

FUJITSU Storage ETERNUS SF AdvancedCopy Manager V16.6

Operation Guide

Solaris

B1FW-6011-07ENZO(00)
December 2017

Preface

Purpose

This manual describes the functionality of FUJITSU Storage ETERNUS SF AdvancedCopy Manager (hereafter abbreviated as "AdvancedCopy Manager").

Intended Readers

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

Organization

This manual is composed as follows:

[Chapter 1 Overview of AdvancedCopy Manager](#)

This chapter provides an overview of AdvancedCopy Manager.

[Chapter 2 Starting and Stopping Daemons](#)

This chapter describes how to start and stop daemons provided by AdvancedCopy Manager.

[Chapter 3 Backup Operation](#)

This chapter describes how to backup data.

[Chapter 4 Backup Operation Using Suspend/Resume Function](#)

This chapter describes how to use the Suspend/Resume function in the backup operation.

[Chapter 5 Backup Operation Using QuickOPC Function](#)

This chapter describes how to use the QuickOPC function in the backup operation.

[Chapter 6 Replication](#)

This chapter describes the replication operation.

[Chapter 7 Storage Cluster Continuous Copy Function](#)

This chapter describes the replication operation of AdvancedCopy Manager using the Storage Cluster Continuous Copy function.

[Chapter 8 Operation Maintenance](#)

This chapter describes how to perform configuration and system maintenance tasks and troubleshooting on the system management information database.

[Chapter 9 Command References](#)

This chapter describes commands.

[Chapter 10 Notes on Operations](#)

This chapter describes additional operational notes on Backup and Replication.

[Appendix A Pre-processing and Post-processing of Backup and Restoration](#)

This appendix describes the script used for pre-processing and post-processing of backup and restore functionality.

[Appendix B Sample Script for Backup Operations](#)

This appendix describes some sample scripts using AdvancedCopy Manager.

[Appendix C Pre-processing and Post-processing of Replication](#)

This appendix describes the script used for pre-processing and post-processing replication.

[Appendix D System Volume Backup Operation](#)

This appendix describes how to back up and restore the system volume.

Appendix E Backup Operations Using Power-Saving

This appendix describes backup operations using the power-saving features of disks in ETERNUS Disk storage system.

Appendix F Operation in Oracle Solaris Zones

This appendix describes the operation using the Oracle Solaris Zones Support function.

Appendix G Backup and Restoration of ZFS File Systems

This appendix describes backup and restoration of ZFS file systems.

Appendix H Operation in Oracle VM Server for SPARC

This appendix describes the operation in Oracle VM Server for SPARC.

Notation

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

Operating Systems

Formal Name	Abbreviation	
Microsoft(R) Windows Server(R) 2008 Standard (32-bit)(64-bit) Microsoft(R) Windows Server(R) 2008 Standard without Hyper-V(TM) (32-bit)(64-bit) Microsoft(R) Windows Server(R) 2008 Enterprise (32-bit)(64-bit) Microsoft(R) Windows Server(R) 2008 Enterprise without Hyper-V(TM) (32-bit)(64-bit) Microsoft(R) Windows Server(R) 2008 Datacenter (32-bit)(64-bit) Microsoft(R) Windows Server(R) 2008 Datacenter without Hyper-V(TM) (32-bit)(64-bit)	Windows Server 2008	Windows
Microsoft(R) Windows Server(R) 2008 R2 Foundation Microsoft(R) Windows Server(R) 2008 R2 Standard Microsoft(R) Windows Server(R) 2008 R2 Enterprise Microsoft(R) Windows Server(R) 2008 R2 Datacenter	Windows Server 2008 R2	
Microsoft(R) Windows Server(R) 2012 Standard Microsoft(R) Windows Server(R) 2012 Datacenter	Windows Server 2012	
Microsoft(R) Windows Server(R) 2012 R2 Standard Microsoft(R) Windows Server(R) 2012 R2 Datacenter	Windows Server 2012 R2	
Microsoft(R) Windows Server(R) 2016 Standard Microsoft(R) Windows Server(R) 2016 Datacenter	Windows Server 2016	
Windows(R) 7 Home Basic Windows(R) 7 Home Premium Windows(R) 7 Professional Windows(R) 7 Enterprise Windows(R) 7 Ultimate	Windows 7	
Windows(R) 8 Windows(R) 8 Pro	Windows 8	
Windows(R) 8.1 Windows(R) 8.1 Pro	Windows 8.1	
Windows(R) 10 Home Windows(R) 10 Pro Windows(R) 10 Mobile	Windows 10	
Oracle Solaris 10	Solaris 10	Solaris or Solaris OS
Oracle Solaris 11	Solaris 11	

Formal Name	Abbreviation	
Red Hat(R) Enterprise Linux(R) 5 (for x86) Red Hat(R) Enterprise Linux(R) 5 (for Intel64)	RHEL5	Linux
Red Hat(R) Enterprise Linux(R) 6 (for x86) Red Hat(R) Enterprise Linux(R) 6 (for Intel64)	RHEL6	
Red Hat(R) Enterprise Linux(R) 7 (for Intel64)	RHEL7	
SUSE(R) Linux Enterprise Server 12 for AMD64 & Intel64	SUSE Linux Enterprise Server 12	
HP-UX 11i v3	HP-UX	
AIX(R) V6.1 AIX(R) V7.1	AIX	
VMware(R) vSphere 5	VMware vSphere 5	VMware
VMware(R) vSphere 6	VMware vSphere 6	

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

FUJITSU Storage ETERNUS

Formal Name	Abbreviation		
FUJITSU Storage ETERNUS DX60	-	-	ETERNUS DX series or ETERNUS Disk storage system
FUJITSU Storage ETERNUS DX410 FUJITSU Storage ETERNUS DX440	ETERNUS DX400 series		
FUJITSU Storage ETERNUS DX8100 FUJITSU Storage ETERNUS DX8400 FUJITSU Storage ETERNUS DX8700	ETERNUS DX8000 series		
FUJITSU Storage ETERNUS DX60 S2 FUJITSU Storage ETERNUS DX80 S2 FUJITSU Storage ETERNUS DX90 S2	-	ETERNUS DX S2 series	
FUJITSU Storage ETERNUS DX410 S2 FUJITSU Storage ETERNUS DX440 S2	ETERNUS DX400 S2 series		
FUJITSU Storage ETERNUS DX8100 S2 FUJITSU Storage ETERNUS DX8700 S2	ETERNUS DX8000 S2 series		
FUJITSU Storage ETERNUS DX60 S3 FUJITSU Storage ETERNUS DX100 S3 FUJITSU Storage ETERNUS DX200 S3	-	ETERNUS DX S3 series	
FUJITSU Storage ETERNUS DX500 S3 FUJITSU Storage ETERNUS DX600 S3	-		
FUJITSU Storage ETERNUS DX8700 S3 FUJITSU Storage ETERNUS DX8900 S3	ETERNUS DX8000 S3 series		
FUJITSU Storage ETERNUS DX60 S4 FUJITSU Storage ETERNUS DX100 S4 FUJITSU Storage ETERNUS DX200 S4	-	ETERNUS DX S4 series	
FUJITSU Storage ETERNUS DX500 S4 FUJITSU Storage ETERNUS DX600 S4	-		
FUJITSU Storage ETERNUS AF250 FUJITSU Storage ETERNUS AF650	ETERNUS AF series	ETERNUS AF All-Flash Arrays	
FUJITSU Storage ETERNUS AF250 S2 FUJITSU Storage ETERNUS AF650 S2	ETERNUS AF S2 series		

Formal Name	Abbreviation
FUJITSU Storage ETERNUS DX200F	-
FUJITSU Storage ETERNUS LT20/LT20 S2 FUJITSU Storage ETERNUS LT40/LT40 S2 FUJITSU Storage ETERNUS LT60 S2 FUJITSU Storage ETERNUS LT220 FUJITSU Storage ETERNUS LT230 FUJITSU Storage ETERNUS LT250 FUJITSU Storage ETERNUS LT270	ETERNUS Tape library
FUJITSU Storage ETERNUS Multipath Driver	ETERNUS Multipath Driver

Each individual device name in this document is listed with "FUJITSU Storage" omitted from it.

