	The size of the Snap Data Volume area in use is displayed. The display unit is the number of sectors. The capacity displayed includes both the size used for Copy(*) and the size used for the host(*).
Copy	The size of the Snap Data Volume area in use for Copy(*), out of the total area used, is displayed. The display unit is the number of sectors.
Host	The size of the Snap Data Volume area in use for the host(*), out of the total area used, is displayed. The display unit is the number of sectors.
Pool	The size of the area in use in the Snap Data Pool allocated to the Snap Data Volume is displayed. The display unit is the number of sectors.

*) Area in use for Copy: This is the size used for updates performed while SnapOPC/SnapOPC+ is running.

Area in use for Host: This is the size used for updates performed at the Snap Data Volume while SnapOPC/SnapOPC+ is not running.

Snap Data Pool status display

Specify "poolstat" in the subcommand to display the total capacity of the Snap Data Pool, the size of the used area and the size of the unused area.

a. Format

```
/opt/FJSVswsts/bin/swstdsv poolstat -G Volume-Name
```

b. Description of option

Option	Description
-G	Specifies the ETERNUS disk array volume (Volume-Name). However, if a Snap Data Volume or similar for which a capacity shortage has occurred is specified, the command terminates abnormally. In this case, specify an access path or similar for which a capacity shortage has not occurred. Refer to " Creating an access path " for details of access paths.

c. End status

= 0 : Normal end

> 0 : Abnormal end

d. Example

Specify the access path (/dev/dsk/c39t0d1) to display the status of the Snap Data Pool.

```
# /opt/FJSVswsts/bin/swstdsv poolstat -G /dev/dsk/c39t0d1
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 following table shows the content displayed.

Title	Description
BoxID	Displays the Box ID.
Pool-Type	Displays the Snap Data Pool type. Normal: Normal (not encrypted) Encrypted: Encrypted
Rate	Displays the Snap Data Pool usage rate (%).

Title	Description
Total	Displays the total capacity of the Snap Data Pool. The display unit is the number of sectors.
Used	The size of the Snap Data Pool area in use is displayed. The display unit is the number of sectors. The capacity displayed includes both the size used for Copy(*) and the size used for the host(*).
Copy	The size of the Snap Data Pool area in use for Copy(*), out of the total area used, is displayed. The display unit is the number of sectors.
Host	The size of the Snap Data Pool area in use for the host(*), out of the total area used, is displayed. The display unit is the number of sectors.
Free	Displays the size of the unused are in the Snap Data Pool. The display unit is the number of sectors.

* Area in use for Copy: This is the size used for updates performed while SnapOPC/SnapOPC+ is running.

Area in use for Host: This is the size used for updates performed at the Snap Data Volume while SnapOPC/SnapOPC+ is not 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 information on backup and replication.

10.1.1 General notes

This section provides general information on backup and replication.

Storage Server locale

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

Action is required if there are differences between the language environments (LANG) of the Storage Management Server and the Storage Servers. The following table shows the actions required for various combinations:

Table: Action required for various language environments (LANG)

Storage Management Server	Storage Server	Action required
Windows (SJIS)	Windows (SJIS)	None
	Solaris (EUC)	Install the SJIS package on the Storage Server.
	HP-UX (EUC)	Install the SJIS package on the Storage Server.
	Linux (EUC)	None
	Linux (UTF8)	None
	AIX (SJIS)	None
	AIX (EUC)	Add the SJIS language environment on the Storage Server.
In accordance with the Solaris locale setting (standard: EUC)	Windows (SJIS)	None. (If the Storage Management Server is EUC, Storage Server processing is problem free.)
	Solaris (EUC)	None (if the code is the same). If the Management Server is SJIS, install the SJIS package on the Storage Server.
	HP-UX (EUC)	None (if LANG is the same). If the Management Server is SJIS, install the SJIS package on the Storage Server.
	Linux (EUC)	None
	Linux (UTF8)	None
	AIX (SJIS)	None (if LANG is the same). If the Management Server is EUC, add the EUC language environment.
	AIX (EUC)	None (if LANG is the same). If the Management Server is SJIS, add the SJIS language environment.
Linux (EUC)	Windows (SJIS)	None. (If the Storage Management Server is EUC, Storage Server processing is problem free.)
	Solaris (EUC)	None
	HP-UX (EUC)	None
	Linux (EUC)	None
	Linux (UTF8)	None
	AIX (SJIS)	Add the EUC language environment on the Storage Server.
	AIX (EUC)	None
Linux (UTF8)	Windows (SJIS)	None
	Solaris (EUC)	Add the UTF8 language environment on the Storage Server.
	HP-UX (EUC)	Add the UTF8 language environment on the Storage Server.

Storage Management Server	Storage Server	Action required
	Linux (EUC)	None
	Linux (UTF8)	None
	AIX (SJIS)	Add the UTF8 language environment on the Storage Server.
	AIX (EUC)	Add the UTF8 language environment on the Storage Server.

Backup and Replication Exclusions

The following devices must not be subject to backup or replication:

- The device on which the system is stored
- The device on which AdvancedCopy Manager has been installed
- The device on which the resource files of AdvancedCopy Manager reside

Data integrity of data in a volume

Table: Maintaining data integrity

Data in transaction volumes	Maintenance of data integrity	Operation
File system	An AdvancedCopy Manager command unmounts the file system to maintain data integrity.	See "Pre-processing and Post-processing of Backup and Restoration", and "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.

Fetching Device Information

Before starting the backup or replication, fetch the device information that is managed by all the Storage Servers using the GUI client of AdvancedCopy Manager. The time taken to complete this operation will be proportional to the total number of devices defined on the selected Storage Servers. If many devices are involved, perform this operation while the CPU load or I/O load on the system is low.

As a guide, about 0.5 seconds are required per device (ie, disk) when there is no load on the system

Notes for multipath configurations

When the device has a multipath configuration, a path switch is not carried out automatically, even if one of the two paths is blocked.

When one of the two paths is blocked, please perform the backup or replication again after taking the following steps:

1. Execute the following command on an available device to switch the path:

```
/usr/sbin/pvchange -s /dev/dsk/c*t*d*
```

Be careful, as the system will stop if the commands are executed on a failed device

2. Execute the backup or replication command again.

Mounting of disks from other servers

Before backing up, restoring, or replicating a volume that can be mounted from multiple servers, cancel mounting from the other servers.

For disks that need not be mounted from other servers, prevent multiple servers from detecting or accessing the same logical disk by setting up hardware, such as devices in the Fujitsu ETERNUS storage system and fiber channel switches, accordingly.

Notes on copy operations targeted at file systems

For a copy operation targeted at a file system, unmount the applicable volume to prevent data access and maintain the data integrity.

If backup, restore, or replication are performed without unmounting the volume, an error occurs since a volume in use cannot be unmounted.

Note the following conditions under which unmounting is enabled:

- Another application is not using the volume. If another application is using it, suspend the application temporarily.
- The volume is not in use. If users are accessing the volume, prohibit them from using it temporarily.
- Another volume is not mounted under the volume. If it is mounted, unmount it temporarily.
- The volume is not shared using NFS sharing (for Solaris/HP-UX/Linux). If it is shared, unshare it temporarily.

Unmounting is required only during command execution. Operation can be resumed after command execution.

Notes when copying between different OSs

The replication operation can be used for copying between different OSs.

Use the combinations shown below to copy between different OSs.

Table: Combinations for copying between different OSs

		Replication destination	
		Slice	Logical Unit (disk)
Replication source	Slice	Y	N
	Logical Unit (disk)	Y	N

Y: Possible, N: Not Possible

Notes when using the -m option

The -m option can be specified in the following commands in order to avoid performing communication processing with non-operating servers during inter-server replication:

- Replication start command (swsrpstartsync)
- Replica create command (swsrpmake)
- Replication cancel command (swsrpcancel)

If the -m option is specified, any one of the following conditions must be met in order to avoid problems arising through not performing pre-processing and post-processing at non-operating servers:

1. The non-operating server is shut down.
2. Separate pre-processing and post-processing is executed manually for the volumes on the non-operating server.
3. The target volumes are RAW devices (Oracle, etc., which do not require pre-processing and post-processing), rather than file systems.

Copy processing performance of Advanced Copy

Since the copy processing for Advanced Copy is performed by the disk array unit, please direct enquiries concerning the copy processing performance to your disk array device support department.

Restart while executing synchronous processing

When a copy destination is a volume group and a restart is performed while executing the synchronous processing (EC/REC), the volume will become inactive. When the copy destination becomes inactive while the synchronous processing is being executed, it cannot be activated again. Moreover, the operation such as stopping the synchronous processing cannot be executed. If it is deactivated by mistake, the synchronous processing must be stopped forcibly using ETERNUSmgr/GRmgr and then the copy destination volume is activated.

10.1.2 Notes on the snapshot high-speed backup

This section provides notes on the snapshot high-speed backup.

Number of Required Backup Volumes

Snapshot high-speed backup requires as many backup volumes as the number of specified backup generations. Thus, the snapshot backup is disabled if:

- All the specified number of generations have been backed up, and
- No new volume that can be used as a backup volume exists

Simultaneous Backup

Performing the snapshot high speed backup simultaneously with the same transaction volume specified will perform two backups in parallel.

However, performing the snapshot high-speed backup simultaneously more often than the number of preservation generations will cancel the backups starting from the oldest one. In other words, you cannot perform more simultaneous backups than the specified number of generations.

10.1.3 Notes on synchronous high-speed backup

This section provides notes on synchronous high-speed backup.

Backup processing

The backup command cannot be executed until the transaction and backup volumes are in an equivalency maintenance state.

Backup policy settings

Backup policies can still be set up if the number of registered backup volumes is enough for snapshot type high speed backup, even if the number is not enough for synchronous high speed backup. In this case, synchronized backup will not be possible.

Volume group operation under synchronous processing

If the backup volume is a volume group, do not execute the `vgchange (1M)` command under synchronous processing (EC/REC) execution. If this occurs and the backup volume is inactive then it cannot be re-activated. Additionally, it will not be possible to perform commands including those to stop synchronous processing, execute backup and cancel backup synchronous processing. If accidentally inactivated, use `ETERNUSmgr/GRmgr` or similar to forcibly stop synchronous processing, and then activate the backup volume.

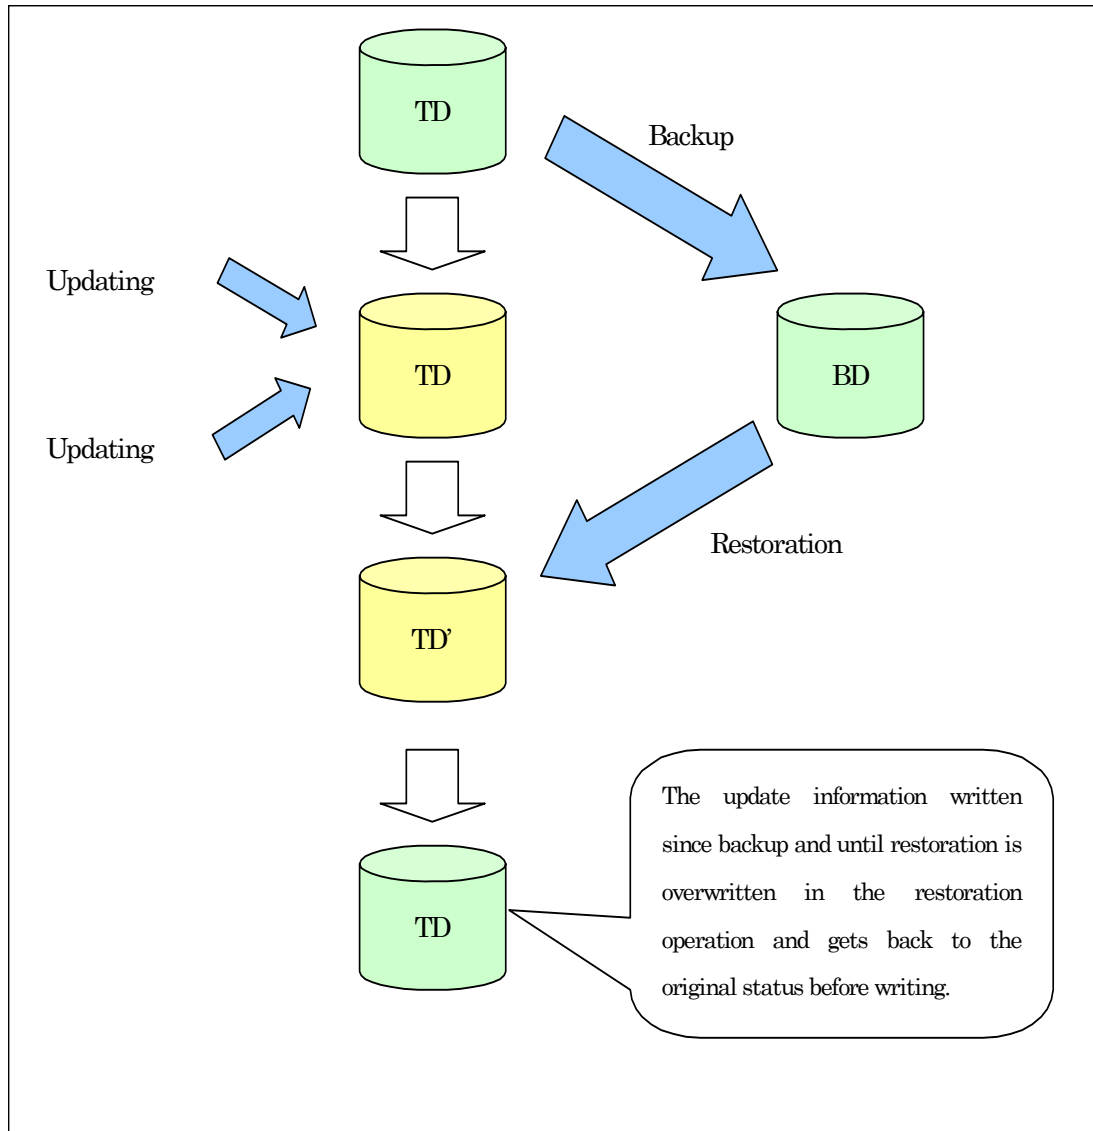
10.1.4 Notes on restoration

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: Notes on restoration of AdvancedCopy Manager



TD: Transaction volume
BD: Backup volume

10.1.5 Notes on synchronous replication

This section provides notes on synchronous replication.

Volume group operation under synchronous processing

If the replica volume is a volume group, do not execute the `vgchange (1M)` command under synchronous processing (EC/REC) execution. If this occurs and the replica volume is inactive then it cannot be reactivated. Additionally, it will not be possible to perform commands including those to stop synchronous processing, execute replication and cancel replication. If accidentally inactivated, use `ETERNUSmgr/GRmgr` or similar to forcibly stop synchronous processing, and then activate the copy destination volume.

10.1.6 Notes on cluster operation

Note the following general points on cluster operation:

- Do not make a `SWSTGNODE` registration that configures an environment variable to the entire system.

- AdvancedCopy Manager does not monitor daemons composing a transaction. If the daemon stops for some reason, AdvancedCopy Manager does not automatically restart it, so you must restart it manually. The stopping of a transaction oriented daemon of AdvancedCopy Manager does not cause the package to fail over.
For details on the daemons composing AdvancedCopy Manager transactions and how to start these daemons, refer to "Starting and Stopping Daemons".

10.1.7 Notes on backups in cluster operation

MC/ServiceGuard

For backups in a cluster 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 the resource adjustment command as a recovery measure.
- You must define a volume group of backup volumes in "VOLUME GROUPS" on the transaction-oriented package control script.
Example: VG[1]=vg02
- In cluster operation, you must perform backups on the active system of the package because AdvancedCopy Manager is placed in the package and operates as part of the package. You cannot make a backup from a standby node or another package.

If the backup volume belongs to a volume group and failover occurs while synchronization (EC/REC) is being handled, package startup is disabled because the backup volume cannot be activated on the standby node.

Therefore, you should perform OPC-based, snapshot type, high speed backups in a cluster environment.

If package startup is disabled during an EC-based, synchronous high speed backup, forcibly stop synchronization by using ETERNUSmgr, GRmgr, or similar, and then start the package.

VERITAS Cluster Server

For backup in a cluster 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 the resource adjustment 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 system of the service group because AdvancedCopy Manager is placed into the service group and operates as part of the service group. You cannot make a backup from a standby node or another service group.

To perform backup/restore using AdvancedCopy Manager, you usually need to unmount transaction volumes beforehand and mount the volumes later.

If the mount points of transaction 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 backup/restore pre-processing and post-processing scripts so that the transaction volumes are set offline or online instead of being unmounted or mounted.

Please note that once the online/offline command is issued there will be a delay before the mount/unmount command can be executed. The "sleep" and "df" commands can be used to delay the mount/unmount until after the process that sets the volumes online or offline has completed

- For details and examples on the pre-processing and post-processing scripts, see "Backup/Restore Pre-processing/Post-processing". For details on the commands used to set resources offline and online, see the relevant cluster system manual.

10.1.7.1 Backup and Restoration during package stoppage

MC/ServiceGuard

When the package to which AdvancedCopy Manager belongs is operating, backup is enabled only on the active node. The standby node environment is not ready to provide the necessary resources, thus disabling backup operation. Likewise, if the package has been stopped, the environment is not ready and backup is disabled.

However, the backup operation is re-enabled only when the package has been stopped, by temporarily arranging the minimum resources.



Note

- If either of the following types of shared disks cannot be enabled (ie, be 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 a backup:
 - Backup/restore using the GUI
 - Manipulation of a backup operation from a Storage Management Server (host specification with the “-h” option)
 - Addition or deletion of a transaction volume/backup volume, or a change in policy
- If display-related commands are executed by Storage Management Server transactions or Storage Server transactions, options must be specified in the commands. Refer to “Backup management commands” for details of backup management display-related commands.

If the package has been stopped, perform a backup by taking the following steps:

1. Confirm that the package has been stopped on both nodes.
For details on how to stop the package, see the MC/ServiceGuard manual.
2. Log in to the node on which you want to make a backup, for example through telnet.
You cannot use the logical IP address. Directly use a node that is accessible 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 only one of the nodes, not both.
4. Perform the backup.
You can perform the backup in the same way as for regular cluster operations.
 5. Release all the shared disks you enabled in Step 3.
Unmount the file systems you mounted.
Stop (set offline) the shared disk.
 6. Start (set online) the package.
Start the package, if necessary.
For details on how to start the package, see the MC/ServiceGuard 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 the necessary resources, thus disabling backup. Likewise, if the service group has been stopped, the environment is not ready and backup is disabled.

However, backup operation is re-enabled only when the service group has been stopped, by temporarily arranging the minimum resources.



Note

- Backup is disabled if either of the following types of shared disks cannot be enabled (ie, set online or mounted in a system):
 - Shared disk for AdvancedCopy Manager shared data
 - Backup operation disk (transaction volume/backup volume to operate)
- The following operations are not allowed in a backup:
 - Backup/restore using the Web GUI
 - Manipulation of the backup operation from a Storage Management Server (host specification with the “-h” option)
 - Addition or deletion of a transaction volume/backup volume, or a change in policy

- In order to execute a display command for backup management to a Storage Management Server transaction that also functions as a Storage Server transaction, you must specify options for the command. For details on the display command for backup management, see "Commands for Backup Management".

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, see the VERITAS Cluster Server manual.
2. Log in to the node on which you want to make a backup, for example through telnet. You cannot use the logical IP address. Directly use a node that is accessible 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. Perform the backup.
You can perform the backup in the same way as for regular operations in a cluster system.
You can use the restore command to perform a restore.
5. Release all the shared disks you enabled in Step 3.
Unmount the file systems you mounted.
Stop (set offline) the shared disk.
6. Start (set online) the service group.
Start the service group, if necessary.
For details on how to start a service group, see the VERITAS Cluster Server manual.

10.1.8 Notes on replication in a cluster system

MC/ServiceGuard

For replication in a cluster 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 the resource adjustment command as a recovery measure.
- You must define a volume group of backup volumes in "VOLUME GROUPS" on the transaction-oriented package control script.
Example: VG[1]=vg02
- You must perform replication on the active system of the package because AdvancedCopy Manager is placed in the package and operates as part of the package.
You cannot perform replication from a standby node or another package.

In addition, if the destination volume belongs to a volume group and failover occurs while synchronization (EC/REC) is being handled, package startup is disabled because the destination volume cannot be activated on the standby node.

Therefore, in a cluster environment you should use OPC based, snapshot type replications. If package startup is disabled during EC-based synchronous replication, forcibly stop the synchronization by using ETERNUSmgr, GRmgr, or similar, and then start the package.

VERITAS Cluster Server

For replication in a cluster 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 the resource adjustment command as a recovery measure.
- When you use synchronous replication, do not register any destination volume to a cluster transaction.

- You must perform replication operations on the active system of the service groups because AdvancedCopy Manager is placed in the service groups and operates as part of the service groups.

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.

Please note that once the online/offline command is issued there will be a delay before the mount/unmount command can be executed. The sleep and df commands can be used to delay the mount/ unmount until after the process that sets the volumes online or offline has completed

For details and examples of the types of commands that require the volumes to be unmounted and mounted, and the pre-processing and post-processing scripts, see "Pre-processing and Post-processing of Replication". For details of the commands used to set resources online and offline, see the relevant cluster system manual.

10.1.8.1 Replication during stoppage of the package

MC/ServiceGuard

When the package to which AdvancedCopy Manager belongs is operating, replication is enabled only on the active node. The standby node environment is not ready to provide the necessary resources, thus disabling replication. Likewise, if the package has been stopped, the environment is not ready and it disables replication.

However, replication is re-enabled only when the package has been stopped, by temporarily arranging the minimum essential resources.



Note

- If either 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 disks (source volume/destination volume to operate)
- The following are not allowed in replication operations:
 - Manipulation for replication from a Storage Management Server (host specification with the "-h" option)
 - Addition or deletion of an original/replica volume
 - Inter-server replication without the "-m" option specified
- Command options must be specified if replication management display commands are executed for Storage Management Server transactions in conjunction with Storage Server transactions. Refer to "Replication management commands" for information concerning the replication management display commands.

Perform replication during stoppage of a package by taking the following steps:

1. Confirm that the package has been stopped on both nodes.
For details on how to stop the package, see the MC/ServiceGuard manual.
2. Log in to the node on which you want to perform replication, for example through telnet.
You cannot use the logical IP address. Directly use a node that is accessible 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

- Be sure to enable the shared disk from either node. Do not enable the shared disk from both nodes.

4. Perform the replication.
You can perform the replication in the same way as for regular cluster operations.
5. Release all the shared disks you enabled in Step 3.
Unmount the file systems you mounted.
Stop (set offline) the shared disk.
6. Start (set online) the package.
Start the package, if necessary.
For details on how to start the package, see the MC/ServiceGuard 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 the necessary resources, thus disabling replication. Likewise, if the service group has been stopped, the environment is not ready and it disables replication.

However, replication is re-enabled only when the service group has been stopped by temporarily arranging the minimum essential resources.



Note

- If either of the following types of shared disks cannot be enabled (ie, set online or mounted in a system), replication is disabled:
 - Shared disk for AdvancedCopy Manager shared data
 - Replication disks (source volume/destination volume)
- The following are not allowed in replication operations:
 - Manipulation for a replication operation from a Storage Management Server (host specification with the “-h” option)
 - Addition or deletion of a source/replica volume
 - Inter-server replication without the “-m” option specified
- To execute a display command for replication management to a Storage Management Server transaction that also functions as a Storage Server transaction, you must specify options for the command. For details on the display command for replication management, see “Commands for Replication Management” in this manual.

Perform replication during stoppage of the service group by following the steps below.

1. Confirm that the service group has been stopped on both nodes.
For details on how to stop a service group, see the VERITAS Cluster Server manual.
2. Log in to the node on which you want to perform replication, for example through telnet.
You cannot use the logical IP address. Directly use a node that is accessible 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 source 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, not from both nodes.

4. Perform replication.
You can perform replication in the same way as for regular cluster operations.
You can also perform a restore from the replica volume to the original volume.
5. Release all the shared disks you enabled in Step 3.
Unmount the file systems you mounted.
Stop (set offline) the shared disk.
6. Start (set online) the service group.
Start the service group, if necessary.

For details on how to start a service group, see the VERITAS Cluster Server manual.

10.1.9 Notes on using ETERNUS VS900 (virtualization switch)

Points to take note of with regard to the copying of virtual disks created in ETERNUS VS900 are explained in this section.

Copy function

EC/REC is the only copy function that can be used in ETERNUS VS900. For this reason, copy virtual disks using the synchronous-type replication of the replication function.

The OPC function is not supported in ETERNUS VS900, therefore backup the function (that uses OPC for restore) cannot be used.

REC operation modes

Operation modes (Transfer mode, Recovery mode, Split mode) that can be used in ETERNUS VS900 are shown in the table below.

Table: Operation modes

Operation mode type	Operation mode that can be used
Transfer mode	Synchronous mode only
Recovery mode	Manual Recovery mode only
Split mode	Automatic Split mode only

Accordingly, the options for the operation mode that can be specified in the replication start command (swsrpstartsync) are different from the ETERNUS disk array.

Table: Transfer mode-related options

Option	ETERNUS disk array	ETERNUS VS900
-i (Consistency mode specified)	Operation is possible	Operation is not possible
-k (Stack mode specified)	Operation is possible	Operation is not possible
-y (Synchronous mode specified)	Operation is possible	Operation is possible (*)
None of -i, -k, or -y options specified	• Operation is possible (in Asynchronous mode)	Operation is not possible

Table: Recovery mode-related options

Option	ETERNUS disk array	ETERNUS VS900
-g (Manual Recovery mode specified)	Operation is possible	Operation is possible (*)
-g option not specified	Operation is possible (in Automatic Recovery mode)	Operation is not possible

Table: Split mode-related options

Option	ETERNUS disk array	ETERNUS VS900
-a (Automatic Split mode specified)	Operation is possible	Operation is possible
-v (Manual Split mode specified)	Operation is possible	Operation is not possible
-a and -v options not specified	Operation is possible (in Automatic Split mode)	Operation is possible (in Automatic Split mode)

(*) In ETERNUS VS900, the Transfer mode and Recovery mode must be specified explicitly. Accordingly, specify the following format to start/restart REC when executing the replication start command (swsrpstartsync). The “-a” option can be omitted.

```
swsrpstartsync - y - a - g From-Volume-Name To-Volume-Name
```

The REC operation modes that can be used in ETERNUS VS900 (transfer mode, Recovery mode, Split mode) are fixed, however, therefore the operation mode cannot be modified using the synchronous processing mode modification command (swsrpchsnc).

EC/REC reversal

EC/REC reversal cannot be executed in ETERNUS VS900.

Displaying/modifying REC transfer buffer information

There is no REC transfer buffer in ETERNUS VS900. For this reason, the following commands cannot be executed for ETERNUS VS900:

- REC transfer buffer information display command (swsrprebuffstat)
- REC transfer buffer settings modification command (swsrprebuffset)

Appendix A Pre-processing and Post-processing of Backup and Restore

This appendix provides information on shell scripts used for backup and restore pre-processing and post-processing in HP-UX AdvancedCopy Manager.

A.1 Overview

Shell scripts used for pre-processing and post-processing are started before and after a backup or restore when the backup or restore command is executed.

These shell scripts describe processing required by AdvancedCopy Manager to back up or restore a transaction volume.

The method to configure the pre/post-processing settings and customize the shell script is explained in this chapter.

A.2 Pre-processing and Post-processing of Backup

Backup using AdvancedCopy Manager must be performed while other processes are prevented from accessing the transaction volume.

Therefore, when a transaction volume is a physical disk, the pre-processing acquires the mount status of a transaction volume using the transaction volume name and then performs the following:

Table: Pre-processing of transaction volume

Transaction volume status	Pre-processing
Mounted	Unmount a transaction volume. (*1)
Unmounted	Take no action.

*1 If the transaction volume cannot be unmounted, customize the shell script for pre-processing. For information on this customization, refer to the section "When you do not want to unmount a transaction volume".

The post-processing that takes place is dependent on what has been done in the pre-processing.

Table: Pre-processing and post-processing of backup volume

Pre-processing	Post-processing
A transaction volume was unmounted.	Remount the transaction volume.
No action taken.	Take no action.

If no transaction volume has been mounted since the start of system operation (for example 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 required processing to a shell script.

When customizing a script, strictly observe the following rules regarding error codes:

Table: Error code rules

Error code	Usage
0-99	Unusable (reserved for AdvancedCopy Manager)
100-255	Usable



Note

For backing up the volume group that includes the logical volume in which the file system has been built, refer to "Backup a volume group". The pre-processing and post-processing scripts have to be customized to execute unmount/mount processing for all logical volumes across which the file system exists.



Note

If post-processing fails, execute the Resource match command (swstsrsemch) because the resource information may be inconsistent.

A.2.1 Backup Pre-processing

The name of the shell script for pre-processing before 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
```

The shell script contents are as follows.

```

1: #!/bin/sh
2:
3: # AdvancedCopy Manager
4: # All Rights Reserved, Copyright FUJITSU LIMITED, 2002-2007
5:
6: #
7: #   Preprocessing of backup processing
8: #
9: #       Argument: $1 Device or VG name of transaction disk
10: #           $2 Mount point of transaction disk
11: #           $3 Device or VG name of backup disk
12: #
13: #       Error number
14: #           2: Argument error
15: #           10: umount error
16: #           30: VG configuration file not found error
17: #           99: Script not customize
18:
19:
20: # Argument check
21: case $# in
22: 1)
23:     ;;
24: 2)
25:     ;;
26: 3)
27:     ;;
28: *)
29:     exit 2
30:     ;;
31: esac
32:
33: device=$1
34: mount_point=$2
35: bk_device=$3
36:
37: # Determination postprocessing file name
38: if [ "$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: # Device type check
48: trans=`echo $device | /usr/bin/grep "/dev/dsk/"`
49: lvmtrans=""`echo $device | /usr/bin/grep "/dev/"`
50: vxpvtrans=""`echo $device | /usr/bin/grep "/dev/vx/dmp/"`
51: if [ "$trans" != "" ]
52: then
53:     dev_type="physical"
54:     dev=""`echo $device | /usr/bin/sed "s/¥/dev¥/dsk¥//"`
55:     post_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".pre"
56:
57:     fstype_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".fstype"
58:     bd_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".bd"
59: elif [ "$vxpvtrans" != "" ]
60: then
61:     dev_type="vxvm_pv"
62:     # /dev/vx/dmp/XXXX -> XXXX