Software Products

Formal Name	Abbreviation
Windows(R) Internet Explorer(R)	Internet Explorer
Mozilla(R) Firefox(R)	Firefox
Google Chrome(TM)	Chrome
Microsoft(R) Windows Server(R) Failover Clustering	WSFC
Microsoft(R) Exchange Server	Exchange Server
Microsoft(R) SQL Server(R)	SQL Server
VMware(R) ESX(R)	VMware ESX
VMware(R) ESXi(TM)	VMware ESXi
VMware(R) vCenter(TM) Server	VMware vCenter Server or vCenter Server
FUJITSU Storage ETERNUS SF Express	Express
FUJITSU Storage ETERNUS SF Storage Cruiser	Storage Cruiser
FUJITSU Software PRIMECLUSTER GDS FUJITSU Software PRIMECLUSTER GD	PRIMECLUSTER GD
FUJITSU Software PRIMECLUSTER GDS Snapshot FUJITSU Software PRIMECLUSTER GD Snapshot	PRIMECLUSTER GD Snapshot
FUJITSU Software PRIMECLUSTER GFS	PRIMECLUSTER GFS

Manuals

Formal Name	Abbreviation
FUJITSU Storage ETERNUS SF Express / Storage Cruiser / AdvancedCopy Manager Quick Reference	ETERNUS SF Quick Reference
FUJITSU Storage ETERNUS SF Express / Storage Cruiser / AdvancedCopy Manager Release Notes	ETERNUS SF Release Notes
FUJITSU Storage ETERNUS SF Express / Storage Cruiser / AdvancedCopy Manager Installation and Setup Guide	ETERNUS SF Installation and Setup Guide
FUJITSU Storage ETERNUS SF Express / Storage Cruiser / AdvancedCopy Manager Migration Guide	ETERNUS SF Migration Guide
FUJITSU Storage ETERNUS SF Express / Storage Cruiser / AdvancedCopy Manager Web Console Guide	ETERNUS SF Web Console Guide

Formal Name	Abbreviation
FUJITSU Storage ETERNUS SF Storage Cruiser / AdvancedCopy Manager Cluster Environment Setup Guide	ETERNUS SF Cluster Environment Setup Guide
FUJITSU Storage ETERNUS SF Storage Cruiser / AdvancedCopy Manager Operation Guide for VMware vSphere Virtual Volumes	ETERNUS SF Operation Guide for VMware vSphere Virtual Volumes
FUJITSU Storage ETERNUS SF Express / Storage Cruiser / AdvancedCopy Manager Messages	ETERNUS SF Messages
FUJITSU Storage ETERNUS SF Express / Storage Cruiser Event Guide	ETERNUS SF Event Guide
FUJITSU Storage ETERNUS SF Express / Storage Cruiser / AdvancedCopy Manager Glossary	ETERNUS SF Glossary

Besides the above-mentioned, each individual manual name in this document is listed with "FUJITSU Storage" omitted from it.

Others

- In this manual, the following products are collectively described as "Windows Server 2012 or later".
 - Windows Server 2012
 - Windows Server 2012 R2
 - Windows Server 2016
 - In this manual, the following products are collectively described as "Windows Server 2012 R2 or later".
 - Windows Server 2012 R2
 - Windows Server 2016
 - In this manual, unless there is a special disclaimer, ETERNUS Disk storage system includes the ETERNUS AF All-Flash Arrays and the ETERNUS DX200F All-Flash Array.
However, the following functions are not supported in the ETERNUS AF All-Flash Arrays and the ETERNUS DX200F All-Flash Array:
 - Energy saving operation for storage device
 - NAS operation for storage device
- The following function is not supported in the ETERNUS DX200F All-Flash Array:
- Automated Storage Tiering operation at multiple tiers (two or more)

Export Controls

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

Trademarks

- Microsoft, Windows, Windows Server, and Internet Explorer are registered trademarks of Microsoft Corporation in the United States and other countries.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.
- Linux is a registered trademark of Linus Torvalds.
- Red Hat and RPM are registered trademarks of Red Hat, Inc. in the U.S. and other countries.
- Novell is a registered trademark of Novell Inc., and SUSE and the SUSE logo is a trademark, of SUSE LLC, in the United States and other countries.
- HP-UX is a registered trademark of Hewlett-Packard Company in the United States and other countries.

- AIX is a trademark or a registered trademark of International Business Machines Corporation in the United States and other countries.
- VMware, VMware logo, Virtual SMP, and vMotion are the registered trademarks or trademarks of VMware, Inc. in the United States and other countries.
- All other trademarks and product names are the property of their respective owners.

Shipment Date and Revision History

Shipment date	Revision	Document Part Number	
		PDF	HTML
December 2013	1	B1FW-6011-01ENZO(00)	B1FW-6011-01ENZ2(00)
May 2014	1.1	B1FW-6011-01ENZO(01)	B1FW-6011-01ENZ2(01)
June 2014	2	B1FW-6011-02ENZO(00)	B1FW-6011-02ENZ2(00)
January 2015	3	B1FW-6011-03ENZO(00)	B1FW-6011-03ENZ2(00)
March 2015	3.1	B1FW-6011-03ENZO(01)	B1FW-6011-03ENZ2(01)
July 2015	3.2	B1FW-6011-03ENZO(02)	B1FW-6011-03ENZ2(02)
October 2015	4	B1FW-6011-04ENZO(00)	B1FW-6011-04ENZ2(00)
October 2016	5	B1FW-6011-05ENZO(00)	B1FW-6011-05ENZ2(00)
May 2017	6	B1FW-6011-06ENZO(00)	B1FW-6011-06ENZ2(00)
December 2017	7	B1FW-6011-07ENZO(00)	B1FW-6011-07ENZ2(00)

Notes

- No part of this manual may be reproduced without permission.
- This manual is subject to change without advance notice.

Copyright

Copyright 2013-2017 FUJITSU LIMITED

Update History

Content of Update	Updated Section	Revision
The information related to ETERNUS DX200F is described.	"Others" in Preface	1.1
The explanations for the Oracle Database systems are added.	1.2.1, 1.2.2, 1.3.1, 1.3.2	2
Notes regarding the number of simultaneous executions of Concurrent OPC and Concurrent Suspend processes are modified.	6.5.1.1, 6.5.2.3	
The section regarding the transfer mode of the inter-box synchronous replication is modified.	6.5.2.2	
The explanations of backup of databases are modified.	7.1.2.2	
The steps for backing up the data base are added.	7.1.2.2	
The explanations to restore the data base are modified.	7.1.2.3	
Notes of the collecting error survey information in the cluster environment are modified.	7.2.1.1	

Content of Update	Updated Section	Revision
The explanations in regard to how to correspond to certain situations are modified.	7.3.2.2, 7.4.2.1, 8.3.1.1	
The explanations on changing the port number of the Management Server / Managed Server are added.	7.5.3	
The operating environments of AdvancedCopy Manager for the swstresrst command and the swstsrsemth command are modified.	8.1	
Precautions for when operating on a clustered environment are added.	9.1.6	
The explanations in regard to energy-saving operation (eco-mode) are altered.	Appendix E	
Information related to the ETERNUS DX200F is modified.	"Others" in Preface	3
Notes for the Advanced Copy function that uses TFOV's are added.	1.2.1, Chapter3, Chapter4, Chapter5	
The explanation of Storage Cluster Continuous Copy function is added.	1.2.2	
Solaris 11 is added to the software condition of ETERNUS Multipath Driver.	1.3.2, 1.4.1	
The explanations of preparations are modified.	4.4, 5.4	
The explanations of the replication operation of AdvancedCopy Manager using the Storage Cluster Continuous Copy function are added.	1.2.1, 1.2.2, Chapter3, Chapter4, Chapter5, Chapter7, 9.1, 9.1.1.3, 9.1.2.1, 9.1.2.3, 9.1.2.8, 9.3.1.1, 9.3.1.4, 9.3.1.5, 9.3.2.1, 9.3.2.2, 9.3.2.3, 9.3.2.4, 9.3.2.6, 9.5.1, 9.5.2, F.2	
The corrective procedure when hardware failure occurs in volumes is modified.	8.3.2.2, 8.4.2.1	
The explanation when migrating from ETERNUS SF Manager V16.1 or later is added to options.	9.2.16	
The support device of the guest domain is added.	H4.2	
The term "VMware Virtual Volumes" is changed to "VMware vSphere Virtual Volumes" according to the terms used by VMware, Inc.	Throughout this manual	
Information related to VMware vSphere(R) 6 is added.	"Notation" in Preface	
The command name of Figure 6.30 is corrected.	6.5.2.2	
The Copyright is modified.	A.2.2.1, A.3.1.1, A.3.2.1, C.2.2.2, C.2.3.1, C.2.3.2	
Information related to ETERNUS DX8700 S3/DX8900 S3 is added.	"Notation" in Preface	3.2
The reference destination manual is modified.	H.4.2	
Information related to Windows Server 2003 is deleted.	"Notation" in Preface	4
Information related to Windows 10 is added.	"Notation" in Preface	
A note is added.	3.4.3, 4.4.3, 5.4.3, 6.4.3, F.7.4	
The explanation related to the stgxfwcmmodsrv command execution is modified.	8.5.1 - 8.5.4, 8.5.6.2 - 8.5.6.4	
The explanation related to the operating environment during a command execution is modified.	9.1, 9.3	