```



```

62:         dev="`echo $device | /usr/bin/awk -F/ ' { print $5 }'`"
63:         post_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".pre"
64:
fstype_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".fstype"
65:         bd_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".bd"
66:     elif [ "$lvmtrans" != "" ]
67:     then
68:         dev_type="logical"
69:         # /dev/XXXX -> XXXX
70:         dev="`echo $device | /usr/bin/awk -F/ ' { print $3 }'`"
71:         post_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".pre"
72:
fstype_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".fstype"
73:         bd_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".bd"
74:     else
75:         exit 0
76:     fi
77:
78: # Save VG name of backup disk
79: bk_trans="`echo $bk_device | /usr/bin/grep "/dev/dsk/"`"
80: bk_lvmtrans="`echo $bk_device | /usr/bin/grep "/dev/"`"
81: bk_vxpvtrans="`echo $bk_device | /usr/bin/grep "/dev/vx/dmp/"`"
82: if [ "$bk_trans" = "" -a "$bk_lvmtrans" != "" -a "$bk_vxpvtrans" = "" ]
83: then
84:     bk_vg_name="`echo $bk_device | /usr/bin/awk -F/ ' { print $3 }'`"
85:     cfg_file="/etc/lvmconf/"$bk_vg_name".conf"
86:     if [ ! -r $cfg_file ]
87:     then
88:         exit 30
89:     fi
90:     echo $bk_device > $bd_file
91: fi
92:
93: # When the transaction disk is a physical device #####
94: if [ "$dev_type" = "physical" ]
95: then
96:
97:     if [ "$mount_point" != "" ]
98:     then
99: # When device cannot be unmounted --- 1
100: #
101: #         if [ "$device" = "/dev/dsk/cXtXdX" ]
102: #         then
103: #             if [ "$bk_device" != "" ]
104: #             then
105: #                 echo $bk_device > $bd_file
106: #             fi
107: #             df -ln $mount_point 2>/dev/null | awk -F: ' { print $2 }'
| awk ' { print $1 }' > $fstype_file
108: #             sync
109: #             sync
110: #             echo "fsck" > $post_file
111: #
112: # When device can be unmounted --- 2
113: #             else
114: #                 /usr/bin/df -ln $mount_point 2>/dev/null | /usr/bin/awk
-F: ' { print $2 }' | /usr/bin/awk ' { print $1 }' > $fstype_file
115: #                 /usr/sbin/umount $mount_point 2>/dev/null
116: #                 if [ $? != 0 ]
117: #                 then
118: #
119:
retry_count=3
119:                 sleep_time=1

```

```

120:                                     result_flag=1
121:
122:                                     while [ $retry_count -gt 0 ]
123:                                     do
124:                                         /usr/sbin/umount $mount_point >
125: $err_log_path/$dev.umount 2>&1
126:                                     if [ $? != 0 ]
127:                                     then
128:                                         retry_count=`expr $retry_count - 1`
129:                                         /usr/bin/sleep $sleep_time
130:                                     else
131:                                         result_flag=0
132:                                         break
133:                                     fi
134:                                     done
135:
136:                                     if [ $result_flag != 0 ]
137:                                     then
138:                                         /usr/sbin/fuser -cu $mount_point>
139: $err_log_path/$dev.fuser 2>&1
140:                                         /usr/bin/ps -ef > $err_log_path/$dev.ps 2>&1
141:                                     exit 10
142:                                     else
143:                                         /usr/bin/rm -f $err_log_path/$dev.umount
144:                                         fi
145:                                     fi
146:                                     echo "mount" > $post_file
147: #                                     fi
148: # When device was not mounted --- 3
149: #
150:                                     else
151:                                         echo "none" > $post_file
152:                                     fi
153: # When the transaction disk is a volume group #####
154: elif [ "$dev_type" = "logical" ]
155: then
156:
157: # Devices is volume group and script not customize
158:     exit 99
159:
160: # When devices of volume group cannot be unmounted --- 4
161: #
162: #     # Specify the name of volume group not to unmount
163: #     if [ "$device" = "/dev/vgXX" ]
164: #     then
165: #         if [ "$bk_device" != "" ]
166: #         then
167: #             echo $bk_device > $bd_file
168: #         fi
169: #         sync
170: #         sync
171: #         echo "fsck" > $post_file
172: #     fi
173: #
174: # When devices of volume group can be unmounted --- 5
175: #
176: #     # Specify the name of volume group to unmount
177: #     if [ "$device" = "/dev/vgXX" ]
178: #     then
179: #         # Unmount all logical volumes of the volume group
180: #         mount_point="/XX"

```

```

181: #           /usr/sbin/umount $mount_point 2>/dev/null
182: #           if [ $? != 0 ]
183: #           then
184: #               retry_count=3
185: #               sleep_time=1
186: #               result_flag=1
187: #
188: #               while [ $retry_count -gt 0 ]
189: #               do
190: #
191: #                   # $dev is volume group name
192: #
193: #                   /usr/sbin/umount $mount_point >
194: #                   $err_log_path/$dev.umount 2>&1
195: #                   if [ $? != 0 ]
196: #                   then
197: #                       retry_count=`expr $retry_count - 1`
198: #                       sleep $sleep_time
199: #                   else
200: #                       rm -f $err_log_path/$dev.umount
201: #                       result_flag=0
202: #                       break
203: #                   fi
204: #               done
205: #
206: #               if [ $result_flag != 0 ]
207: #               then
208: #                   /usr/sbin/fuser -cu $mount_point>
209: #                   $err_log_path/$dev.fuser 2>&1
210: #                   /usr/bin/ps -ef > $err_log_path/$dev.ps 2>&1
211: #                   exit 10
212: #               fi
213: #               echo "mount" > $post_file
214: #           fi
215: #
216: # When devices of volume group was not mounted --- 6
217: #
218: # Specify the name of volume group to do nothing
219: # if [ "$device" = "/dev/vgXX" ]
220: # then
221: #     echo "none" > $post_file
222: # fi
223: #
224: # When the transaction disk is a VxVM physical volume --- 7 #####
225: elif [ "$dev_type" = "vxvm_pv" ]
226: then
227:     # Nothing is done to VxVM PV.
228:     echo "none" > $post_file
229: fi
230:
231: exit 0

```

A. 2. 2 Backup Post-processing

The name of the shell script for the post-processing of a backup is as follows.

- In the case of non-cluster operation

```
/etc/opt/FJSVswsts/sh/0pcBackup.post
```

- In the case of cluster operation

```
/etc/opt/FJSVswsts/<logic node name>/sh/0pcBackup.post
```

The shell script contents are as follows.

```

1: #!/bin/sh
2:
3: # AdvancedCopy Manager
4: # All Rights Reserved, Copyright FUJITSU LIMITED, 2002-2007
5:
6: #
7: #   Postprocessing of backup processing
8: #
9: #           Argument: $1 Device or VG name of transaction disk
10: #                $2 Mount point of transaction disk
11: #
12: #   Error number
13: #           2: Argument error
14: #           11: mount error
15: #           31: vgcfgrestore error
16: #           12: fsck error
17:
18: # Argument check
19: case $# in
20: 1)
21:     ;;
22: 2)
23:     ;;
24: *)
25:     exit 2
26:     ;;
27: esac
28:
29: device=$1
30: mount_point=$2
31:
32: # Determination of postprocessing file name
33:
34: if [ "$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: # Device type check
44: trans="`echo $device | /usr/bin/grep "/dev/dsk/"`"
45: lvmtrans="`echo $device | /usr/bin/grep "/dev/"`"
46: vxpvtrans="`echo $device | /usr/bin/grep "/dev/vx/dmp/"`"
47: if [ "$trans" != "" ]
48: then
49:     dev_type="physical"
50:     dev="`echo $device | /usr/bin/sed "s/¥/dev¥/dsk¥//"/`"
51:     post_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".pre"
52:
53:     fstype_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".fstype"
54:     bd_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".bd"
55: elif [ "$vxpvtrans" != "" ]
56: then
57:     dev_type="vxvm_pv"
58:     # /dev/vx/dmp/XXXX -> XXXX
59:     dev="`echo $device | /usr/bin/awk -F¥/ ' { print $5 }'`"
60:     post_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".pre"
61:
62:     fstype_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".fstype"

```

```

61:         bd_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".bd"
62:     elif [ "$lvmtrans" != "" ]
63:     then
64:         dev_type="logical"
65:         # /dev/vgXX -> vgXX
66:         dev="`echo $device | /usr/bin/awk -F/ ' { print $3 } '`"
67:         post_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".pre"
68:
fstype_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".fstype"
69:         bd_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".bd"
70:     else
71:         exit 0
72:     fi
73:
74: # Confirmation of postprocessing
75: if [ ! -r $post_file ]
76: then
77:     exit 0
78: fi
79: post="`/usr/bin/cat $post_file`"
80:
81: # Confirmation of FStype
82: if [ ! -r $fstype_file ]
83: then
84:     fs=""
85: else
86:     fs="`/usr/bin/cat $fstype_file`"
87: fi
88:
89: # Restore VG Configuration
90: if [ -r $bd_file ]
91: then
92:     bk_device="`/usr/bin/cat $bd_file`"
93:     bk_trans="`echo $bk_device | /usr/bin/grep "/dev/dsk/"`"
94:     bk_lvmtrans="`echo $bk_device | /usr/bin/grep "/dev/"`"
95:     bk_vxpvtrans="`echo $bk_device | /usr/bin/grep "/dev/vx/dmp/"`"
96:     if [ "$bk_trans" = "" -a "$bk_lvmtrans" != "" -a "$bk_vxpvtrans" = "" ]
97:     then
98:         bk_vg_name="`echo $bk_device | /usr/bin/awk -F/ ' { print $3 } '`"
99:         cfg_file="/etc/lvmconf/"$bk_vg_name".conf"
100:        if [ ! -r $cfg_file ]
101:        then
102:            exit 31
103:        fi
104:        bk_pv_rname="`/usr/sbin/vgcfgrestore -n $bk_device -l 2>
/dev/null | /usr/bin/awk ' {print $1}' | /usr/bin/grep ¥"/dev/rds¥`"
105:        if [ "$bk_pv_rname" = "" ]
106:        then
107:            echo "NULL: /usr/sbin/vgcfgrestore -n $bk_device -l 2>
/dev/null | /usr/bin/awk ' {print $1}' | /usr/bin/grep ¥"/dev/rds¥" >
$err_log_path/$dev.vgcfgrestore
108:            exit 31
109:        fi
110:        /usr/sbin/vgchange -a n $bk_device > /dev/null 2>&1
111:        reststatus=0
112:        for bk_pv_rname1 in $bk_pv_rname
113:        do
114:            /usr/sbin/vgcfgrestore -n $bk_device $bk_pv_rname1 >
/dev/null 2>&1
115:            if [ $? != 0 ]
116:            then
117:                /usr/sbin/vgcfgrestore -n $bk_device $bk_pv_rname1 >>
$err_log_path/$dev.vgcfgrestore 2>&1

```

```

118:             fi
119:             if [ $? = 0 ]
120:             then
121:                 /usr/bin/rm -f $err_log_path/$dev.vgcfgrestore
122:                 reststatus=1
123:                 break
124:             fi
125:         done
126:         if [ "$SWSTGNODE" != "" ]
127:         then
128:             /usr/sbin/vgchange -c y $bk_device > /dev/null 2>&1
129:             /usr/sbin/vgchange -a e $bk_device > /dev/null 2>&1
130:         else
131:             /usr/sbin/vgchange -a y $bk_device > /dev/null 2>&1
132:         fi
133:         if [ $reststatus = 0 ]
134:         then
135:             exit 31
136:         fi
137:     fi
138: fi
139:
140: # No processing
141: # When device was not mounted --- 3
142: # When devices of volume group was not mounted --- 6
143: # When the transaction disk is a VxVM physical volume --- 7
144: if [ "$post" = "none" ]
145: then
146:     /usr/bin/rm -rf $post_file 2> /dev/null
147:     /usr/bin/rm -rf $fstype_file 2> /dev/null
148:     /usr/bin/rm -rf $bd_file 2> /dev/null
149:     exit 0
150: fi
151:
152: # mount processing
153: if [ "$post" = "mount" ]
154: then
155:
156: # When device can be unmounted --- 2
157:     if [ "$dev_type" = "physical" ]
158:     then
159:         /usr/bin/df -l $device > /dev/null 2>&1
160:         if [ $? != 0 ]
161:         then
162:             if [ ! -r $fstype_file ]
163:             then
164:                 /usr/sbin/mount $device $mount_point 2>
165: /dev/null
166:             else
167:                 if [ "$fs" = "" ]
168:                 then
169:                     /usr/sbin/mount $device $mount_point 2>
170: /dev/null
171:                 else
172:                     /usr/sbin/mount -F $fs $device
173: $mount_point 2> /dev/null
174:                 fi
175:             fi
176:             if [ $? != 0 ]
177:             then
178:                 retry_count=3
179:                 sleep_time=1
180:                 result_flag=1

```

```

178:
179:         while [ $retry_count -gt 0 ]
180:         do
181:             if [ ! -r $fstype_file ]
182:             then
183:                 /usr/sbin/mount $device $mount_point >
184: $err_log_path/$dev.mount 2>&1
185:             else
186:                 if [ "$fs" = "" ]
187:                 then
188:                     /usr/sbin/mount $device
189: $mount_point > $err_log_path/$dev.mount 2>&1
190:                 else
191:                     /usr/sbin/mount -F $fs $device
192: $mount_point > $err_log_path/$dev.mount 2>&1
193:                 fi
194:             fi
195:             if [ $? != 0 ]
196:             then
197:                 retry_count=`expr $retry_count - 1`
198:                 /usr/bin/sleep $sleep_time
199:             else
200:                 /usr/bin/rm -f $err_log_path/$dev.mount
201:                 result_flag=0
202:                 break
203:             fi
204:         done
205:
206:         if [ $result_flag != 0 ]
207:         then
208:             exit 11
209:         fi
210:     fi
211:     elif [ "$dev_type" = "logical" ]
212:     then
213:         # When devices of volume group can be unmounted --- 5
214:         # Specify the name of volume group to mount
215:         if [ "$device" = "/dev/vgXX" ]
216:         then
217:             # Mount all logical volumes of the volume group
218:             fs="hfs"
219:             lvname="/dev/vgXX/XXXXX"
220:             lv_mount_point="/XX"
221:             /usr/bin/df -l $lvname > /dev/null 2>&1
222:             if [ $? != 0 ]
223:             then
224:                 /usr/sbin/mount -F $fs $lvname $lv_mount_point >
225: /dev/null
226:             if [ $? != 0 ]
227:             then
228:                 retry_count=3
229:                 sleep_time=1
230:                 result_flag=1
231:                 while [ $retry_count -gt 0 ]
232:                 do
233:                     /usr/sbin/mount -F $fs $lvname
234: $lv_mount_point > $err_log_path/$dev.mount 2>&1
235:                 if [ $? != 0 ]
236:                 then

```



```

236: #                                retry_count=`expr $retry_count
- 1`
237: #                                /usr/bin/sleep $sleep_time
238: #                                else
239: #                                rm -f $err_log_path/$dev.mount
240: #                                result_flag=0
241: #                                break
242: #                                fi
243: #                                done
244: #
245: #                                if [ $result_flag != 0 ]
246: #                                then
247: #                                    exit 11
248: #                                fi
249: #                                fi
250: #                                fi
251: #                                fi
252:
253:                                fi
254:                                /usr/bin/rm -rf $post_file 2> /dev/null
255:                                /usr/bin/rm -rf $fstype_file 2> /dev/null
256:                                /usr/bin/rm -rf $bd_file 2> /dev/null
257:                                exit 0
258: fi
259:
260: # fsck processing
261: # When device cannot be unmounted          --- 1
262: # When devices of volume group cannot be unmounted --- 4
263: if [ "$post" = "fsck" ]
264: then
265:     if [ -r $bd_file ]
266:     then
267:         bk_device=`/usr/bin/cat $bd_file`
268:         bk_trans=`echo $bk_device | /usr/bin/grep "/dev/dsk/"`
269:         if [ "$bk_trans" != "" ]
270:         then
271:             if [ ! -r $fstype_file ]
272:             then
273:                 /usr/sbin/fsck -y $bk_device > /dev/null 2>&1
274:             else
275:                 if [ "$fs" = "" ]
276:                 then
277:                     /usr/sbin/fsck -y $bk_device > /dev/null
278:                 else
279:                     /usr/sbin/fsck -y -F $fs $bk_device >
280: /dev/null 2>&1
281:                 fi
282:             fi
283:             if [ $? != 0 ]
284:             then
285:                 if [ "$fs" = "" ]
286:                 then
287:                     /usr/sbin/fsck -y $bk_device >
288: $err_log_path/$dev.fsck 2>&1
289:                 else
290:                     /usr/sbin/fsck -y -F $fs $bk_device >
291: $err_log_path/$dev.fsck 2>&1
292:                 fi
293:             fi
294:             if [ $? != 0 ]
295:             then
296:                 exit 12
297:             else