Content of Update	Updated Section	Revision
A point is added.	H.4.2	
An unnecessary information is deleted.	H.4.2	
Information related to HP-UX 11i v1/v2 is deleted.	"Notation" in Preface	5
Information related to VMware vSphere 4 is deleted.	"Notation" in Preface	
Information related to the ETERNUS DX80/DX90 is deleted.	"Notation" in Preface, 9.3.2.2, 10.1.2, 10.1.9, 10.1.10	
Information related to the ETERNUS AF250/AF650 is added.	"Notation" and "Others" in Preface	
Information related to the ETERNUS LT60 is deleted.	"Notation" in Preface	
The description of referring to the <i>ETERNUS SF Web Console Guide</i> is added.	3.4.4, 4.4.4, 5.4.4, 6.4.4	
The explanation related to SnapOPC/SnapOPC+ is modified.	6.1.1.1, 6.1.1.2, 6.2.3.1 - 6.2.3.5, 6.5.1, 8.4.2, 8.4.2.3, 9.3.2.2, 9.5.1	
The operations to change the device configuration are added.	8.5.5	
The explanation of the item "Manager IP address" is modified.	9.2.10	
The explanation is modified.	10.1.1.13, 10.1.8	
Information related to Windows Server 2016 is added.	"Notation" and "Others" in Preface	6
Information related to Windows Vista is deleted.	"Notation" in Preface	
Information related to SUSE Linux Enterprise Server 11 is deleted and information related to SUSE Linux Enterprise Server 12 for AMD64 & Intel64 is added.	"Notation" in Preface	
Information related to the ETERNUS DX60 S4/DX100 S4/DX200 S4 is added.	"Notation" in Preface	
The reference destination related to the combination of the manager program and the agent program is changed.	1.1.2	
In [Note], information is added.	6.1.2	
A new condition is added.	6.4.7.1	
A new section is added.	8.5.8	
In DESCRIPTION, information is modified.	9.2.10	
In NOTES, the factors that cannot set the replication volume information are added.	9.3.1.1	
Information related to Solaris 9 is deleted.	"Notation" in Preface, 1.3.2, F.3	7
Information related to the ETERNUS DX500 S4/DX600 S4 is added.	"Notation" in Preface	
Information related to the ETERNUS AF250 S2/AF650 S2 is added.	"Notation" in Preface, 8.5.8	
[Point] is added.	6.5.2.2, F.7.7	
The explanation in the table 8.1 and the table 8.2 is modified.	8.1	
The explanation is modified.	10.1.1.14	
The MPLB resources are added for the devices that can use the Oracle Solaris Zones Support function.	F.5	
In [See], the explanation is added.	F.7.1	

Manual Organization and Reading Suggestions

Manual Organization

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

When to Read	Related Manuals (Abbreviated)	Related Products (*1)			Explanation
		EXP	SC	ACM	
Before installation	Quick Reference	Yes	Yes	Yes	This manual is common for all products.
	Release Notes	Yes	Yes	Yes	This manual is common for all products.
	Overview	-	-	Yes	This manual is unique for the product.
During installation	Installation and Setup Guide	Yes	Yes	Yes	This manual is common for all products.
	Cluster Environment Setup Guide	-	Yes	Yes	This manual is common for Storage Cruiser and AdvancedCopy Manager.
	Migration Guide	Yes	Yes	Yes	This manual is common for all products.
During operation	Operation Guide	Yes	Yes	Yes	<p>This manual is unique for each product. The following manuals are available:</p> <ul style="list-style-type: none"> - Express Operation Guide - Storage Cruiser Operation Guide - Storage Cruiser Operation Guide for Optimization Function - AdvancedCopy Manager Operation Guide (for Windows) - AdvancedCopy Manager Operation Guide (for Solaris) - AdvancedCopy Manager Operation Guide (for Linux) - AdvancedCopy Manager Operation Guide (for HP-UX) - AdvancedCopy Manager Operation Guide (for AIX) - AdvancedCopy Manager Operation Guide for Copy Control Module - Storage Cruiser / AdvancedCopy Manager Operation Guide for VMware vSphere Virtual Volumes
	Web Console Guide	Yes	Yes	Yes	This manual is common for all products.
Anytime	Event Guide	Yes	Yes	-	This manual is common for Express and Storage Cruiser.
	Messages	Yes	Yes	Yes	This manual is common for all products.
	Glossary	Yes	Yes	Yes	This manual is common for all products.

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

How to Read Manuals

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

Purpose	Manual	Main Contents	How to Read
Acquiring a product overview and basic operation knowledge	Quick Reference	<ul style="list-style-type: none"> - Product overview - Installation decision - Overview of the necessary tasks from installation to first use 	Read if you want to acquire a fundamental knowledge of the product and its operation in order to decide to install it or not.
	AdvancedCopy Manager Overview	<ul style="list-style-type: none"> - Main functions - Linkable applications - Procedure overview for Advanced Copy of ETERNUS Disk storage system 	
Confirming the updated contents	Release Notes	<ul style="list-style-type: none"> - New function overview with - Incompatibilities with previous version - Fixed bugs 	Read if you want to know the updated contents from a previous version and if you perform the upgrade.
Deciding if an upgrade is required	Migration Guide	<ul style="list-style-type: none"> - Notes and cautions about the upgrade - Upgrade procedure 	Read if you want to upgrade from a previous version.
Installing and correctly operating the product Setting up operating environment depending on purpose	Installation and Setup Guide	<ul style="list-style-type: none"> - Operating environment - Installation procedure - Setup procedure - Uninstallation procedure 	Read if you want to install and setup the product.
	Cluster Environment Setup Guide	<ul style="list-style-type: none"> - Supported cluster software - Installation procedure for a clustered system - Setup procedure for a clustered system - Uninstallation procedure for a clustered system 	Read if you want to install and setup the product on a clustered system.
Administration and operation of the installed system	Express Operation Guide	<ul style="list-style-type: none"> - Starting and stopping the software - Device monitoring - Data copy inside the storage system - Necessary tasks after an architectural modification of the system as well as product maintenance 	Read if you want to start or shutdown the system, monitor the operation status, do backup/restore operations, etc.
	Storage Cruiser Operation Guide	<ul style="list-style-type: none"> - Starting and stopping the software - Device monitoring - Necessary tasks after an architectural modification of 	

Purpose	Manual	Main Contents	How to Read
		<ul style="list-style-type: none"> the system as well as product maintenance - Command reference 	
	Storage Cruiser Operation Guide for Optimization Function	<ul style="list-style-type: none"> - Operating environment construction - Operating status monitoring - Necessary tasks after an architectural modification of the system as well as product maintenance - Command reference 	
	AdvancedCopy Manager Operation Guide (for Windows)	<ul style="list-style-type: none"> - Starting and stopping the software 	
	AdvancedCopy Manager Operation Guide (for Solaris)	<ul style="list-style-type: none"> - Data backup/restore inside the storage system 	
	AdvancedCopy Manager Operation Guide (for Linux)	<ul style="list-style-type: none"> - Necessary tasks after an architectural modification of the system as well as product maintenance 	
	AdvancedCopy Manager Operation Guide (for HP-UX)	<ul style="list-style-type: none"> - Command reference 	
	AdvancedCopy Manager Operation Guide (for AIX)		
	AdvancedCopy Manager Operation Guide for Copy Control Module		
	Storage Cruiser / AdvancedCopy Manager Operation Guide for VMware vSphere Virtual Volumes	<ul style="list-style-type: none"> - Operating environment construction - Virtual machine operation using a virtual volume - Backup/restore of the virtual machine - Necessary tasks after an architectural modification of the system as well as product maintenance 	
	Web Console Guide	<ul style="list-style-type: none"> - Operating environment - Screen layout description 	Read if you want to understand the ETERNUS SF Web Console.
Dealing with messages issued by the software	Messages	<ul style="list-style-type: none"> - Messages and their explanations - Parameter (variable information) description - System action - Countermeasures 	Read if you want a practical way of investigating and dealing with messages issued by the software.
Dealing with events issued by the software	Event Guide	<ul style="list-style-type: none"> - Phenomenon of event - Countermeasures 	Read if you need to find a practical way of investigating and dealing with events.

Purpose	Manual	Main Contents	How to Read
Researching the meaning of specific terms related to the products and other important terms	Glossary	<ul style="list-style-type: none"> - Product specific terminology explanation - Explanation of important terminology appearing in the manual - Synonyms and related terms - Proper form of abbreviated terms 	Read if you want to learn the meaning of important terms, product specific terms or abbreviations used in the manuals.

Contents

Chapter 1 Overview of AdvancedCopy Manager.....	1
1.1 Product Configuration of AdvancedCopy Manager.....	1
1.1.1 Functions.....	1
1.1.2 System Configuration.....	3
1.2 Function Overview of AdvancedCopy Manager.....	4
1.2.1 What Is Backup Function?.....	4
1.2.2 What Is Replication Function?.....	5
1.2.3 Web Console.....	6
1.2.4 Copy Control Module.....	6
1.3 Operating Environment of AdvancedCopy Manager.....	6
1.3.1 Software Requirement of AdvancedCopy Manager's Manager.....	6
1.3.2 Software Requirement of AdvancedCopy Manager's Agent.....	7
1.4 Managing Device on AdvancedCopy Manager.....	8
1.4.1 Normal Devices.....	8
1.4.2 SDX Objects of PRIMECLUSTER GD/SynfinityDISK.....	8
1.4.3 Volumes under Control of VERITAS Volume Manager (VxVM).....	9
1.4.3.1 Logical Volume.....	9
1.4.3.2 Physical Disks That Constitute Disk Group.....	12
1.5 Backup to Secondary Storage Media.....	13
1.6 Operation with IPv6 Address.....	13
1.6.1 Software Requirements.....	13
1.6.2 Support Target.....	13
1.6.3 Combination in IPv4/IPv6 Addresses Coexisting Environment.....	14
1.6.4 Combination in Environment Where Earlier Version Number and IPv6 Unsupported Operating System Coexist.....	14
1.6.5 IPv6 Address Notation.....	15
Chapter 2 Starting and Stopping Daemons.....	16
2.1 Daemons of AdvancedCopy Manager.....	16
2.2 Starting Daemons.....	16
2.3 Stopping Daemons.....	17
Chapter 3 Backup Operation.....	19
3.1 Overview.....	19
3.1.1 Snapshot Backup Processing.....	20
3.1.2 Synchronized High-speed Backup Processing.....	20
3.1.3 Restoring Data.....	21
3.2 Operation Design.....	22
3.2.1 Selecting Server on Which Backups Are Performed.....	22
3.2.2 Selecting Backup Source.....	23
3.2.3 Configuring Backup Policy.....	24
3.2.4 Preparing Backup Volume.....	24
3.2.5 Preparing Automatic Backup Operation.....	25
3.3 Flow of Operation.....	28
3.4 Preparing to Start Backup.....	29
3.4.1 Activating Daemons.....	29
3.4.2 Starting Web Console.....	29
3.4.3 Registering Managed Server.....	30
3.4.4 Fetching Device Information on Managed Server.....	30
3.4.5 Configuring Environment for Backup Operation Server.....	30
3.4.6 Setting Operation Type for Device.....	30
3.4.7 Setting Backup Policy.....	31
3.4.8 Customizing Pre-processing and Post-processing.....	31

3.4.9 Preparing Device Map File.....	32
3.4.9.1 Describing Device Map File.....	32
3.5 Operation.....	33
3.5.1 Performing Snapshot High-speed Backup.....	33
3.5.2 Performing Synchronized High-speed Backup.....	33
3.5.3 Restoring Data.....	34
3.5.3.1 Restoring Entire Volume.....	34
3.5.3.2 Restoring Individual File.....	34
3.5.4 Managing Backup History.....	34
3.5.4.1 Displaying Backup History.....	34
3.5.4.2 Deleting Backup History.....	34
3.6 Changing Setting Information.....	35
3.6.1 Adding Device.....	35
3.6.2 Deleting Device.....	35
3.6.2.1 Deleting Transaction Volume.....	35
3.6.2.2 Deleting Backup Volume.....	35
3.6.3 Adding Managed Server.....	36
3.6.4 Deleting Managed Server.....	36
3.7 Stopping Operation.....	37
3.8 SDX Object Operations.....	37
3.8.1 Units of SDX Object Operations.....	37
3.8.1.1 Configuring Unit of SDX Object Operations.....	37
3.8.2 Backup Operation in Units of Logical Volumes (Method with PRIMECLUSTER GD Snapshot).....	38
3.8.2.1 Designing Backup Operation.....	39
3.8.2.2 Transaction Volume.....	39
3.8.2.3 Backup Volumes.....	39
3.8.2.4 Copy Functions Available When Used with PRIMECLUSTER GD Snapshot.....	40
3.8.2.5 Configuration Options.....	41
3.8.2.6 Fetching Device Information from Managed Server.....	41
3.8.2.7 Configuring Operation Type for Device.....	41
3.8.2.8 Configuring Backup Policy.....	41
3.8.2.9 Backup.....	42
3.8.2.10 Restore.....	42
3.8.2.11 Operational Notes.....	43
3.8.3 Backup Operation in Units of Slices (Method without PRIMECLUSTER GD Snapshot).....	43
3.8.3.1 Configuring Backup Operation.....	44
3.8.3.2 Fetching Device Information from Managed Server.....	44
3.8.3.3 Configuring Operation Type for Device.....	44
3.8.3.4 Backup.....	45
3.8.3.5 Restore.....	45
3.9 VxVM Volume Operation.....	45
3.9.1 Backup Operation in Units of Logical Volumes.....	45
3.9.1.1 Configuring Backup Operation.....	46
3.9.1.1.1 Supported VxVM Volume Structures.....	46
3.9.1.1.2 Preparing Backup Volume.....	46
3.9.1.2 Preparations.....	47
3.9.1.2.1 Configuring Operation Type of Device.....	47
3.9.1.2.2 Customizing Pre-processing and Post-processing Scripts.....	48
3.9.1.2.3 Preparing Device Map File.....	49
3.9.1.3 Operations.....	49
3.9.1.3.1 Executing Backup.....	49
3.9.1.3.2 Executing Restoration.....	50
3.9.1.4 Changing Configuration Information.....	50

3.9.2 Backup Operation in Units of Physical Disks.....	50
3.9.2.1 Operational Configuration.....	51
3.9.2.2 Preparations.....	52
3.9.2.2.1 Confirming Disk Group Configuration Information File.....	52
3.9.2.2.2 Setting Transaction Volume and Backup Volume.....	52
3.9.2.2.3 Preparing Device Map File.....	53
3.9.2.3 Backup.....	53
3.9.2.4 Restoration.....	56
Chapter 4 Backup Operation Using Suspend/Resume Function.....	60
4.1 Overview.....	60
4.2 Backup Operation Design.....	61
4.3 Flow of Operations.....	62
4.4 Preparations.....	63
4.4.1 Activating Daemons.....	63
4.4.2 Starting Web Console.....	63
4.4.3 Registering Managed Server.....	63
4.4.4 Fetching Device Information on Managed Server.....	63
4.4.5 Configuring Environment for Backup Operation Server.....	63
4.4.6 Setting Operation Type for Device.....	63
4.4.7 Setting Backup Policy.....	64
4.4.8 Customizing Pre-processing and Post-processing.....	64
4.4.9 Preparing Device Map File.....	64
4.5 Backup Operations.....	65
4.5.1 Starting Synchronous Processing.....	67
4.5.2 Backing Up.....	68
4.5.3 Displaying Synchronous Processing Execution Status.....	69
4.5.4 Cancelling Synchronous Processing.....	70
4.5.5 History Management.....	71
4.5.6 Restoring.....	72
Chapter 5 Backup Operation Using QuickOPC Function.....	73
5.1 Overview.....	73
5.2 Configuring Backup Operation.....	74
5.3 Operation Flow.....	75
5.4 Preparation.....	75
5.4.1 Activating Daemons.....	75
5.4.2 Starting Web Console.....	76
5.4.3 Registering Managed Server.....	76
5.4.4 Fetching Device Information on Managed Server.....	76
5.4.5 Configuring Environment for Backup Operation Server.....	76
5.4.6 Setting Operation Type for Device.....	76
5.4.7 Setting Backup Policy.....	77
5.4.8 Customizing Pre-processing and Post-processing.....	77
5.4.9 Preparing Device Map File.....	77
5.5 Operation.....	78
5.5.1 Invoking Backup.....	78
5.5.2 History Management.....	79
5.5.3 Restoration.....	80
5.5.4 Cancelling Restoration.....	81
5.5.5 Displaying Execution Status of Tracking Processing.....	82
5.5.6 Stopping Tracking Processing.....	82
Chapter 6 Replication.....	83

6.1 Overview.....	83
6.1.1 Snapshot Replication Processing.....	83
6.1.1.1 Types of Snapshot Replication.....	84
6.1.1.2 SnapOPC/SnapOPC+ Copy Destination Volume.....	90
6.1.2 Synchronized Replication Processing.....	91
6.2 Configuration.....	92
6.2.1 Determining Server That Performs Replication.....	93
6.2.2 Determining Replication Target.....	93
6.2.2.1 Logical Unit (Disk) Operation.....	95
6.2.3 Notes on Configuration of SnapOPC/SnapOPC+ Replication.....	95
6.2.3.1 Notes on Session Configurations.....	95
6.2.3.2 Notes on Copy Destination Disk.....	98
6.2.3.3 Monitoring Usage of Copy Destination Volume.....	99
6.2.3.4 Notes When Using SDV as Copy Destination Volume.....	100
6.2.3.5 Notes on Cluster Operations.....	101
6.3 Flow of Operations.....	101
6.4 Preparations.....	102
6.4.1 Starting Daemon.....	102
6.4.2 Starting Web Console.....	102
6.4.3 Registering Managed Server.....	103
6.4.4 Fetching Device Information on Managed Server.....	103
6.4.5 Configuring Source Volume and Replica Volume.....	103
6.4.6 Customizing Pre-processing and Post-processing.....	103
6.4.7 Creating Groups.....	104
6.4.7.1 Conditions for Replication Volumes That Make Up Group.....	104
6.4.7.2 How to Create Groups.....	107
6.5 Operations.....	108
6.5.1 Executing Snapshot Replication.....	108
6.5.1.1 Concurrent OPC Functions.....	121
6.5.2 Executing Synchronized Replication.....	122
6.5.2.1 Intra-box Synchronous Replication.....	123
6.5.2.2 Inter-box Synchronous Replication.....	123
6.5.2.2.1 Initial Copy Skip Function.....	128
6.5.2.3 Concurrent Suspend Function.....	131
6.5.2.3.1 Changing Timeout Value for Concurrent Suspend Function.....	133
6.5.2.3.2 Concurrent Suspend When REC Disk buffer Is Used.....	134
6.5.2.4 Destination Access Permission Function.....	134
6.5.3 Backup and Restoration Using Replication Function.....	135
6.5.4 Execution of Synchronization Reversal.....	136
6.6 Changing Configuration Information.....	136
6.6.1 Changing Source Volume or Replica Volume.....	137
6.6.2 Deleting Source Volume or Replica Volume.....	137
6.6.3 Changing Device Information of Source or Replica Volume.....	137
6.7 Stopping Operation.....	137
6.8 SDX Object Replication.....	137
6.8.1 Units of SDX Object Operations.....	138
6.8.2 Replication Operation in Units of Logical Volumes (Method with PRIMECLUSTER GD Snapshot).....	138
6.8.2.1 Configuring Replication Operation.....	139
6.8.2.2 Replication Volumes.....	140
6.8.2.3 Copy Functions Available for Each PRIMECLUSTER GD Structure.....	142
6.8.2.4 Fetching Device Information from Managed Server.....	143
6.8.2.5 Configuring Replication Volumes.....	143
6.8.2.6 Executing Replication.....	143