```

```

294:                                     /usr/bin/rm -f $err_log_path/$dev.fsck
295:                                     fi
296:                                     fi
297: #                                     else
298: #                                     # Specify the name of volume group to fsck
299: #                                     if [ "$bk_device" = "/dev/vgXX" ]
300: #                                     then
301: #                                     fs="hfs"
302: #                                     lvname="/dev/vgXX/rXXXXX"
303: #
304: #                                     fsck -F $fs -y $lvname > /dev/null 2>&1
305: #                                     if [ $? != 0 ]
306: #                                     then
307: #                                     if [ "$fs" = "" ]
308: #                                     then
309: #                                     fsck -y $lvname >
310: $err_log_path/$dev.fsck 2>&1
311: #                                     else
312: #                                     fsck -y -F $fs $lvname >
313: $err_log_path/$dev.fsck 2>&1
314: #                                     fi
315: #                                     if [ $? != 0 ]
316: #                                     then
317: #                                     exit 12
318: #                                     else
319: #                                     rm -f $err_log_path/$dev.fsck
320: #                                     fi
321: #                                     fi
322: #                                     fi
323: #                                     /usr/bin/rm -rf $post_file 2> /dev/null
324: #                                     /usr/bin/rm -rf $fstype_file 2> /dev/null
325: #                                     /usr/bin/rm -rf $bd_file 2> /dev/null
326: #                                     exit 0
327: fi
328:
329: exit 0

```

A.2.3 Customizing the shell script

The method to customize the shell script for restore pre/post-processing is explained below.

When you do not want to unmount the transaction volume

If you do not want to unmount a mounted transaction volume, eliminate the comment characters (“#”) in lines 101-110, 113 and 146 (in the case of a physical volume) or lines 163-172 (in the case of a volume group), and point to the target device in line 101 or line 163.

In addition, do not perform an update operation on a transaction volume file system after performing the sync command on line 108 or line 169 until backup post-processing. If a file system update is performed in the meantime, the backup is taken with a corrupted file system. In this case the fsck command executed by the post-processing may raise an error.

Backup a volume group

When you are backing up the volume group containing the logical volume used to build the file system, insert a comment character (“#”) into line 158 of the pre-processing script (OpBackup.pre) to remove the exit statement.

- Please remove the comment (“#”) from lines 177-214 of the pre-processing script (OpBackup.pre), as well as from lines 209, 210 and 214-251 of the post-processing script (OpBackup.post). Correct the description in a comment so that unmount/mount processing

is carried out on all the logical volumes across which the file system exists. Change the parts shown as "X" to match the actual environment. If more than one logical volume is being processed, repeat lines 180-212 of the pre-processing script (OpcBackup.exe) and lines 218-250 of the post-processing script (PpcBackup.post) once for each additional logical volume.

- If the operating volume was not unmounted in the pre-processing script then the fsck command will need to be executed (on all the logical volumes contained in the backup volume) in post-processing. Remove the comment ("#") in lines 297-320 of the post-processing script (OpcBackup.post), and correct the description in a comment so that fsck processing is carried out on all the logical volumes across which the file system exists. Change the parts shown as "X" to match the actual environment. If more than one logical volume is being processed, repeat lines 309-319 once for each additional logical volume.



Note

Note that unmount/mount processing to a volume group is not carried out if the scripts have not been customized.



Note

When the backup volume is a volume group, in order to enable use of the backup volume as LVM, volume group configuration information is restored using the `vgcfgrestore` command in the post-processing. The information is restored from the standard backup file `"/etc/lvmconf/volume-group-name.conf"`. If volume group configuration information is backed up by another file, please customize the script.

In cluster operation, the respective nodes comprising the cluster must have volume group configuration information.



Note

Correct the parameters and options, of the OS commands, such as the `mount` command and the `fsck` command, appropriately for their use.



Note For shared-mode volume groups:

When reconfiguring shared-mode volume groups, comment out line numbers 90-138, 144, 145 and 150 in the post-processing script (OpcBackup.post) to prevent the volume group from being reconfigured by the script.

Reconfigure the volume group by manually performing the following steps after the backup has been executed.

1. Stop the volume group (on all nodes involved in the transaction).

```
# /usr/sbin/vgchange -a n <vg_name>
#
```

2. Restore volume group configuration information on the node where the volume group was created.

```
# /usr/sbin/vgcfgrestore -n <vg_name> <pv_path>
#
```

3. Mark the volume group as shareable on the node where the volume group was created.

```
# /usr/sbin/vgchange -S y -c y <vg_name>

#
```

4. Start the volume group on all nodes involved in the transaction.

```
# /usr/sbin/vgchange -a s <vg_name>

#
```

File system name

Change the file system names on lines 217 and 301 of the post-processing shell script (OpcBackup.post) that is executed when backup is performed as appropriate for the operation.

Cluster operation by VERITAS Cluster Server

Customization of the pre-processing and post-processing scripts is required when performing cluster operations using VERITAS Cluster Server, and when the mount point of the transaction volume is registered in the cluster transaction.

Change unmount/mount processing in the pre-processing and post-processing scripts into offline/online processing of the mount point resource.

Please note that once the online/offline command is issued there will be a delay before the mount/unmount command can be executed. The sleep and df commands can be used to delay the mount/unmount until after the process that sets the volumes online or offline has completed

The examples below show how the pre-processing and post-processing scripts are modified.

[Example] Changing the unmount processing of the backup pre-processing script (OpcBackup.pre)0
[lines 115, 124, 181, 193rd]

(Before change)	/usr/sbin/umount \$mount_point
(after change)	/opt/VRTSvcs/bin/hares -offline resource-name -sys system-name

[Example] Adding unmount wait processing to the backup pre-processing script (OpcBackup.pre)
[line 145, 213rd]

(addition)	while /usr/bin/df -l "\$device">/dev/null 2>&1; do ;; done
------------	--

[Example] Changing the mount processing of the backup post-processing script (OpcBackup.post)
[lines 164, 168, 170, 183, 187, 189, 224, 233rd]

(Before change)	/usr/sbin/mount \$device \$mount_point /usr/sbin/mount -F \$fs \$device \$mount_point /usr/sbin/mount -F \$fs \$lvname \$lv_mount_point
(after change)	/opt/VRTSvcs/bin/hares -online resource-name -sys system-name

[Example] Adding mount wait processing to the backup post-processing script (OpcBackup.post)
[lines 209, 252nd]

(addition)	while ! /usr/bin/df -l "\$device">/dev/null 2>&1; do ;; done
------------	--

A.3 Restore Pre-processing and Post-processing

In AdvancedCopy Manager, restore processing must be performed while transaction volumes are unmounted.

Thus, when the transaction volume is a physical disk, the mount state of the transaction volume is determined using its name and the following is then performed in pre-processing:

Table: Pre-processing of transaction volume

Transaction volume status	Pre-processing
Mounted	Unmount the transaction volume.
Unmounted	Take no action.

What to do in the post-processing must be determined depending on what has been done in the pre-processing.

Table: Restoration pre-processing and post-processing

Pre-processing	Post-processing
A transaction volume was unmounted.	Remount the transaction volume.
No action taken.	Take no action.

If special pre-processing or post-processing is required, you need to add it to a shell script. When customizing a script, strictly observe the following rules regarding error codes:

Table: Error code rules

Error code	Usage
0-99	Unusable (reserved for AdvancedCopy Manager)
100-255	Usable



Note

When restoring a file system in a volume group, refer to "Restore a volume group", and customize the pre-processing and post-processing scripts so that the unmount/mount command is carried out for all logical volumes across which a file system exists.



Note

If post-processing fails, execute the Resource match command (swstsrsemch) because the resource information may be inconsistent.

A.3.1 Restore Pre-processing

The name of the shell script for pre-processing before a restore 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
```

The shell script contents are as follows.

```

1: #!/bin/sh
2:
3: # AdvancedCopy Manager
4: # All Rights Reserved, Copyright FUJITSU LIMITED, 2002-2007
5:
6: #
7: #   Preprocessing of restoration processing
8: #
9: #           Argument: $1 Device or VG name of transaction disk
10: #                $2 Mount point of transaction disk
11: #
12: #   Error number
13: #           2: Argument error (system error)
14: #           10: umount error
15: #           30: VG configuration file not found error
16: #           99: Script not customize
17:
18:
19: # Argument check
20: case $# in
21: 1)
22:     ;;
23: 2)
24:     ;;
25: *)
26:     exit 2
27:     ;;
28: esac
29:
30: device=$1
31: mount_point=$2
32:
33: # Determination of postprocessing file name
34:
35: if [ "$SWSTGNODE" != "" ]
36: then
37:     swstg_node="/`echo $SWSTGNODE`"
38: else
39:     swstg_node=""
40: fi
41:
42: err_log_path="/var/opt/FJSVswsts"$swstg_node"/log"
43:
44: # Device type check
45: trans=`echo $device | /usr/bin/grep "/dev/dsk/"`
46: lvmtrans=`echo $device | /usr/bin/grep "/dev/"`
47: vxpvtrans=`echo $device | /usr/bin/grep "/dev/vx/dmp/"`
48: if [ "$trans" != "" ]
49: then
50:     dev_type="physical"
51:     dev=`echo $device | /usr/bin/sed "s/¥/dev¥/dsk¥//"`
52:     post_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".pre"
53:
54: fstype_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".fstype"
55: elif [ "$vxpvtrans" != "" ]
56: then
57:     dev_type="vxvm_pv"
58:     # /dev/vx/dmp/XXXX -> XXXX
59:     dev=`echo $device | /usr/bin/awk -F/ ' { print $5 } '`
60:     post_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".pre"
61:
62: fstype_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".fstype"

```

```

61: elif [ "$lvmtrans" != "" ]
62: then
63:     dev_type="logical"
64:     # /dev/vgXX -> vgXX
65:     dev="`echo $device | /usr/bin/awk -F/ '{ print $3 }`"
66:     post_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".pre"
67:
fstype_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".fstype"
68: else
69:     exit 0
70: fi
71:
72: # When the transaction disk is a physical device #####
73: if [ "$dev_type" = "physical" ]
74: then
75:
76: # When device was mounted --- 1
77:     if [ "$mount_point" != "" ]
78:     then
79:         /usr/bin/df -ln $mount_point 2>/dev/null | /usr/bin/awk -F:
' { print $2 }' | /usr/bin/awk '{ print $1 }' > $fstype_file
80:         /usr/sbin/umount $mount_point 2> /dev/null
81:         if [ $? != 0 ]
82:         then
83:             retry_count=3
84:             sleep_time=1
85:             result_flag=1
86:
87:             while [ $retry_count -gt 0 ]
88:             do
89:                 /usr/sbin/umount $mount_point >
$err_log_path/$dev.umount 2>&1
90:                 if [ $? != 0 ]
91:                 then
92:                     retry_count=`expr $retry_count - 1`
93:                     /usr/bin/sleep $sleep_time
94:                 else
95:                     /usr/bin/rm -f $err_log_path/$dev.umount
96:                     result_flag=0
97:                     break
98:                 fi
99:             done
100:
101:             if [ $result_flag != 0 ]
102:             then
103:                 /usr/sbin/fuser -cu $mount_point>
$err_log_path/$dev.fuser 2>&1
104:                 /usr/bin/ps -ef > $err_log_path/$dev.ps 2>&1
105:                 exit 10
106:             fi
107:         fi
108:         echo "mount" > $post_file
109:
110: # When device was not mounted --- 2
111:     else
112:         echo "none" > $post_file
113:     fi
114:
115: # When the transaction disk is a volume group #####
116: elif [ "$dev_type" = "logical" ]
117: then

```

```

120:
121:     # Check VG configuration file exists
122:     vg_name=`echo $device | /usr/bin/awk -F/ ' { print $3 } '`
123:     cfg_file="/etc/lvmconf/"$vg_name".conf"
124:     if [ ! -r $cfg_file ]
125:     then
126:         exit 30
127:     fi
128:
129: # Devices is volume group and script not customize
130:     exit 99
131:
132: # When devices of volume group was mounted --- 3
133: #     # Specify the name of volume group to unmount
134: #     if [ "$device" = "/dev/vgXX" ]
135: #     then
136: #
137: #         # Unmount all logical volumes of the volume group
138: #         mount_point="/XX"
139: #         /usr/sbin/umount $mount_point 2>/dev/null
140: #         if [ $? != 0 ]
141: #         then
142: #             retry_count=3
143: #             sleep_time=1
144: #             result_flag=1
145: #
146: #             while [ $retry_count -gt 0 ]
147: #             do
148: #                 /usr/sbin/umount $mount_point >
149: # $err_log_path/$dev.umount 2>&1
150: #                 if [ $? != 0 ]
151: #                 then
152: #                     retry_count=`expr $retry_count - 1`
153: #                     sleep $sleep_time
154: #                 else
155: #                     rm -f $err_log_path/$dev.umount
156: #                     result_flag=0
157: #                     break
158: #                 fi
159: #             done
160: #
161: #             if [ $result_flag != 0 ]
162: #             then
163: #                 /usr/sbin/fuser -cu $mount_point>
164: # $err_log_path/$dev.fuser 2>&1
165: #                 /usr/bin/ps -ef > $err_log_path/$dev.ps 2>&1
166: #                 exit 10
167: #             fi
168: #         fi
169: #         echo "mount" > $post_file
170: #     fi
171:
172: # When devices of volume group was not mounted --- 4
173: #     # Specify the name of volume group to do nothing
174: #     if [ "$device" = "/dev/vgXX" ]
175: #     then
176: #         echo "none" > $post_file
177: #     fi
178: #
179: # When the transaction disk is a VxVM physical volume --- 5 #####
180: elif [ "$dev_type" = "vxvm_pv" ]

```



```
181: then
182:     # Nothing is done to VxVM PV.
183:     #
184:     echo "none" > $post_file
185: fi
186: 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 Restore Post-processing

The name of the shell script for post-processing after a restore 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
```

The shell script contents are as follows.