6.8.2.7 Operational Notes.....	144
6.8.3 Replication Operation in Units of Slices (Method without PRIMECLUSTER GD Snapshot).....	144
6.8.3.1 Configuring Replication Operation.....	145
6.8.3.1.1 Supported Logical Volume Structures.....	145
6.8.3.2 Fetching Device Information from Managed Server.....	145
6.8.3.3 Configuring Replication Volumes.....	145
6.8.3.4 Executing Replication.....	145
6.9 VxVM Volume Operation.....	146
6.9.1 Replication Operation in Units of Logical Volume.....	146
6.9.1.1 Configuration of Operations Involving VERITAS Volume Structures.....	146
6.9.1.1.1 Supported VxVM Volume Structures.....	146
6.9.1.1.2 Preparing Replica Volume.....	146
6.9.1.2 Preparations.....	148
6.9.1.2.1 Setting Original Volume and Replica Volume.....	148
6.9.1.3 Executing Replication.....	148
6.9.1.4 Changing Configuration Information.....	150
6.9.2 Replication Operation in Units of Physical Disks.....	150
6.9.2.1 Configuration.....	151
6.9.2.2 Preparations.....	151
6.9.2.2.1 Confirming Disk Group Configuration Information File.....	151
6.9.2.2.2 Configuring Original and Replica Volumes.....	152
6.9.2.3 Replication.....	152
Chapter 7 Storage Cluster Continuous Copy Function.....	156
7.1 Overview.....	156
7.1.1 System Configuration.....	157
7.1.2 Operating Requirements.....	157
7.1.3 Support Function.....	158
7.2 Operation Design.....	158
7.3 Flow of Operation.....	158
7.4 Preparations.....	159
7.5 Operations.....	160
7.6 Backup Operation in Oracle Solaris Zones.....	161
7.6.1 Operation Design.....	161
7.6.2 Preparations.....	162
7.6.2.1 Configuring of Storage Cluster.....	163
7.6.2.2 Setting Up Source and Replica Volumes.....	163
7.6.3 Operations.....	163
7.6.4 Reconfiguration.....	164
7.6.5 Available Commands.....	164
7.6.6 Specifications of Settings File.....	165
7.7 Backup and Restoration of ZFS File Systems.....	165
7.7.1 Online Backup Using ZFS Snapshot with AdvancedCopy Manager Function.....	165
7.7.1.1 Preparations.....	168
7.7.1.2 Operation.....	170
7.7.2 Online Backup Using ZFS Snapshot with ZFS Uses zfs send and zfs receive Commands.....	174
7.7.2.1 Preparations.....	176
7.7.2.2 Operation.....	176
7.7.3 Offline Backup Using Export and Import Functions of ZFS (System Stop or Disconnection of Storage Pool).....	177
7.7.3.1 Preparations.....	180
7.7.3.2 Operation.....	180
7.7.4 Maintenance.....	180
7.8 Oracle VM Server for SPARC.....	181
7.8.1 Operation Design.....	181

7.8.2 Preparations.....	181
7.8.2.1 Configuring of Storage Cluster.....	182
7.8.2.2 Setting Up Source and Replica Volumes.....	182
7.8.3 Operation.....	182
7.9 Operating During Maintenance.....	183
7.9.1 Recovery from Hardware Failure.....	186
7.10 When Recovering Storage Cluster Continuous Copy Function.....	190
7.11 When Changing Configurations/Settings.....	192
7.11.1 Deleting Storage Cluster Environment.....	192
7.11.2 When Changing Storage Cluster Target Volume.....	193
7.12 Cascade Copies.....	193
Chapter 8 Operation Maintenance.....	195
8.1 Maintenance of Database.....	195
8.1.1 Making Individual Backups of Databases.....	196
8.1.1.1 Maintaining Backup Management List.....	196
8.1.1.1.1 Backing up Backup Management List.....	196
8.1.1.1.2 Restoring Backup Management List.....	197
8.1.1.2 Maintaining Replication Management List.....	197
8.1.1.2.1 Backing up Replication Management List.....	197
8.1.1.2.2 Restoring Replication Management List.....	197
8.1.1.3 Maintaining Database.....	197
8.1.1.3.1 Backing up Database.....	197
8.1.1.3.2 Restoring Database.....	198
8.1.1.4 Maintaining Registry.....	198
8.1.1.4.1 Backing Up Registry.....	198
8.1.1.4.2 Restoring Registry.....	198
8.1.2 Making Package Backup of Databases.....	199
8.1.2.1 Preliminaries.....	200
8.1.2.2 Backup Operations.....	202
8.1.2.3 Restoration Operations.....	204
8.2 Troubleshooting.....	207
8.2.1 Collecting Troubleshooting Information.....	207
8.2.1.1 Collecting Error Survey Information.....	207
8.3 Troubleshooting: Hardware Error During Backup.....	207
8.3.1 Overview.....	208
8.3.2 Troubleshooting.....	208
8.3.2.1 Hardware Error on Transaction Volume.....	209
8.3.2.2 Hardware Error on Backup Volume.....	210
8.3.2.3 Troubleshooting If Bad Sector Occurred in Transaction Volume.....	210
8.4 Troubleshooting: Hardware or Other Error During Replication.....	211
8.4.1 Overview.....	211
8.4.2 Troubleshooting.....	211
8.4.2.1 Hardware Error on Replication Volume.....	213
8.4.2.2 Troubleshooting If Bad Sector Occurred in Copy Source Volume.....	213
8.4.2.3 Troubleshooting When Lack of Free Physical Space Has Occurred in Copy Destination Volume.....	214
8.4.2.4 Error (halt) on Remote Copy Processing.....	215
8.5 Changing Operating Environment.....	216
8.5.1 Changing IP Address of Management Server.....	216
8.5.2 Changing IP Address of Managed Server.....	217
8.5.3 Changing Port Number of Management Server or Managed Server.....	218
8.5.4 Changing Server Name of Management Server or Managed Server.....	218
8.5.5 Changing Device Configuration.....	219
8.5.5.1 Backup Management.....	219

8.5.5.2 Replication Management.....	219
8.5.6 Changing Configuration in Cluster Operation.....	220
8.5.6.1 Changing IP Address of Management Server Transaction.....	220
8.5.6.2 Changing IP Address of Managed Server Transaction.....	220
8.5.6.3 Changing Port Number of Management Server Transaction or Managed Server Transaction.....	220
8.5.6.4 Changing Server Name of Management Server Transaction or Managed Server Transaction.....	221
8.5.6.5 Changing Logical Node Name.....	221
8.5.6.6 Changing Transaction Name.....	222
8.5.6.7 Changing Shared Disk for Shared Data.....	223
8.5.7 Operation on Model Upgrade for ETERNUS Disk Storage System.....	225
8.5.8 Changing the SSL Version Settings for Communicating with Supported Devices.....	225
8.6 Action When Managed Server Information Cannot Be Deleted.....	225
8.7 Maintaining Database of ETERNUS SF Manager.....	227
8.7.1 Saving Storage Cruiser's Manager Environment.....	227
8.7.2 Maintaining Common Control and AdvancedCopy Manager Repository.....	227
8.7.2.1 Backup of Common Control and AdvancedCopy Manager Repository.....	227
8.7.2.2 Restore of Common Control and AdvancedCopy Manager Repository.....	228
8.7.3 Maintaining License Management Database.....	229
8.7.3.1 Backup of License Management Database.....	229
8.7.3.2 Restore of License Management Database.....	229
8.7.4 Saving AdvancedCopy Manager CCM Environment Settings Files.....	229
Chapter 9 Command References.....	231
9.1 Backup Management Commands.....	231
9.1.1 Environment Definition Commands.....	235
9.1.1.1 swstsvrset (Managed Server Configuration Information Setting Command).....	235
9.1.1.2 swstsvrdisp (Managed Server Configuration Information Display Command).....	236
9.1.1.3 swstdevinfo (Device Information Setting Command).....	237
9.1.1.4 swstdevdisp (Device Usage Status Display Command).....	239
9.1.1.5 swstbkpolset (Backup Policy Setting Command).....	241
9.1.1.6 swstbkpoldisp (Backup Policy Display Command).....	243
9.1.1.7 swstbkpoldel (Backup Policy Deletion Command).....	244
9.1.2 Operation Commands.....	245
9.1.2.1 swstbackup (Backup Execution Command).....	245
9.1.2.2 swstbackstat (Backup Execution Status Display Command).....	248
9.1.2.3 swstrestore (Restore Execution Command).....	249
9.1.2.4 swstcancelrest (Restore Cancellation Command).....	253
9.1.2.5 swstreststat (Restore Execution Status Display Command).....	255
9.1.2.6 swsthistdisp (History Information Display Command).....	257
9.1.2.7 swsthistdel (History Information Deletion Command).....	259
9.1.2.8 swststartsync (Backup Synchronous Processing Start Command).....	261
9.1.2.9 swstcancelsync (Backup Synchronous Processing Cancel Command).....	263
9.1.2.10 swstsyncstat (Backup Synchronous Processing Progress Display Command).....	265
9.1.2.11 swstexecstat (Execution Status Display Command).....	267
9.1.2.12 swstcanceltrk (Tracking Cancel Command).....	271
9.1.2.13 swstrkstat (Tracking Status Display Command).....	272
9.1.3 Maintenance Commands.....	274
9.1.3.1 swstsrsem (Resource Match Command).....	274
9.1.3.2 swstresback (Resource Backup Command).....	275
9.1.3.3 swstresrst (Resource Restore Command).....	276
9.2 Configuration Management Commands.....	278
9.2.1 stgcmmodnode (Management Server Information Change Command).....	280
9.2.2 stgxfwcmaddsrv (Server Information Addition Command).....	281
9.2.3 stgxfwcmmodsrv (Server Information Change Command).....	283

9.2.4 stgxfwcmdelsrv (Server Information Deletion Command)	286
9.2.5 stgxfwcmsetdev (Device Information Collection/Reflection Command)	287
9.2.6 stgxfwcmdeldev (Device Information Deletion Command)	290
9.2.7 stgxfwcmdispsrv (Server Information Display Command)	291
9.2.8 stgxfwcmdispdev (Device Information Display Command)	293
9.2.9 stgxfwcmdisppat (Partition Information Display Command)	296
9.2.10 stgenvdisp (Environment Information Display Command)	298
9.2.11 stgmgrinfoset (Management Information Package Backup Server Registration/Deletion Command)	303
9.2.12 stgmgrinfobkup (Management Information Package Backup Command)	304
9.2.13 stgmgrinfodisp (Management Information Package Backup Status Display Command)	306
9.2.14 stgxfwcmsetsdx (SDX Object Management Unit Configuration Command)	308
9.2.15 stgxfwcmsetmode (Information Collection Mode Configuration Command)	309
9.2.16 stgrecpcnv (Repository Data Migration Command)	309
9.2.17 stgforcedelsrv (Managed Server Forcibly Deletion Command)	310
9.3 Replication Management Commands	311
9.3.1 Environment Definition Commands	313
9.3.1.1 swsrpsetvol (Replication Volume Information Setting Command)	313
9.3.1.2 swsrpvolinfo (Replication Volume Information Display Command)	317
9.3.1.3 swsrpdelv (Replication Volume Information Deletion Command)	319
9.3.1.4 swsrprebuffstat (REC Buffer Information Display Command)	321
9.3.1.5 swsrprebuffset (REC Buffer Settings Change Command)	324
9.3.2 Operation Commands	327
9.3.2.1 swsrpstartsync (Synchronous Processing Start Command)	327
9.3.2.2 swsrpmake (Replication Creation Command)	333
9.3.2.3 swsrpstat (Operation Status Display Command)	342
9.3.2.4 swsrpcancel (Replication Cancellation Command)	348
9.3.2.5 swsrpchsnc (Synchronization Mode Change Command)	354
9.3.2.6 swsrprevsync (Synchronization Reverse Command)	356
9.3.3 Maintenance Commands	358
9.3.3.1 swsrprecoveres (Resource Adjustment Command)	358
9.4 Start/Stop Daemon Commands	360
9.4.1 Starting and Stopping Communication Daemon	360
9.4.2 Starting and Stopping AdvancedCopy Manager Daemons	360
9.5 SnapOPC/SnapOPC+ Support Commands	361
9.5.1 swstestupdate (Update Volume Measurement Command)	362
9.5.2 swstsdv (Snap Data Volume Operation/Reference Command)	363
9.6 Troubleshooting Data Collection Command	366
9.6.1 esfsnap (AdvancedCopy Manager's Manager Troubleshooting Data Collection Command)	366
9.6.2 acmcapture (AdvancedCopy Manager's Agent Troubleshooting Data Collection Command)	367
Chapter 10 Notes on Operations	369
10.1 Notes on Backup and Replication Operation	369
10.1.1 General Notes	369
10.1.1.1 Backup and Replication Exclusions	369
10.1.1.2 Consistency of Data on Volumes	369
10.1.1.3 Fetching Device Information	369
10.1.1.4 Running Software	369
10.1.1.5 Maximum Number of Simultaneously Executing Operations	369
10.1.1.6 Slice Containing VTOC	370
10.1.1.7 SynfinityFILE and PRIMECLUSTER GFS Local File Systems	370
10.1.1.8 SynfinityFILE/Global File System and PRIMECLUSTER GFS Shared File System	370
10.1.1.9 Mounting of Disks from Other Servers	370
10.1.1.10 Notes on Copy Operations Targeted at File Systems	370
10.1.1.11 Notes on Executing Logical Unit (Disk) Unit Copy	371

10.1.1.12 Notes on Executing Copy Between Different Operating Systems.....	371
10.1.1.13 Notes on Executing Copy from Logical Unit (Disk) to Slice (Partition) Using Replication Function.....	372
10.1.1.14 Operations in ETERNUS Multipath Driver Environment.....	372
10.1.1.15 Notes When Using -m Option.....	372
10.1.1.16 Notes When Using -t Option.....	372
10.1.1.17 Copy Processing Performance of Advanced Copy.....	373
10.1.1.18 Moving of VM Guest of VMware.....	373
10.1.2 Notes on Snapshot High-speed Backup.....	373
10.1.3 Notes on Synchronous High-speed Backup.....	375
10.1.4 Notes on Restoration.....	375
10.1.5 Notes on SDX Object Operations.....	375
10.1.5.1 Notes on SDX Object Operations (Backup Management).....	375
10.1.5.2 Notes on SDX Object Operations (Replication Management).....	378
10.1.6 Notes on Cluster Operation.....	381
10.1.7 Notes on Backup in Clustered System.....	382
10.1.7.1 Backup and Restoration During Stoppage of Cluster.....	384
10.1.8 Notes on Replication in Clustered System.....	386
10.1.8.1 Replication During Stoppage of Cluster.....	389
10.1.9 Use in Combination with ODX (Offloaded Data Transfer).....	391
10.1.10 Use in Combination with XCOPY (Extended Copy).....	392
Appendix A Pre-processing and Post-processing of Backup and Restoration.....	394
A.1 Overview.....	394
A.2 Pre-processing and Post-processing of Backup.....	394
A.2.1 Pre-processing of Backup.....	395
A.2.2 Post-processing of Backup.....	397
A.2.3 Customizing Shell Script.....	401
A.3 Pre-processing and Post-processing of Restoration.....	402
A.3.1 Pre-processing of Restoration.....	403
A.3.2 Post-processing of Restoration.....	405
A.3.3 Customizing Shell Script.....	408
Appendix B Sample Script for Backup Operations.....	410
B.1 Backup and Restoration of Transaction Volume.....	410
B.1.1 Backup Procedure.....	410
B.1.2 Restoration Procedure.....	410
B.2 Sample Scripts for Oracle.....	411
B.2.1 Overview of Sample Scripts.....	411
B.2.1.1 Collecting DB Information.....	411
B.2.1.2 File System Backup.....	413
B.2.1.3 Raw Device Backup.....	414
B.2.2 Applying Sample Shell Scripts.....	416
B.2.2.1 DB Information Collection.....	416
B.2.2.2 File System Backup.....	416
B.2.2.3 Raw Device Backup.....	416
B.2.3 Restore Method.....	417
B.2.3.1 File system.....	417
B.2.3.2 Raw Device.....	417
B.2.4 Notes on Using Sample Shell Scripts.....	417
Appendix C Pre-processing and Post-processing of Replication.....	419
C.1 Overview.....	419
C.2 Pre-processing and Post-processing of Replication.....	419
C.2.1 Pre-processing and Post-processing for File System.....	419