```

1: #!/bin/sh
2:
3: # AdvancedCopy Manager
4: # All Rights Reserved, Copyright FUJITSU LIMITED, 2002-2007
5:
6: #
7: #   Postprocessing of restoration processing
8: #
9: #           Argument: $1 Device or VG name of transaction disk
10: #                   $2 Mount point of transaction disk
11: #
12: #           Error number
13: #                   2: Argument error
14: #                   11: mount error
15: #                   31: vfcfgrestore error
16: #                   12: fsck error
17:
18: # Argument check
19: case $# in
20: 1)
21:     ;;
22: 2)
23:     ;;
24: *)
25:     exit 2
26:     ;;
27: esac
28:
29: device=$1
30: mount_point=$2
31:
32: # Determination of postprocessing file name
33:
34: if [ "$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: trans="`echo $device | /usr/bin/grep "/dev/dsk/"`"
44: lvmtrans="`echo $device | /usr/bin/grep "/dev/"`"
45: vxpvtrans="`echo $device | /usr/bin/grep "/dev/vx/dmp/"`"
46: if [ "$trans" != "" ]
47: then
48:     dev_type="physical"
49:     dev="`echo $device | /usr/bin/sed "s/¥/dev¥/dsk¥//"`"
50:     post_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".pre"
51:
52: fstype_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".fstype"
53: elif [ "$vxpvtrans" != "" ]
54: then
55:     dev_type="vxvm_pv"
56:     # /dev/vx/dmp/XXXX -> XXXX
57:     dev="`echo $device | /usr/bin/awk -F¥/ ' { print $5 }`"
58:     post_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".pre"
59:
60: fstype_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".fstype"
61: elif [ "$lvmtrans" != "" ]
62: then
63:     dev_type="logical"

```

```

62:         # /dev/vgXX -> vgXX
63:         dev="`echo $device | /usr/bin/awk -F/ ' { print $3 }'`"
64:         post_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".pre"
65:         fstype_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".fstype"
66:     else
67:         exit 0
68:     fi
69:
70: # Confirmation of postprocessing
71: if [ ! -r $post_file ]
72: then
73:     exit 0
74: fi
75: post="`/usr/bin/cat $post_file`"
76:
77: # Confirmation of FStype
78: if [ ! -r $fstype_file ]
79: then
80:     fs=""
81: else
82:     fs="`/usr/bin/cat $fstype_file`"
83: fi
84:
85: # Restore VG Configuration
86: if [ "$dev_type" = "logical" ]
87: then
88:     vg_name="`echo $device | /usr/bin/awk -F/ ' { print $3 }'`"
89:     cfg_file="/etc/lvmconf/"$vg_name".conf"
90:     if [ ! -r $cfg_file ]
91:     then
92:         exit 31
93:     fi
94:     pv_rname="`/usr/sbin/vgcfgrestore -n $device -l 2> /dev/null |
/usr/bin/awk ' {print $1}' | /usr/bin/grep %"/dev/rdsk%"`"
95:     if [ "$pv_rname" = "" ]
96:     then
97:         echo "NULL: /usr/sbin/vgcfgrestore -n $device -l 2> /dev/null |
/usr/bin/awk ' {print $1}' | /usr/bin/grep %"/dev/rdsk%" >
98:         $err_log_path/$dev.vgcfgrestore 2>&1
99:         exit 31
100:     fi
101:     /usr/sbin/vgchange -a n $device > /dev/null 2>&1
102:     reststatus=0
103:     for pv_rname1 in $pv_rname
104:     do
105:         /usr/sbin/vgcfgrestore -n $device $pv_rname1 > /dev/null 2>&1
106:         if [ $? != 0 ]
107:         then
108:             /usr/sbin/vgcfgrestore -n $device $pv_rname1 >>
109:             $err_log_path/$dev.vgcfgrestore 2>&1
110:         fi
111:         if [ $? = 0 ]
112:         then
113:             /usr/bin/rm -f $err_log_path/$dev.vgcfgrestore
114:             reststatus=1
115:             break
116:         fi
117:     done
118:     if [ "$SWSTGNODE" != "" ]
119:     then
120:         /usr/sbin/vgchange -c y $device > /dev/null 2>&1
121:         /usr/sbin/vgchange -a e $device > /dev/null 2>&1

```

```

120:     else
121:         /usr/sbin/vgchange -a y $device > /dev/null 2>&1
122:     fi
123:     if [ $reststatus = 0 ]
124:     then
125:         exit 31
126:     fi
127: fi
128:
129: # No processing
130: # When device was not mounted --- 2
131: # When devices of volume group was not mounted --- 4
132: # When the transaction disk is a VxVM physical volume --- 5
133: if [ "$post" = "none" ]
134: then
135:     /usr/bin/rm -rf $post_file 2> /dev/null
136:     /usr/bin/rm -rf $fstype_file 2> /dev/null
137:     exit 0
138: fi
139:
140: # mount processing
141: if [ "$post" = "mount" ]
142: then
143:
144: # When device was mounted --- 1
145:     if [ "$dev_type" = "physical" ]
146:     then
147:         /usr/bin/df -l $device > /dev/null 2>&1
148:         if [ $? != 0 ]
149:         then
150:             if [ ! -r $fstype_file ]
151:             then
152:                 /usr/sbin/mount $device $mount_point 2>
/dev/null
153:             else
154:                 if [ "$fs" = "" ]
155:                 then
156:                     /usr/sbin/mount $device $mount_point 2>
/dev/null
157:                 else
158:                     /usr/sbin/mount -F $fs $device
$mount_point 2> /dev/null
159:                 fi
160:             fi
161:         if [ $? != 0 ]
162:         then
163:             retry_count=3
164:             sleep_time=1
165:             result_flag=1
166:
167:             while [ $retry_count -gt 0 ]
168:             do
169:                 if [ ! -r $fstype_file ]
170:                 then
171:                     /usr/sbin/mount $device $mount_point >
$err_log_path/$dev.mount 2>&1
172:                 else
173:                     if [ "$fs" = "" ]
174:                     then
175:                         /usr/sbin/mount $device
$mount_point > $err_log_path/$dev.mount 2>&1
176:                     else
177:                         /usr/sbin/mount -F $fs $device

```

```

$mount_point > $err_log_path/$dev.mount 2>&1
178:                                     fi
179:                                     fi
180:                                     if [ $? != 0 ]
181:                                     then
182:                                         retry_count=`expr $retry_count - 1`
183:                                         /usr/bin/sleep $sleep_time
184:                                     else
185:                                         /usr/bin/rm -f $err_log_path/$dev.mount
186:                                         result_flag=0
187:                                     break
188:                                     fi
189:                                     done
190:
191:                                     if [ $result_flag != 0 ]
192:                                     then
193:                                         exit 11
194:                                     fi
195:                                     fi
196:                                     fi
197: #     elif [ "$dev_type" = "logical" ]
198: #     then
199: #
200: # When devices of volume group was mounted --- 3
201: #         # Specify the name of volume group to mount
202: #         if [ "$device" = "/dev/vgXX" ]
203: #         then
204: #             # Mount all logical volumes of the volume group
205: #             fs="hfs"
206: #             lvname="/dev/vgXX/XXXXX"
207: #             lv_mount_point="/XX"
208: #
209: #             fsck -F $fs -y $lvname > /dev/null 2>&1
210: #             if [ $? != 0 ]
211: #             then
212: #                 fsck -F $fs -y $lvname > $err_log_path/$dev.fsck
213: #                 if [ $? != 0 ]
214: #                 then
215: #                     exit 12
216: #                 else
217: #                     rm -f $err_log_path/$dev.fsck
218: #                 fi
219: #             fi
220: #             /usr/bin/df -l $lvname > /dev/null 2>&1
221: #             if [ $? != 0 ]
222: #             then
223: #                 /usr/sbin/mount -F $fs $lvname $lv_mount_point 2>
224: #                 /dev/null
225: #                 if [ $? != 0 ]
226: #                 then
227: #                     retry_count=3
228: #                     sleep_time=1
229: #                     result_flag=1
230: #                     while [ $retry_count -gt 0 ]
231: #                     do
232: #                         /usr/sbin/mount -F $fs $lvname
233: #                         $lv_mount_point > $err_log_path/$dev.mount 2>&1
234: #                         if [ $? != 0 ]
235: #                         then
236: #                             retry_count=`expr $retry_count - 1`
237: #                             /usr/bin/sleep $sleep_time

```

```

237: #                                     else
238: #                                     rm -f $err_log_path/$dev.mount
239: #                                     result_flag=0
240: #                                     break
241: #                                     fi
242: #                                     done
243: #
244: #                                     if [ $result_flag != 0 ]
245: #                                     then
246: #                                         exit 11
247: #                                     fi
248: #                                     fi
249: #                                     fi
250: #                                     fi
251: #                                     fi
252: #                                     /usr/bin/rm -rf $post_file 2> /dev/null
253: #                                     /usr/bin/rm -rf $fstype_file 2> /dev/null
254: #                                     exit 0
255: fi
256:
257: exit 0

```

A.3.3 Customizing the shell script

The method to customize the shell script for restore pre/post-processing is explained below.

Restore a volume group

When restoring a file system that is part of a logical volume in a volume group, insert a comment character (“#”) into line 130 of the pre-processing script (OpcRestore.pre) to remove the exit statement

- Eliminate the comment (“#”) from lines 134-170 of the pre-processing script (OpcRestore.pre), and also from lines 197, 198 and 202-250 of the post-processing script (OpcRestore.post), Correct the description in the comment so that unmount/mount processing is carried out on all the logical volumes used to build the file system. Modify the parts shown as “X” to match the actual environment. If more than one logical volume is being processed, repeat lines 138-168 of the pre-processing script (OpcRestore.exe) and lines 206-249 of the post-processing script (OpcRestore.post) once for each additional logical volume.
- If the operating volume was not unmounted in the pre-processing script, and if fsck is not being carried out to the backup volume, then the fsck command will need to be executed (to all the logical volumes contained in the backup volume) in post-processing



Note

Note that unmount/mount processing to a volume group is not carried out if the scripts have not been customized.



Note

When the transaction volume is a volume group, in order to enable use of the backup volume as LVM, volume group configuration information is restored using the `vgcfgrestore` command in the restore post-processing. In the above-mentioned script, volume configuration information is restored from the standard backup file `"/etc/lvmconf/volume-group-name.conf"`. If volume group configuration information is backed up by another file, please customize the script. In cluster operation, the respective nodes comprising the cluster must have volume group configuration information.



Note

Please correct parameters, options, etc. of the OS command, such as the “mount” command and the “fsck” command, according to their employment.



Note For shared-mode volume groups:

When reconfiguring shared-mode volume groups, comment out line numbers 96-127, 133, 134 and 138 in the post-processing script (OpcBackup.post) to prevent the volume group from being reconfigured by the post-processing script.

Reconfigure the volume group by manually performing the following steps after the Replication has been executed.

1. Stop the volume group (on all nodes involved in the transaction).

```
# /usr/sbin/vgchange -a n <vg_name>
#
```

2. Restore volume group configuration information (on the node where the volume group was created).

```
# /usr/sbin/vgcfgrestore -n <vg_name> <pv_path>
#
```

3. Mark the volume group as shareable (on the node where the volume group was created).

```
# /usr/sbin/vgchange -S y -c y <vg_name>
#
```

4. Start the volume group (on all nodes involved in the transaction).

```
# /usr/sbin/vgchange -a s <vg_name>
#
```

File system name

Change the file system name on line 205 of the post-processing shell script (OpcRestore.post) that is executed when restore is performed as appropriate for the operation.

Cluster operations on VERITAS Cluster Server

Customization of the pre-processing and post-processing scripts is required, when performing cluster operations on VERITAS Cluster Server, and when the mount point of the transaction volume is registered in the cluster transaction.

Change unmount/mount processing in the pre-processing and post-processing script into offline/online processing of the mount point resource.

Please note that once the online/offline command is issued there will be a delay before the mount/unmount command can be executed. The “sleep” and “df” commands can be used to delay the mount/ unmount until after the process that sets the volumes online or offline has completed

The examples below show how the pre-processing and post-processing scripts are modified.

[Example] Changing the unmount processing of the restore pre-processing script
(OpcRestore.pre)
[line 80, 89, 139, 148th]

(Before change)	/usr/sbin/umount \$mount_point
(after change)	/opt/VRTSvcs/bin/hares -offline resource-name -sys <i>system-name</i>

[Example] Adding unmount wait processing to the restore pre-processing script
(OpcRestore.pre)
[line 110, 169th]

(addition)	while /usr/bin/df -l "\$device">/dev/null 2>&1; do ;; done
------------	--

[Example] Changing mount processing of the restore post-processing script
(OpcRestore.post)
[line 152, 156, 158, 171, 175, 177, 223, 232nd]

(Before change)	/usr/sbin/mount \$device \$mount_point /usr/sbin/mount -F \$fs \$device \$mount_point /usr/sbin/mount -F \$fs \$lvname \$lv_mount_point
(after change)	/opt/VRTSvcs/bin/hares -online resource-name -sys <i>system-name</i>

[Example] Adding mount wait processing to the restore post-processing script
(OpcRestore.post)
[line 197, 251st]

(addition)	while ! /usr/bin/df -l "\$device">/dev/null 2>&1; do ;; done
------------	--

Appendix B Replication Pre-processing and Post-processing

This appendix explains the shell scripts that perform the replication pre-processing and post-processing in AdvancedCopy Manager.

B.1 Overview

The shell scripts that perform the replication pre-and post-processing start before and after the replication (copy) processing.

These shell scripts contain code required by AdvancedCopy Manager for the copy source volume and copy destination volume.

This chapter describes the 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.

B.2 Replication Pre-processing and Post-processing

The replication pre-processing and post-processing in AdvancedCopy Manager needs to be performed for both the replication source volume and the replication destination volume. However, the purpose of each is different.

- 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 any application from accessing a destination volume during replication processing. 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 performed in pre-processing and post-processing differ according to the contents of the volume. The next section describes the pre-processing and post-processing of a file system.

B.2.1 Pre-processing and post-processing for a file system

The pre-processing script should unmount the file system and the post-processing script should mount it. (For Windows, the processing consists of allocating and canceling a drive letter.) If no volume is mounted, the pre-processing and post-processing are not performed. When the replication target volume is a physical disk, 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: Pre-processing and post-processing for a file system

Command	Copy status	Source/t arget	Pre-proc essing	Post-pro cessing
Replication start command (swsrpstartsync)	Copy not started or replication establishment	Source	X	X
		Destinat ion	0--*1	X
Replication creation command (for synchronous replication) (swsrpmake)	equivalency maintenance	Source	0	0
		Destinat ion	X	0--*2
Replication cancellation command (swsrpcancel)	equivalency maintenance	Source	0	0
		Destinat ion	X	0--*2
	Replication establishment	Source	X	X
		Destinat ion	X	X
	copying	Source	X	X
		Destinat ion	X	X
Replication creation command (for snapshot replication) (swsrpmake)	Copy not started	Source	0	0
		Destinat ion	0	0

0: executed

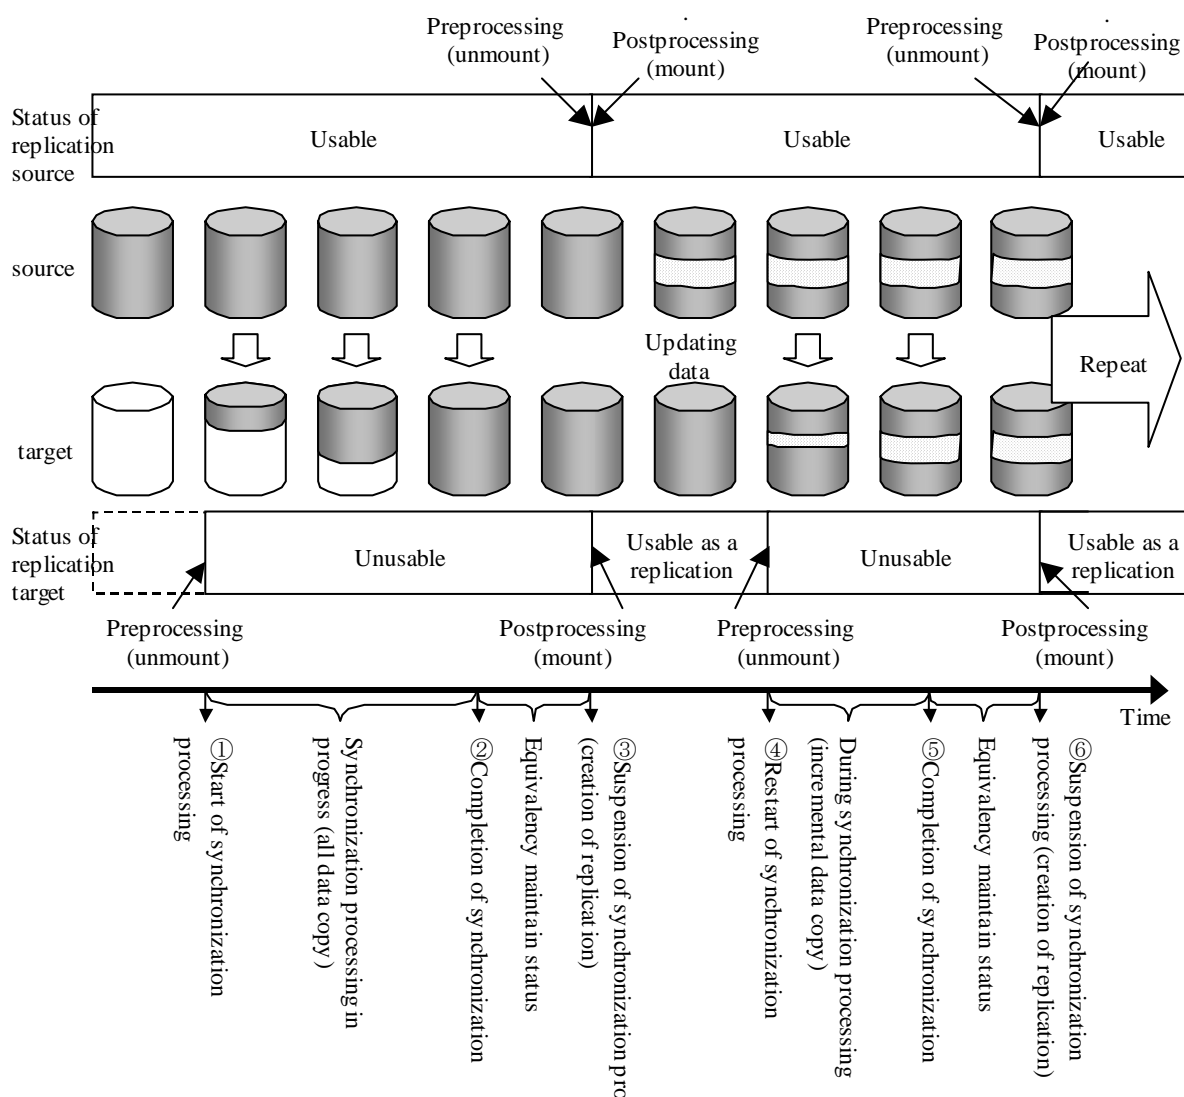
X: not executed

The post-processing for the pre-processing performed in (1) is performed in (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 below.

Figure: Snapshot replication processing of snapshot



If processing other than the standard pre-processing and post-processing is required, customize the scripts.

When customizing a script, strictly observe the following rules regarding error codes:

Table: Error code rules

Error code	Usage
0-99	Unusable (reserved for AdvancedCopy Manager)
100-255	Usable



Note

When the replication target is a volume group that includes a logical volume in which the file system has been built, refer to "When the replication target is a volume group". The pre-processing and post-processing scripts have to be customized to execute unmount/mount processing for all logical volumes in which the file system has been built.



Note

If the replication target volume cannot be unmounted, see "When you do not want to unmount the replication target volume" and customize the script for pre-processing.



Note

AdvancedCopy Manager does not support any functions for the pre-processing and post-processing of a database. You will need to perform your own pre-processing and post-processing before and after the execution of an AdvancedCopy Manager command.

B.2.2 Pre-processing when replication is executed

The pre-processing scripts for the replication source volume (RepSrcPre.js) and the replication destination volume (RepDstPre.js) are stored in the following directory. Customize the scripts according to the processing requirements.

- In the case of non-cluster operation:

```
/etc/opt/FJSVswsrp/sh directory subordinate
```

- In the case of cluster operation:

```
/etc/opt/FJSVswsrp/<logic node name>/sh directory subordinate
```

- RepSrc.pre (copy source volume script)

```

1: #!/bin/sh
2:
3: # AdvancedCopy Manager
4: # All Rights Reserved, Copyright FUJITSU LIMITED, 2002-2007
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: #           99: Script not customize
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: # Device type check
30: # Determination postprocessing file name
31: trans=`echo $device | /usr/bin/grep "/dev/dsk/"`
32: lvmtrans=`echo $device | /usr/bin/grep "/dev/"`
33: vxpvtrans=`echo $device | /usr/bin/grep "/dev/vx/dmp/"`
34: if [ "$trans" != "" ]
35: then
36:     dev_type="physical"
37:     # /dev/dsk/c?t?d? -> c?t?d?
38:     dev=`echo $device | /usr/bin/sed "s/¥/dev¥/dsk¥//"`
39:     post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".spre"
40:
41:     fstype_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".sfstype"
42: elif [ "$vxpvtrans" != "" ]
43: then
44:     dev_type="vxvm_pv"
45:     # /dev/vx/dmp/XXXX -> XXXX
46:     dev=`echo $device | /usr/bin/awk -F/ ' { print $5 } '`
47:     post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".spre"
48:
49:     fstype_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".sfstype"
50: elif [ "$lvmtrans" != "" ]
51: then
52:     dev_type="logical"
53:     # /dev/vgXX -> vgXX
54:     dev=`echo $device | /usr/bin/awk -F/ ' { print $3 } '`
55:     post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".spre"
56: else
57:     exit 0
58: fi
59:
60: err_log_path="/var/opt/FJSVswsrp/"$SWSTGNODE"/log"
61: # When the disk is a physical device #####

```

```

61: if [ "$dev_type" = "physical" ]
62: then
63:
64:     if [ "$mount_point" != "" ]
65:     then
66: # When device cannot be unmounted --- 1
67: #         # Specify the name of volume group not to unmount
68: #         if [ "$device" = "/dev/dsk/cXtXdX" ]
69: #         then
70: #             sync
71: #             sync
72: #             echo "none" > $post_file
73: #
74: # When device can be unmounted --- 2
75: #         else
76: #             /usr/bin/df -ln $mount_point 2>/dev/null | /usr/bin/awk
-F: '{ print $2 }' | /usr/bin/awk '{ print $1 }' > $fstype_file
77: #             /usr/sbin/umount $mount_point 2>/dev/null
78: #
79: #             if [ $? != 0 ]
80: #             then
81: #                 retry_count=3
82: #                 sleep_time=1
83: #                 result_flag=1
84: #
85: #                 while [ $retry_count -gt 0 ]
86: #                 do
87: #                     /usr/sbin/umount $mount_point >
$err_log_path/$dev.umount 2>&1
88: #                     if [ $? != 0 ]
89: #                     then
90: #                         retry_count=`expr $retry_count - 1`
91: #                         /usr/bin/sleep $sleep_time
92: #                     else
93: #                         /usr/bin/rm -f $err_log_path/$dev.umount
94: #                         result_flag=0
95: #                         break
96: #                     fi
97: #                 done
98: #
99: #                 if [ $result_flag != 0 ]
100: #                 then
101: #                     /usr/sbin/fuser -cu $mount_point>
$err_log_path/$dev.fuser 2>&1
102: #                     /usr/bin/ps -ef > $err_log_path/$dev.ps 2>&1
103: #                     exit 10
104: #                 fi
105: #             fi
106: #             echo "mount,$mount_point" > $post_file
107: #         fi
108: # When device was not mounted --- 3
109: #         else
110: #             echo "none" > $post_file
111: #         fi
112: #
113: # When the disk is a volume group #####
114: elif [ "$dev_type" = "logical" ]
115: then
116:
117: # Devices is volume group and script not customize
118:     exit 99
119:

```

```

120: # When devices of volume group cannot be unmounted --- 4
121: #
122: #     # Specify the name of volume group not to unmount
123: #     if [ "$device" = "/dev/vgXX" ]
124: #     then
125: #         sync
126: #         sync
127: #         echo "none" > $post_file
128: #     fi
129: #
130: # When devices of volume group can be unmounted --- 5
131: #
132: #     # Specify the name of volume group to unmount
133: #     if [ "$device" = "/dev/vgXX" ]
134: #     then
135: #         # Unmount all logical volumes of the volume group
136: #         mount_point="/XX"
137: #         /usr/sbin/umount $mount_point 2>/dev/null
138: #         if [ $? != 0 ]
139: #         then
140: #             retry_count=3
141: #             sleep_time=1
142: #             result_flag=1
143: #
144: #             while [ $retry_count -gt 0 ]
145: #             do
146: #                 /usr/sbin/umount $mount_point >
147: # $err_log_path/$dev.umount 2>&1
148: #                 if [ $? != 0 ]
149: #                 then
150: #                     retry_count=`expr $retry_count - 1`
151: #                     sleep $sleep_time
152: #                 else
153: #                     rm -f $err_log_path/$dev.umount
154: #                     result_flag=0
155: #                     break
156: #                 fi
157: #             done
158: #
159: #             if [ $result_flag != 0 ]
160: #             then
161: #                 /usr/sbin/fuser -cu $mount_point>
162: # $err_log_path/$dev.fuser 2>&1
163: #                 /usr/bin/ps -ef > $err_log_path/$dev.ps 2>&1
164: #                 exit 10
165: #             fi
166: #         fi
167: #         echo "mount" > $post_file
168: #     fi
169: #
170: # When devices of volume group was not mounted --- 6
171: #
172: #     # Specify the name of volume group to do nothing
173: #     if [ "$device" = "/dev/vgXX" ]
174: #     then
175: #         echo "none" > $post_file
176: #     fi
177: #
178: # When Src devices is a VxVM physical device --- 7 #####
179: elif [ "$dev_type" = "vxvm_pv" ]
then
# Nothing is done to VxVM PV.

```



```
180:      #
181:      echo "none" > $post_file
182:
183: fi
184:
185: exit 0
```

- RepDst.pre (copy destination volume script)

```

1: #!/bin/sh
2:
3: # AdvancedCopy Manager
4: # All Rights Reserved, Copyright FUJITSU LIMITED, 2002-2007
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: #             99: Script not customize
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: # Device type check
30: # Determination postprocessing file name
31: trans=`echo $device | /usr/bin/grep "/dev/dsk/"`
32: lvmtrans=`echo $device | /usr/bin/grep "/dev/"`
33: vxpvtrans=`echo $device | /usr/bin/grep "/dev/vx/dmp/"`
34: if [ "$trans" != "" ]
35: then
36:     dev_type="physical"
37:     # /dev/dsk/c?t?d? -> c?t?d?
38:     dev=`echo $device | /usr/bin/sed "s/¥/dev¥/dsk¥//"`
39:     post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".dpre"
40:
41:     fstype_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".dfstype"
42: elif [ "$vxpvtrans" != "" ]
43: then
44:     dev_type="vxvm_pv"
45:     # /dev/vx/dmp/XXXX -> XXXX
46:     dev=`echo $device | /usr/bin/awk -F¥/ ' { print $5 } '`
47:     post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".dpre"
48:
49:     fstype_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".dfstype"
50: elif [ "$lvmtrans" != "" ]
51: then
52:     dev_type="logical"
53:     # /dev/vgXX -> vgXX
54:     dev=`echo $device | /usr/bin/awk -F¥/ ' { print $3 } '`
55:     post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".dpre"
56:     conf_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".conf"
57: else
58:     exit 0
59: fi

```

```

60: err_log_path="/var/opt/FJSVswsrp/"$SWSTGNODE"/log"
61:
62:
63: # When the disk is a physical device #####
64: if [ "$dev_type" = "physical" ]
65: then
66:
67:     if [ "$mount_point" != "" ]
68:     then
69: # When device can be unmounted --- 2
70:         /usr/bin/df -ln $mount_point 2>/dev/null | /usr/bin/awk -F:
' { print $2 }' | /usr/bin/awk ' { print $1 }' > $fstype_file
71:         /usr/sbin/umount $mount_point 2>/dev/null
72:
73:         if [ $? != 0 ]
74:         then
75:             retry_count=3
76:             sleep_time=1
77:             result_flag=1
78:
79:             while [ $retry_count -gt 0 ]
80:             do
81:                 /usr/sbin/umount $mount_point >
$err_log_path/$dev.umount 2>&1
82:                 if [ $? != 0 ]
83:                 then
84:                     retry_count=`expr $retry_count - 1`
85:                     /usr/bin/sleep $sleep_time
86:                 else
87:                     /usr/bin/rm -f $err_log_path/$dev.umount
88:                     result_flag=0
89:                     break
90:                 fi
91:             done
92:
93:             if [ $result_flag != 0 ]
94:             then
95:                 /usr/sbin/fuser -cu $mount_point>
$err_log_path/$dev.fuser 2>&1
96:                 /usr/bin/ps -ef > $err_log_path/$dev.ps 2>&1
97:                 exit 10
98:             fi
99:         fi
100:         echo "mount,$mount_point" > $post_file
101: #     fi
102: # When device was not mounted --- 3
103:     else
104:         echo "none" > $post_file
105:     fi
106:
107: # When the disk is a volume group #####
108: elif [ "$dev_type" = "logical" ]
109: then
110:
111: # Devices is volume group and script not customize
112:     exit 99
113:
114: # When devices of volume group can be unmounted --- 5
115: #
116: #     # Specify the name of volume group to unmount
117: #     if [ "$device" = "/dev/vgXX" ]
118: #     then

```

```

119: #           # Unmount all logical volumes of the volume group
120: #           mount_point="/XX"
121: #           /usr/sbin/umount $mount_point 2> /dev/null
122: #           if [ $? != 0 ]
123: #           then
124: #               retry_count=3
125: #               sleep_time=1
126: #               result_flag=1
127: #
128: #               while [ $retry_count -gt 0 ]
129: #               do
130: #                   /usr/sbin/umount $mount_point >
131: #                   $err_log_path/$dev.umount 2>&1
132: #                   if [ $? != 0 ]
133: #                   then
134: #                       retry_count=`expr $retry_count - 1`
135: #                       sleep $sleep_time
136: #                   else
137: #                       rm -f $err_log_path/$dev.umount
138: #                       result_flag=0
139: #                       break
140: #                   fi
141: #               done
142: #               if [ $result_flag != 0 ]
143: #               then
144: #                   /usr/sbin/fuser -cu $mount_point>
145: #                   $err_log_path/$dev.fuser 2>&1
146: #                   /usr/bin/ps -ef > $err_log_path/$dev.ps 2>&1
147: #                   exit 10
148: #               fi
149: #               echo "mount" > $post_file
150: #           fi
151: #
152: # When devices of volume group was not mounted --- 6
153: #
154: #           # Specify the name of volume group to do nothing
155: #           if [ "$device" = "/dev/vgXX" ]
156: #           then
157: #               echo "none" > $post_file
158: #           fi
159: #
160: # When Src/Dst devices is a VxVM physical volume --- 7 #####
161: elif [ "$dev_type" = "vxvm_pv" ]
162: then
163:     # Nothing is done to VxVM PV.
164:     #
165:     echo "none" > $post_file
166:
167: fi
168:
169: exit 0

```

B.2.3 Post-processing when replication is executed

The replication source volume (RepSrcPost.js) and the replication destination volume (RepDstPost.js) are prepared by a post-processing script, which is stored in the following directory. This script can be customized according to the processing requirements.

- In the case of non-cluster operation:

```
/etc/opt/FJSVswsrp/sh directory subordinate
```

- In the case of cluster operation:

```
/etc/opt/FJSVswsrp/<logic node name>/sh directory subordinate
```

- RepSrc.post (copy source volume script)

```

1: #!/bin/sh
2:
3: # AdvancedCopy Manager
4: # All Rights Reserved, Copyright FUJITSU LIMITED, 2002-2007
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: # Device type check
27: # Determination postprocessing file name
28: trans=`echo $device | /usr/bin/grep "/dev/dsk/"`
29: lvmtrans=`echo $device | /usr/bin/grep "/dev/"`
30: vxpvtrans=`echo $device | /usr/bin/grep "/dev/vx/dmp/"`
31: if [ "$trans" != "" ]
32: then
33:     dev_type="physical"
34:     # /dev/dsk/c?t?d? -> c?t?d?
35:     dev=`echo $device | /usr/bin/sed "s/¥/dev¥/dsk¥//"`
36:     post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".spre"
37:
38:     fstype_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".sfstype"
39: elif [ "$vxpvtrans" != "" ]
40: then
41:     dev_type="vxvm_pv"
42:     # /dev/vx/dmp/XXXX -> XXXX
43:     dev=`echo $device | /usr/bin/awk -F¥/ ' { print $5 } '`
44:     post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".spre"
45:
46:     fstype_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".sfstype"
47: elif [ "$lvmtrans" != "" ]
48: then
49:     dev_type="logical"
50:     # /dev/vgXX -> vgXX
51:     dev=`echo $device | /usr/bin/awk -F¥/ ' { print $3 } '`
52:     post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".spre"
53:
54:     fstype_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".sfstype"
55: else
56:     exit 0
57: fi
58: err_log_path="/var/opt/FJSVswsrp/"$SWSTGNODE"/log"
59: # Confirmation of postprocessing

```

```

59: if [ ! -r $post_file ]
60: then
61:     exit 0
62: fi
63: post="/usr/bin/cat $post_file | /usr/bin/cut -d',' -f1`"
64: mount_point="/usr/bin/cat $post_file | /usr/bin/cut -d',' -f2`"
65:
66: # Confirmation of FStype
67: if [ ! -r $fstype_file ]
68: then
69:     fs=""
70: else
71:     fs="/usr/bin/cat $fstype_file`"
72: fi
73:
74: # No processing
75: # When Src device cannot be unmounted          --- 1
76: # When Src device was not mounted              --- 3
77: # When Src devices of volume group cannot be unmounted --- 4
78: # When Src devices of volume group was not mounted --- 6
79: # When Src devices is a VxVM physical device   --- 7
80: if [ "$post" = "none" ]
81: then
82:     /usr/bin/rm -rf $post_file 2> /dev/null
83:     /usr/bin/rm -rf $fstype_file 2> /dev/null
84:     exit 0
85: fi
86:
87: # mount processing
88: if [ "$post" = "mount" ]
89: then
90: # When Src device can be unmounted --- 2
91:     if [ "$dev_type" = "physical" ]
92:     then
93:         /usr/bin/df -l $device > /dev/null 2>&1
94:         if [ $? != 0 ]
95:         then
96:             if [ ! -r $fstype_file ]
97:             then
98:                 /usr/sbin/mount $device $mount_point 2>
/dev/null
99:             else
100:                 if [ "$fs" = "" ]
101:                 then
102:                     /usr/sbin/mount $device $mount_point 2>
/dev/null
103:                 else
104:                     /usr/sbin/mount -F $fs $device
$mount_point 2> /dev/null
105:                 fi
106:             fi
107:             if [ $? != 0 ]
108:             then
109:                 retry_count=3
110:                 sleep_time=1
111:                 result_flag=1
112:
113:                 while [ $retry_count -gt 0 ]
114:                 do
115:                     if [ ! -r $fstype_file ]
116:                     then
117:                         /usr/sbin/mount $device $mount_point >

```

```

    $err_log_path/$dev.mount 2>&1
118:                                     else
119:                                     if [ "$fs" = "" ]
120:                                     then
121:                                     /usr/sbin/mount $device
$mount_point > $err_log_path/$dev.mount 2>&1
122:                                     else
123:                                     /usr/sbin/mount -F $fs $device
$mount_point > $err_log_path/$dev.mount 2>&1
124:                                     fi
125:                                     fi
126:                                     if [ $? != 0 ]
127:                                     then
128:                                     retry_count=`expr $retry_count - 1`
129:                                     /usr/bin/sleep $sleep_time
130:                                     else
131:                                     /usr/bin/rm -f $err_log_path/$dev.mount
132:                                     result_flag=0
133:                                     break
134:                                     fi
135:                                     done
136:
137:                                     if [ $result_flag != 0 ]
138:                                     then
139:                                     exit 11
140:                                     fi
141:                                     fi
142:                                     fi
143: # When devices of volume group can be unmounted --- 5
144: #     elif [ "$dev_type" = "logical" ]
145: #     then
146: #         #Specify the name of volume group to mount
147: #         if [ "$device" = "/dev/vgXX" ]
148: #         then
149: #             # Mount all logical volumes of the volume group
150: #             fs="hfs"
151: #             lvname="/dev/vgXX/XXXXX"
152: #             lv_mount_point="/XX"
153: #
154: #             /usr/bin/df -l $lvname > /dev/null 2>&1
155: #             if [ $? != 0 ]
156: #             then
157: #                 /usr/sbin/mount -F $fs $lvname $lv_mount_point
2>/dev/null
158: #                 if [ $? != 0 ]
159: #                 then
160: #                     retry_count=3
161: #                     sleep_time=1
162: #                     result_flag=1
163: #
164: #                     while [ $retry_count -gt 0 ]
165: #                     do
166: #                         /usr/sbin/mount -F $fs $lvname
$lv_mount_point > $err_log_path/$dev.mount 2>&1
167: #                     if [ $? != 0 ]
168: #                     then
169: #                         retry_count=`expr $retry_count
- 1`
170: #                         /usr/bin/sleep $sleep_time
171: #                     else
172: #                         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: #                               fi
185: #                               fi
186: #                               /usr/bin/rm -rf $post_file 2> /dev/null
187: #                               /usr/bin/rm -rf $fstype_file 2> /dev/null
188: #                               exit 0
189: #                               fi
190: #
191: #                               exit 0
```

- RepDst.post (Copy destination volume script)

```

1: #!/bin/sh
2:
3: # AdvancedCopy Manager
4: # All Rights Reserved, Copyright FUJITSU LIMITED, 2002-2007
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: #           30: VG configuration file not found error
15: #           31: vgcfgrestore error
16: #           12: fsck error
17:
18: # Argument check
19: case $# in
20: 1)
21:     ;;
22: *)
23:     exit 2
24:     ;;
25: esac
26:
27: device=$1
28:
29: # Device type check
30: # Determination postprocessing file name
31: trans="`echo $device | /usr/bin/grep "/dev/dsk/"`"
32: lvmtrans="`echo $device | /usr/bin/grep "/dev/"`"
33: vxpvtrans="`echo $device | /usr/bin/grep "/dev/vx/dmp/"`"
34: if [ "$trans" != "" ]
35: then
36:     dev_type="physical"
37:     # /dev/dsk/c?t?d? -> c?t?d?
38:     dev="`echo $device | /usr/bin/sed "s/¥/dev¥/dsk¥//"`"
39:     post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".dpre"
40:
fstype_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".dfstype"
41: elif [ "$vxpvtrans" != "" ]
42: then
43:     dev_type="vxvm_pv"
44:     # /dev/vx/dmp/XXXX -> XXXX
45:     dev="`echo $device | /usr/bin/awk -F¥/ ' { print $5 }'`"
46:     post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".dpre"
47:
fstype_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".dfstype"
48: elif [ "$lvmtrans" != "" ]
49: then
50:     dev_type="logical"
51:     # /dev/vgXX -> vgXX
52:     dev="`echo $device | /usr/bin/awk -F¥/ ' { print $3 }'`"
53:     post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".dpre"
54:
fstype_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".dfstype"
55: else
56:     exit 0
57: fi
58:

```

```

59: err_log_path="/var/opt/FJSVswsrp/"$SWSTGNODE"/log"
60:
61: # Confirmation of postprocessing
62: if [ ! -r $post_file ]
63: then
64:     exit 0
65: fi
66: post="/usr/bin/cat $post_file | /usr/bin/cut -d',' -f1`"
67: mount_point="/usr/bin/cat $post_file | /usr/bin/cut -d',' -f2`"
68:
69: # Confirmation of FStype
70: if [ ! -r $fstype_file ]
71: then
72:     fs=""
73: else
74:     fs="/usr/bin/cat $fstype_file`"
75: fi
76:
77: # Restore VG Configuration
78: if [ "$dev_type" = "logical" ]
79: then
80:     vg_name="/usr/bin/echo $device | /usr/bin/awk -F/ ' { print $3 }`"
81:     cfg_file="/etc/lvmconf/"$vg_name".conf"
82:     if [ ! -r $cfg_file ]
83:     then
84:         exit 30
85:     fi
86:     pv_rname="/usr/sbin/vgcfgrestore -n $vg_name -l 2> /dev/null |
/usr/bin/awk ' {print $1}' | /usr/bin/grep %"/dev/rds%`"
87:     if [ "$pv_rname" = "" ]
88:     then
89:         echo "/usr/sbin/vgcfgrestore -n $vg_name -l 2> /dev/null |
/usr/bin/awk ' {print $1}' | /usr/bin/grep %"/dev/rds%`" >
$err_log_path/$dev.vgcfgrestore 2>&1
90:         exit 31
91:     fi
92:     /usr/sbin/vgchange -a n $vg_name > /dev/null 2>&1
93:     reststatus=0
94:     for pv_rname1 in $pv_rname
95:     do
96:         /usr/sbin/vgcfgrestore -n $vg_name $pv_rname1 > /dev/null 2>&1
97:         if [ $? = 0 ]
98:         then
99:             /usr/sbin/vgcfgrestore -n $vg_name $pv_rname1 >
$err_log_path/$dev.vgcfgrestore 2>&1
100:             if [ $? = 0 ]
101:             then
102:                 reststatus=1
103:                 break
104:             fi
105:         fi
106:     done
107:     if [ "$SWSTGNODE" != "" ]
108:     then
109:         /usr/sbin/vgchange -c y $vg_name > /dev/null 2>&1
110:         /usr/sbin/vgchange -a e $vg_name > /dev/null 2>&1
111:     else
112:         /usr/sbin/vgchange -a y $vg_name > /dev/null 2>&1
113:     fi
114:     if [ $reststatus = 0 ]
115:     then
116:         exit 31

```

```

117:     fi
118: fi
119:
120: # No processing
121: if [ "$post" = "none" ]
122: then
123: # When Src device cannot be unmounted --- 1
124: #     if [ "$device" = "/dev/dsk/cXtXdX" ]
125: #         then
126: #             rdevice=`echo $device | sed "s/¥/dsk¥//¥/rdsk¥//"`
127: #             src_fs=ufs
128: #             fsck -F $src_fs -y $rdevice > /dev/null 2>&1
129: #             if [ $? != 0 ]
130: #                 then
131: #                     exit 12
132: #                 fi
133: #         fi
134: #
135: # When Src devices of volume group cannot be unmounted --- 4
136: #     if [ "$device" = "/dev/vgXX" ]
137: #         then
138: #             rdevice="/dev/vgXX/XXXXX"
139: #             src_fs=ufs
140: #             fsck -F $src_fs -y $rdevice > /dev/null 2>&1
141: #             if [ $? != 0 ]
142: #                 then
143: #                     exit 12
144: #                 fi
145: #         fi
146: # When Src/Dst device was not mounted --- 3
147: # When Src/Dst devices of volume group was not mounted --- 6
148: # When Src/Dst devices is a VxVM physical volume --- 7
149:     /usr/bin/rm -rf $post_file 2> /dev/null
150:     /usr/bin/rm -rf $fstype_file 2> /dev/null
151:     exit 0
152: fi
153:
154: # mount processing
155: if [ "$post" = "mount" ]
156: then
157:     if [ "$dev_type" = "physical" ]
158:     then
159: # When Dst device can be unmounted --- 2
160:         /usr/bin/df -l $device > /dev/null 2>&1
161:         if [ $? != 0 ]
162:         then
163:             if [ ! -r $fstype_file ]
164:             then
165:                 /usr/sbin/mount $device $mount_point 2>
/dev/null
166:             else
167:                 if [ "$fs" = "" ]
168:                 then
169:                     /usr/sbin/mount $device $mount_point 2>
/dev/null
170:                 else
171:                     /usr/sbin/mount -F $fs $device
$mount_point 2> /dev/null
172:                 fi
173:             fi
174:             if [ $? != 0 ]
175:             then

```

```

176:         retry_count=3
177:         sleep_time=1
178:         result_flag=1
179:
180:         while [ $retry_count -gt 0 ]
181:         do
182:             if [ ! -r $fstype_file ]
183:             then
184:                 /usr/sbin/mount $device $mount_point >
$err_log_path/$dev.mount 2>&1
185:             else
186:                 if [ "$fs" = "" ]
187:                 then
188:                     /usr/sbin/mount $device
$mount_point > $err_log_path/$dev.mount 2>&1
189:                 else
190:                     /usr/sbin/mount -F $fs $device
$mount_point > $err_log_path/$dev.mount 2>&1
191:                 fi
192:             fi
193:             if [ $? != 0 ]
194:             then
195:                 retry_count=`expr $retry_count - 1`
196:                 /usr/bin/sleep $sleep_time
197:             else
198:                 /usr/bin/rm -f $err_log_path/$dev.mount
199:                 result_flag=0
200:                 break
201:             fi
202:         done
203:
204:         if [ $result_flag != 0 ]
205:         then
206:             exit 11
207:         fi
208:     fi
209: fi
210: #     elif [ "$dev_type" = "logical" ]
211: #     then
212: # When Dst devices of volume group can be unmounted --- 5
213: #         # Specify the name of volume group to mount
214: #         if [ "$device" = "/dev/vgXX" ]
215: #         then
216: #             # Mount all logical volumes of the volume group
217: #             fs="hfs"
218: #             lvname="/dev/vgXX/XXXXX"
219: #             lv_mount_point="/XX"
220: #
221: #             /usr/bin/df -l $lvname > /dev/null 2>&1
222: #             if [ $? != 0 ]
223: #             then
224: #                 /usr/sbin/mount -F $fs $lvname $lv_mount_point
2>/dev/null
225: #             if [ $? != 0 ]
226: #             then
227: #                 retry_count=3
228: #                 sleep_time=1
229: #                 result_flag=1
230: #
231: #                 while [ $retry_count -gt 0 ]
232: #                 do
233: #                     /usr/sbin/mount -F $fs $lvname

```

```

$lv_mount_point > $err_log_path/$dev.mount 2>&1
234: #                               if [ $? != 0 ]
235: #                               then
236: #                                   retry_count=`expr $retry_count
- 1`
237: #                                   /usr/bin/sleep $sleep_time
238: #                               else
239: #                                   rm -f $err_log_path/$dev.mount
240: #                                   result_flag=0
241: #                                   break
242: #                               fi
243: #                               done
244: #
245: #                               if [ $result_flag != 0 ]
246: #                               then
247: #                                   exit 11
248: #                               fi
249: #                               fi
250: #                               fi
251: #                               fi
252: #                               fi
253: #                               /usr/bin/rm -rf $post_file 2> /dev/null
254: #                               /usr/bin/rm -rf $fstype_file 2> /dev/null
255: #                               exit 0
256: fi
257:
258: exit 0

```

B.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 a mounted volume, remove the comment characters (“#”) from lines 68-72, 75 and 107 (in the case of a physical volume) or lines 123-128 (in the case of a volume group) in the pre-processing script (RepSrc.pre), and point to the target device in line 70 or 125. In addition, do not perform an update operation on a replication source volume file system after performing the sync command on line 68 or 125 until post-processing for original volume. If a file system update is performed in the meantime, the replica is extracted with an imperfect file system. In this case the fsck command which is carried out by the post-processing script (RepDst.post) may serve as an error.
- The replication source volume must be unmounted.

Replicate a volume group

When replicating a file system which is part of a logical volume in a volume group, the processing scripts will need to be modified as follows:

- Remove the comment (“#”) from lines 133-166 in the source volume pre-processing script (RepSrc.pre), and lines 144-184 in the source volume post-processing script (RepSrc.post). Correct the description in a comment so that the unmount/mount processing is carried out for all logical volumes in which a file system is constructed. Change the sections indicated with an upper case “X” to match the actual environment. If more than one logical volume is being processed, repeat lines 136-164 of the pre-processing script and lines 151-183 of the post-processing script once for each logical volume. Remove the comment (“#”) from lines 117-150 in the destination volume pre-processing script (RepDst.pre), and lines 210, 211 and 214-251 in the destination volume post-processing script (RepDst.post). Correct the description in a comment so that the unmount/mount processing is carried out for all logical volumes in which a file system is constructed. Change the sections indicated with an upper case “X” to match the actual environment. If more than one logical volume is being processed, repeat lines 120-149 of the pre-processing script and lines 218-250 of the post-processing script once for each logical volume.
- When not unmounting copied material volume, it is necessary to carry out fsck to all the logical volumes by which the file system in copy place volume was built with the post-processing script of copy place volume.

Insert a comment (“#”) into line 118 of the source volume pre-processing script and line 112 of the destination pre-processing script.



Note

Note that unmount/mount processing to a volume group is not carried out if the scripts have not been customized.



Note

When a destination volume is a volume group, in order to enable use of the volume for LVM, volume group configuration information is restored using the “vgcfgrestore” command in the restore post-processing. Volume configuration information is restored from the standard backup file “/etc/lvmconf/volume-group-name.conf”. If volume group configuration information is backed up by another file the script will need to be customized.

In cluster operation, the respective nodes comprising the cluster must have volume group configuration information.



Note

Modify the parameters and options of the OS commands (such as the mount command and the fsck command), appropriately for their use.



Note For shared-mode volume groups:

When reconfiguring shared-mode volume groups, comment out lines 78-118, 121, 122 and 152 in the post-processing script (OpcBackup.post) to prevent the volume group from being reconfigured by the post-processing script.

Reconfigure the volume group by manually performing the following steps after the Replication has been executed.

1. Stop the volume group (on all nodes involved in the transaction).

```
# /usr/sbin/vgchange -a n <vg_name>
#
```

- Restore volume group configuration information (on the node where the volume group was created).

```
# /usr/sbin/vgcfgrestore -n <vg_name> <pv_path>
#
```

- Mark the volume group as shareable (on the node where the volume group was created).

```
# /usr/sbin/vgchange -S y -c y <vg_name>
#
```

- Start the volume group (on all nodes involved in the transaction).

```
# /usr/sbin/vgchange -a s <vg_name>
#
```

File system name

Change the file system names on line 150 of the replication source volume post-processing shell script (RepSrc.post) and lines 127, 139, and 217 of the replication destination volume post-processing shell script (RepDst.post) as appropriate for the operation.

Cluster operation using VERITAS Cluster Server

Customization of the pre-processing and post-processing scripts is required when performing cluster operations using VERITAS Cluster Server, and when the mount point of the original/source volume is registered in the cluster transaction.

Change unmount/mount processing in the pre-processing and post-processing scripts into offline/online processing of the mount point resource.

Please note that once the offline/online command is issued there will be a delay before the mount/unmount command can be executed. The “sleep” and “df” commands can be used to delay the mount/unmount until after the process that sets the volumes online or offline has completed. The examples below show how the pre-processing and post-processing scripts are modified.

[Example] Changing the unmount processing of the source volume pre-processing script (RepSrc.pre)

[lines 77, 87, 137 and 146]

(Before change)	/usr/sbin/umount \$mount_point
(after change)	/opt/VRTSvcs/bin/hares -offline resource-name -sys system-name

[Example] Adding unmount wait processing to the source volume pre-processing script (RepSrc.pre)

[lines 106 and 165]

(addition)	while /usr/bin/df -l "\$device">/dev/null 2>&1; do ;; done
------------	--

[Example] Changing the unmount processing of the destination volume pre-processing script (RepDst.pre)

[lines 71, 81, 121 and 130]

(Before change)	/usr/sbin/umount \$mount_point
(after change)	/opt/VRTSvcs/bin/hares -offline resource-name -sys system-name

[Example] Adding unmount wait processing to the destination volume pre-processing script (RepDst.pre)

[lines 100 and 149]

(addition)	while /usr/bin/df -l "\$device">/dev/null 2>&1; do ;; done
------------	--

[Example] Changing the mount processing of the source volume post-processing script (RepSrc.post)

[lines 98, 102, 104, 117, 121, 123, 157 and 166]

(Before change)	/usr/sbin/mount \$device \$mount_point /usr/sbin/mount -F \$fs \$device \$mount_point /usr/sbin/mount -F \$fs \$lvname \$lv_mount_point
(after change)	/opt/VRTSvcs/bin/hares -online resource-name -sys <i>system-name</i>

[Example] Adding mount wait processing to the source volume post-processing script (RepSrc.post)

[lines 143 and 185]

(addition)	while ! /usr/bin/df -l "\$device">/dev/null 2>&1; do ;; done
------------	--

[Example] Changing the mount processing of the destination volume post-processing script

[lines 165, 169, 171, 188, 190, 224 and 233]

(Before change)	/usr/sbin/mount \$device \$mount_point /usr/sbin/mount -F \$fs \$device \$mount_point /usr/sbin/mount -F \$fs \$lvname \$lv_mount_point
(after change)	/opt/VRTSvcs/bin/hares -online resource-name -sys <i>system-name</i>

[Example] Adding mount wait processing to the destination volume post-processing script

[lines 210 and 252]

(addition)	while ! /usr/bin/df -l "\$device">/dev/null 2>&1; do ;; done
------------	--

Appendix C Sample Shell Script for Oracle Integration 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 using AdvancedCopy Manager.



This section explains how the sample shell 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.

C.1 Overview

AdvancedCopy Manager provides some sample shell scripts for backup processing. The backup administrator must modify the sample shell scripts according to their specific requirements and then apply them to backup operations.


C.2 Backup and Restoration of a Transaction volume

C.2.1 Backup procedure

The following is an overview of the backup procedure.

1. Check the work operation status
Check whether processing is in progress or not.
2. Stop Processing
If the transaction volume is in use, stop all update processes to make the transaction volume consistent and ready for backup.
3. Disk backup
Create backup copies of transaction volumes using the backup function of AdvancedCopy Manager.
4. Restart Processing
Restart any processing that was in progress in Step 1.

Figure: Backup method

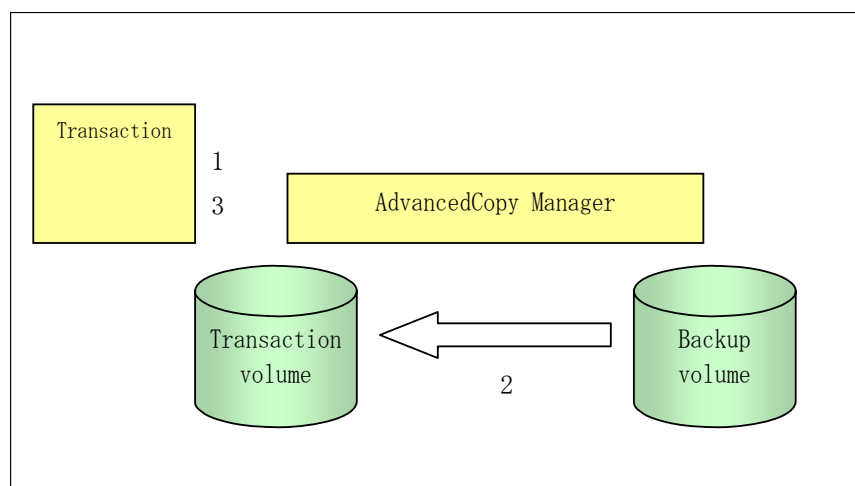


C.2.2 Restore procedure

The following is an overview of performing restore from a backup tape.

1. Check the work operation status
If work is in operation, completely stop the processing.
2. Restore
Use the restore function of AdvancedCopy Manager to restore from the backup volume to the transaction volume.
3. Restart processing
Restart any processing that was in progress in Step 1.

Figure: Restore method



C.3 Sample Shell Scripts

Sample scripts provided for AdvancedCopy Manager are created assuming that the following versions of Oracle are used. If you use any other version, correct the sample scripts as required.

- Oracle 10g (V 10.1.0.2.0)

C.3.1 Overview of sample shell scripts

The sample shell scripts that are provided for AdvancedCopy Manager are summarized below. These scripts are stored in the `/etc/opt/FJVSwsts/samp` directory.

Figure: Sample shell scripts

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

C.3.1.1 Collecting DB information

swst_or_iv outputs the following information from an Oracle database. The output file names and their contents are:

- spdata.lst: List of table space names and storage destinations
- 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
- spreadolog.lst: REDO Log file information of an Oracle database

Use the database information above as a reference for creating shell scripts for backup processing.

a. Format

```
swst_or_iv
```

b. Usage example

```
# swst_or_iv
swst_or_iv completed.
```

c. Result example

```
spdata.lst
  TABLESPACE-NAME                DEVICE-NAME
-----
ACMTS1                            /dev/vg170/rlvol170
SYSaux                             /var/adm/crash/u01/app/oracle/OraHome_1/oradata/acmtest1/sysaux01.dbf
SYSTEM                             /var/adm/crash/u01/app/oracle/OraHome_1/oradata/acmtest1/system01.dbf
```

```

UNDOTBS1
/var/adm/crash/u01/app/oracle/OraHome_1/oradata/acmtest1/undotbs01.dbf
USERS
/var/adm/crash/u01/app/oracle/OraHome_1/oradata/acmtest1/users01.dbf
spdata.bgn
alter tablespace ACMTS1 begin backup;
alter tablespace SYSAUX begin backup;
alter tablespace SYSTEM begin backup;
alter tablespace UNDOTBS1 begin backup;
alter tablespace USERS begin backup;
spdata.end
alter tablespace ACMTS1 end backup;
alter tablespace SYSAUX end backup;
alter tablespace SYSTEM end backup;
alter tablespace UNDOTBS1 end backup;
alter tablespace USERS end backup;
spcont.lst
NAME                                TYPE  VALUE
-----
control_files                        string
/var/adm/crash/u01/app/oracle/
OraHome_1/oradata/acmtest1/control01.ctl, /var/adm/crash/u01
/app/oracle/OraHome_1/oradata/
acmtest1/control02.ctl, /var/a
dm/crash/u01/app/oracle/OraHom
e_1/oradata/acmtest1/control03
.ctl

splog.lst
NAME                                TYPE  VALUE
-----
db_recovery_file_dest                string
/var/adm/crash/u01/app/oracle/
OraHome_1/flash_recovery_area

db_recovery_file_dest_size           big integer 2G

spredolog.lst
GROUP#    MEMBER
-----
3         /var/adm/crash/u01/app/oracle/OraHome_1/oradata/acmtest1/redo03.log
2         /var/adm/crash/u01/app/oracle/OraHome_1/oradata/acmtest1/redo02.log
1         /var/adm/crash/u01/app/oracle/OraHome_1/oradata/acmtest1/redo01.log

```

d. Processing outline

```

#!/bin/sh
reterr() --- An error message is output before an error is returned.
sqldata()

{
sqlplus /nolog <<!
connect / as sysdba
SELECT --- Acquire Oracle table space information.
}
sqllog()

{
sqlplus /nolog <<!
connect / as sysdba

```

```

show --- Acquire Oracle log information.
}
sqlcont()

{
sqlplus /nolog <<!
connect / as sysdba
show --- Acquire Oracle control information.
}
sqlredolog()

{
sqlplus /nolog <<!
connect / as sysdba
show --- Acquire Oracle REDO log information.
}

# main()

# 0) Environmental variable setting
PATH=$PATH:/usr/bin:/usr/sbin:/opt/FJJSwsts/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

# 4) Acquisition of Oracle REDO log information
sqlredolog() invocation --- Acquire Oracle REDO log information

exit 0

```

C. 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 "Backup procedure".

- a. Start

```
swst_or_fs_bk {-a | -i} <mount-point-name>
```

- b. Explanation of operands

```

-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.
mount-point-name: Specify the name of the mount point targeted for backup.

```

- c. 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/vg170 at 23:00 on 2005.10.01 )
```

d. 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 /nolog <<!  
connect / as sysdba  
alter system switch logfile;  
alter system flush shared_pool;  
alter tablespace CCC begin backup; --- Notify Oracle of the start of backup.  
}  
sqlend()  
  
{  
sqlplus /nolog <<!  
connect / as sysdba  
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 $VG --- 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 $VG`" --- Fetch a backup volume as a character device.
BDDATE="`getdate $VG`" --- Fetch the backup date.
BDTIME="`gettime $VG`" --- Fetch the backup time.

exit 0
```

C. 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 "Backup procedure".

a. Start

```
swst_or_rd_bk {-a | -i} <device name>
```

b. Explanation of operands

```
-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.
device name: Specify the name of the raw device targeted for backup.
```

c. Example using swst_or_rd_bk

```
While the database is online, back up /dev/vg170 in which an Oracle database is located.
# swst_or_rd_bk -a /dev/vg170
swst_or_rd_bk completed. (/dev/vg170 -> /dev/vg171 at 23:01 on 2005.10.01 )
```

d. 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.
get_chr_blk() --- Conversion from a character device to a block device
get_blk_chr() --- Conversion from a block device to a character device
sqlbegin()

{
sqlplus /nolog <<!
connect / as sysdba
alter system switch logfile;
alter system flush shared_pool;
alter tablespace CCC begin backup; --- Notify Oracle of the start of backup.
}
sqlend()

{
sqlplus /nolog <<!
connect / as sysdba
```

```

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/FJVSswsts/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.
DST="`get_blk_chr $DSTBLK`"
BKDATE="`getdate $SRC`"
BKTIME="`gettime $SRC`"

exit 0

```

C.3.2 Applying the sample shell scripts

C.3.2.1 DB information collection

Sample shell script update

Update the script `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.

C. 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.

C. 3. 2. 3 Raw device backup

AdvancedCopy Manager setup

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.

C. 3. 3 Restore Method

C. 3. 3. 1 File system

Use the restore function of AdvancedCopy Manager for restore processing.

C. 3. 3. 2 Raw device

Use the restore function of AdvancedCopy Manager for restore processing.

C.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 slices.

Backup of Oracle resources

Note the following when performing this function.

1. Initialization parameter file
Every time a change occurs, back up this file using the cp command of the OS.
2. Control files
 - a. You are recommended to use mirroring using Oracle.
 - b. If a change is made in the DB environment, use the following command to create a trace file:

```
alter database backup control file to trace;
```


If a control file is corrupted, execute the SQL statement in the trace file to recover the control file.
3. Data files
 - a. The recommended location of data files is on a raw device.
 - b. 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.
 - c. Create backup copies in AdvancedCopy Manager.
 - d. 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.
 - e. 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;
```
4. REDO files
 - a. You are recommended to mirror the online REDO log files and archive log files using Oracle or the OS.
 - b. No online REDO log file needs to be backed up.



Note

For online backup 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.