C.2.2 Pre-processing When Replication Is Executed.....	422
C.2.3 Post-processing When Replication Is Executed.....	426
C.2.4 Customizing Shell Script.....	432
Appendix D System Volume Backup Operation.....	435
D.1 Overview.....	435
D.2 Preparation.....	435
D.2.1 Registration Managed Server.....	436
D.2.2 Registration of Device Information.....	436
D.2.3 Settings of Replication Volume Information.....	436
D.3 Backing up System Volume.....	437
D.4 Restoring System Volume.....	438
Appendix E Backup Operations Using Power-Saving.....	440
E.1 Environment Settings.....	440
E.1.1 Settings for ETERNUS Disk Storage System.....	440
E.1.2 Settings for AdvancedCopy Manager.....	443
E.1.3 Settings for Storage Cruiser.....	443
E.2 Operation.....	443
E.2.1 Start Disk Drive.....	443
E.2.2 Operate AdvancedCopy Manager.....	443
E.2.3 Stop Disk Drive.....	443
Appendix F Operation in Oracle Solaris Zones.....	444
F.1 Overview.....	444
F.2 Configuration.....	444
F.3 Software Requirements.....	445
F.4 Available Functions.....	446
F.4.1 Functions of AdvancedCopy Manager.....	446
F.4.2 Cluster Operation.....	446
F.5 Supported Devices.....	447
F.6 Flow of Operations.....	447
F.7 Preparations.....	448
F.7.1 Setting up Non-global Zones.....	448
F.7.2 Starting Daemons.....	450
F.7.3 Starting Web Console.....	450
F.7.4 Registering Managed Server.....	450
F.7.5 Setting Environment for Non-global Zones.....	450
F.7.6 Setting of Virtual Environment Support Function of Managed Servers.....	451
F.7.7 Fetching Device Information from Managed Servers.....	451
F.7.8 Other Preparations.....	452
F.8 Operation.....	452
F.9 Changing Operating Environment.....	452
F.9.1 Changing IP Addresses of Managed Servers.....	452
F.9.2 Changing Port Numbers of Managed Servers.....	452
F.10 Available Commands.....	452
F.11 Specifications of Settings File.....	453
F.11.1 Host Information Settings File (vstxapi_hostacm.ini).....	453
Appendix G Backup and Restoration of ZFS File Systems.....	454
G.1 Overview.....	454
G.1.1 Units of Backing up.....	454
G.1.2 Backup Methods.....	455
G.1.3 Functions of AdvancedCopy Manager.....	459
G.1.4 Software Requirements.....	460

G.2 Online Backup Using ZFS Snapshot with AdvancedCopy Manager Function.....	460
G.2.1 Operation Design.....	460
G.2.2 System Configuration.....	462
G.2.3 Preparations.....	464
G.2.4 Operation.....	467
G.3 Online Backup Using ZFS Snapshot with ZFS Uses zfs send and zfs receive Commands.....	472
G.3.1 Operation Design.....	472
G.3.2 System Configuration.....	474
G.3.3 Preparations.....	477
G.3.4 Operation.....	479
G.4 Offline Backup.....	488
G.4.1 Operation Design.....	488
G.4.2 System Configuration.....	492
G.4.3 Preparations.....	494
G.4.4 Operation.....	497
G.5 Maintenance.....	503
G.5.1 Changing Devices in Storage Pool Configuration.....	503
Appendix H Operation in Oracle VM Server for SPARC.....	504
H.1 Overview.....	504
H.2 Software Requirements.....	504
H.3 Available Functions.....	504
H.3.1 Functions of AdvancedCopy Manager.....	505
H.3.2 Cluster Operation.....	505
H.4 Supported Devices.....	505
H.4.1 Supported Devices for Control Domain and I/O Root Domain.....	505
H.4.2 Supported Device for Guest Domain.....	505
H.5 Operation.....	506

Chapter 1 Overview of AdvancedCopy Manager

1.1 Product Configuration of AdvancedCopy Manager

1.1.1 Functions

AdvancedCopy Manager consists of the following functions.

Table 1.1 Functions

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



Note

The Manager program has also the Agent function.

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

Managed Server

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

Management Server


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

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

- When only one Managed Server is installed

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

Figure 1.1 When Only One Managed Server Is Installed



- When two or more Managed Servers are installed

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

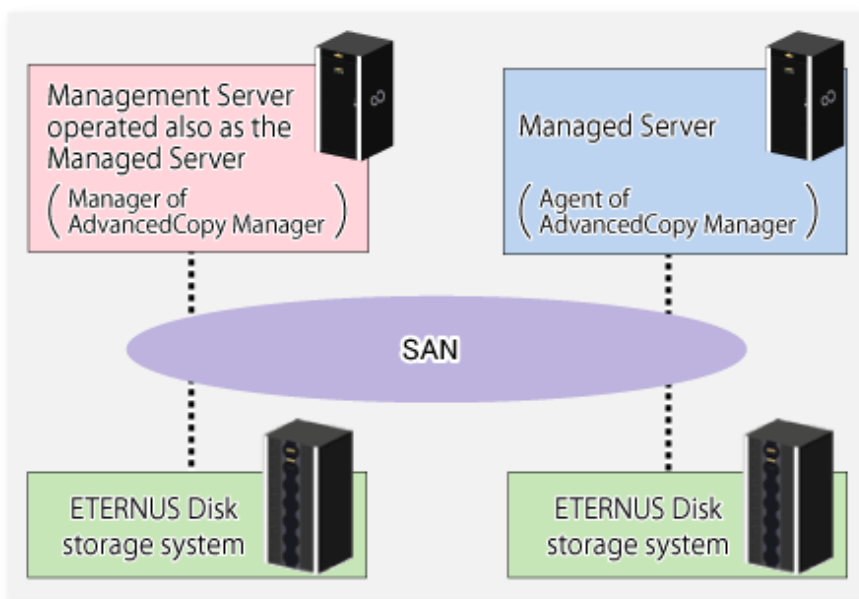
- When a Managed Server is operated as a Management Server

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

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

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


Figure 1.2 When Managed Server Is Also Operated as Management Server



- When setting up a dedicated Management Server

This setup comprises two or more Managed Servers with a dedicated Management Server.


Figure 1.3 When Dedicated Management Server Manages Managed Servers



1.1.2 System Configuration

The system configuration of AdvancedCopy Manager is shown below.

Figure 1.4 System Configuration of Management Server and Multiple Managed Servers



See

Refer to "Notes on Combining Different Versions" in the *ETERNUS SF Release Guide* for the supported combinations of the manager and the agent.

1.2 Function Overview of AdvancedCopy Manager

1.2.1 What Is Backup Function?

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

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

High-speed Backup Operation

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

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

Backup Operation in Conjunction with Oracle (Sample Shell Scripts)

The sample shell scripts provide AdvancedCopy Manager commands and Oracle operations combined together.

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

Policy Management Function

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

Storage Configuration Management Function

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

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

Backup and Restore Functions

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

History Information Management Function

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



Note

In a backup operation, the Advanced Copy function for Storage Cluster-enabled volumes (hereinafter referred to as "TFOV") is not available. When using a TFOV with the Advanced Copy function, use the Replication function of AdvancedCopy Manager or AdvancedCopy Manager Copy Control Module.

Refer to "[1.2.2 What Is Replication Function?](#)" or "Supported Functions" in "Overview of AdvancedCopy Manager CCM" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module* for details.

1.2.2 What Is Replication Function?

The replication function of AdvancedCopy Manager provides operation management of replications by using the Advanced Copy function of ETERNUS Disk storage system.

As a result, the following replication operations become possible.

High-speed Replication Operation

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

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

Replication Operation in Conjunction with Oracle (Sample Shell Scripts)

The sample shell scripts provide AdvancedCopy Manager commands and Oracle operations combined together.

Main Differences with Backup Function

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

Storage Cluster Continuous Copy Function

In a Storage Cluster environment, if a failover occurs in the Primary Storage while executing Advanced Copy, this function can continue the copy operation with the Secondary Storage by automatically taking over the Advanced Copy process.

Refer to "[Chapter 7 Storage Cluster Continuous Copy Function](#)" for details.

1.2.3 Web Console

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

- Security management
Operational restrictions of users of AdvancedCopy Manager can be performed using the GUI interface.
- Backup management
AdvancedCopy Manager Backup management tasks (such as definition of a server or a device, backup, restoration, etc.) can be performed using the GUI interface.
- Replication management
AdvancedCopy Manager replication management tasks (such as defining volumes and executing replications) can be performed using the GUI interface.

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

1.2.4 Copy Control Module

When the AdvancedCopy Manager Copy Control Module is used, it is possible to perform back ups from the server where the Copy Control Module is installed, using the Advanced Copy function of ETERNUS Disk storage system.

As a result, the AdvancedCopy Manager's agent need not be installed.

For details, refer to the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module*.

1.3 Operating Environment of AdvancedCopy Manager

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

For the hardware requirements, refer to "Operating Environment" in the *ETERNUS SF Installation and Setup Guide*.

For browsers on which Web Console can run, refer to "Supported Web Browsers" in the *ETERNUS SF Web Console Guide*.

Point

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

1.3.1 Software Requirement of AdvancedCopy Manager's Manager

The following table lists the software requirements of AdvancedCopy Manager's manager.

For about the supported operating systems, refer to "Operating Environment of ETERNUS SF Manager" in the *ETERNUS SF Installation and Setup Guide*.

However, for information on the supported operating systems of Oracle Solaris Zones, refer to "[F.3 Software Requirements](#)".

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

Table 1.2 Software Requirements of AdvancedCopy Manager's Manager

Product Name	Remarks
ETERNUS Multipath Driver	Required if all of the following conditions are fulfilled: <ul style="list-style-type: none">- A Management Server operates as a Managed Server.- ETERNUS Multipath Driver carries out multipath disk control and load distribution of a path is carried out further.
PRIMECLUSTER GD Snapshot (4.1A40 or later)	When a Management Server operates as a Managed Server, and when the backup operation or replication operation is performed in linkage with PRIMECLUSTER GD Snapshot.

Product Name	Remarks
VERITAS Volume Manager 4.0 VERITAS Volume Manager 4.1 VERITAS Volume Manager 5.0 VERITAS Volume Manager 5.1 VERITAS Volume Manager 6.0	When a Management Server operates as a Managed Server, and when a logical volume of VERITAS Volume Manager is the target of the operation performed.
Oracle 8i later	When a Management Server operates as a Managed Server, it is necessary to link it to Oracle.

Note

- For Solaris 11 operating system:
VERITAS Volume Manager is not supported.
- For multi-path drivers:
Use the ETERNUS Multipath Drivers for synchronous replication operations that target areas that include VTOC.

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

1.3.2 Software Requirement of AdvancedCopy Manager's Agent

The following table lists the software requirements of AdvancedCopy Manager's agent.

For about the supported operating systems, refer to "Operating Environment of AdvancedCopy Manager's Agent" in the *ETERNUS SF Installation and Setup Guide*.

However, for information on the supported operating systems of Oracle Solaris Zones, refer to "[F.3 Software Requirements](#)".

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

Table 1.3 Software Requirements of AdvancedCopy Manager's Agent

Product Name	Remarks
ETERNUS Multipath Driver	Required if all of the following conditions are fulfilled: <ul style="list-style-type: none"> - A Management Server OS is Solaris 11 or earlier. - ETERNUS Multipath Driver carries out multipath disk control and load distribution of a path is carried out further.
PRIMECLUSTER GD Snapshot (4.1A30 or later)	When the backup operation or replication operation is performed in linkage with PRIMECLUSTER GD Snapshot.
VERITAS Volume Manager 4.0 VERITAS Volume Manager 4.1 VERITAS Volume Manager 5.0 VERITAS Volume Manager 5.1 VERITAS Volume Manager 6.0	When a logical volume of VERITAS Volume Manager is the target of the operation performed.
Oracle 8i later	It is necessary to link to Oracle.

Note

- For Solaris 11 operating system:
VERITAS Cluster Server and VERITAS Volume Manager are not supported.

- For multi-path drivers:

Use the ETERNUS Multipath Drivers for synchronous replication operations that target areas that include VTOC.

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

1.4 Managing Device on AdvancedCopy Manager

This section describes how to manage the devices with AdvancedCopy Manager.

Refer to "3.2.2 Selecting Backup Source" for devices that can be used for the backup operation and to "6.2.2 Determining Replication Target" for devices that can be used for the replication operation.

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

For the supported device on Oracle Solaris Zones, refer to "F.5 Supported Devices".

1.4.1 Normal Devices

The following devices are supported in backup management and replication management.

- sd driver resources
- hddv driver resources
- MPxIO driver resources
- MPLB resources

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



- If Logical Units (disk units) are used for operations, refer to "10.1.1.11 Notes on Executing Logical Unit (Disk) Unit Copy".
- Under Solaris 10 or Solaris 11, if operations are performed under an ETERNUS Multipath Driver (hereafter referred to as ETMPD) environment, some functions provided by AdvancedCopy Manager may not be able to be used, depending on the ETMPD environment settings.
Refer to "10.1.1.14 Operations in ETERNUS Multipath Driver Environment" under "Chapter 10 Notes on Operations" for details.

1.4.2 SDX Objects of PRIMECLUSTER GD/SynfinityDISK

PRIMECLUSTER GD/SynfinityDISK is software whose main function is mirroring disk units.

AdvancedCopy Manager can backup and replicate the SDX objects of PRIMECLUSTER GD/SynfinityDISK.

SDX objects are virtual resources managed by PRIMECLUSTER GD/SynfinityDISK, and they include classes, groups, disks, volumes, and slices.

AdvancedCopy Manager can backup and replicate SDX objects in units of logical volumes or slices.

The format used to specify a device name in AdvancedCopy Manager is different depending on whether units of logical volumes or units of slices are to be used. Specify a device name as follows:

Logical Volume Unit Operations

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

Slice Unit Operations

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

Example:

The class name is class1, the volume name is volume1, and the DEVNAM value of sdxinfo is c1t0d10.

Logical Volume Unit Operations

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

Slice Unit Operations

```
/dev/sfdsk/class1/dsk/volume1:c1t0d10
```



Note

- PRIMECLUSTER GD and PRIMECLUSTER GD Snapshot are required to perform separate operations for each logical volume.
- Operation in units of logical volumes and operation in units of slices at the same time on a single Managed Server is not possible. Note that for cluster operations, operating in units of logical volumes or in units of slices is possible in separate Managed Server transactions.
- In an operation of a logical volume unit, the Storage Cluster Continuous Copy function is not available. When using the Storage Cluster Continuous Copy function, use an operation of a slice unit.

For details on backup of SDX objects, refer to ["3.8 SDX Object Operations"](#). For details on replication of SDX objects, refer to ["6.8 SDX Object Replication"](#).

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

1.4.3 Volumes under Control of VERITAS Volume Manager (VxVM)

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

- Physical slices in a logical volume
- Physical disk that constitutes a volume group

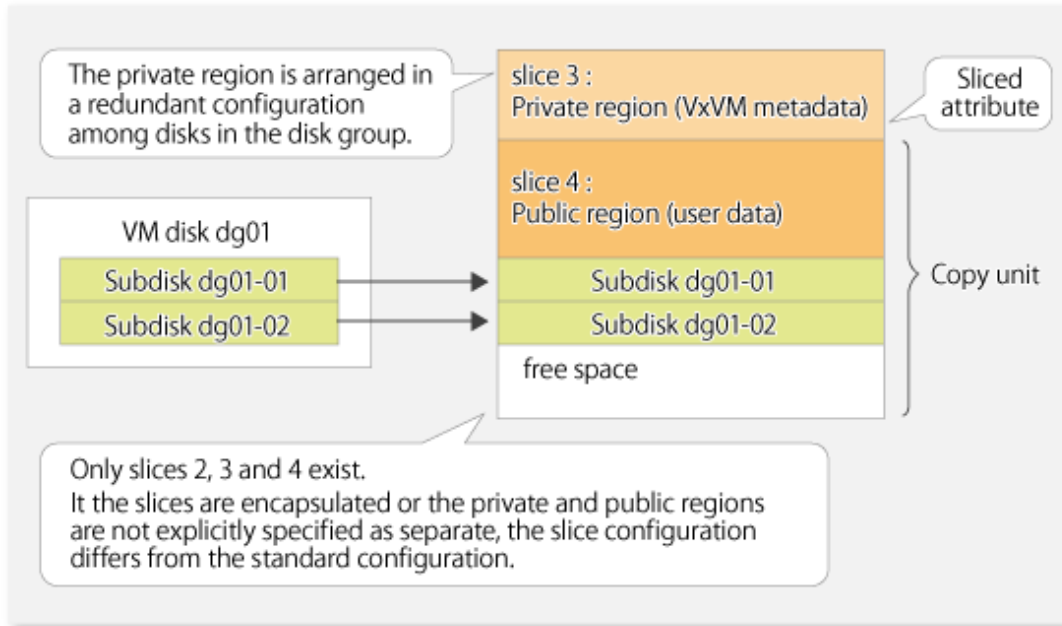
For details on backup of a subordinate volume of VxVM, refer to the section titled ["3.9 VxVM Volume Operation"](#). For details on replication of a subordinate volume of VxVM, refer to the section titled ["6.9 VxVM Volume Operation."](#)

1.4.3.1 Logical Volume

AdvancedCopy Manager copies a logical volume by copying the entire public region of the VM disk (physical disk) having the sliced attribute (the private region is separate from the public region) instead of copying in units of sub-disks.

Specify a device name as the VxVM logical volume name.

Figure 1.5 VM Disk (Physical Disk) Having Sliced Attribute (Private Region Is Separate from Public Region)




The above means that AdvancedCopy Manager copies the VxVM logical volume in units of physical slices (/dev/dsk/cXtXdXsX) instead of units of logical volumes.

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

- All sub-disks making up the logical volume are in one VM disk, and the VM disk does not have any sub-disk of another logical volume.
- The VM disk has the sliced attribute (the private region and public region are located in different slices).
- It is not a non-layered volume.
- The logical volume is not encapsulated.

Figure 1.6 Relationships Between Volumes, VM Disks and Physical Disks (1)



A configuration with more than one logical volume in a single VM disk (the logical volumes and VM disk are defined as N:1 (N>1)), such as shown below, is also supported but note the following points:

- For logical volumes with a constructed file system, the unmount/mount operation must be performed before and after backup or replication.
- For backup of an Oracle database, all logical volumes of the VM disk must be table areas.

Figure 1.7 Relationships Between Volumes, VM Disks and Physical Disks (2)


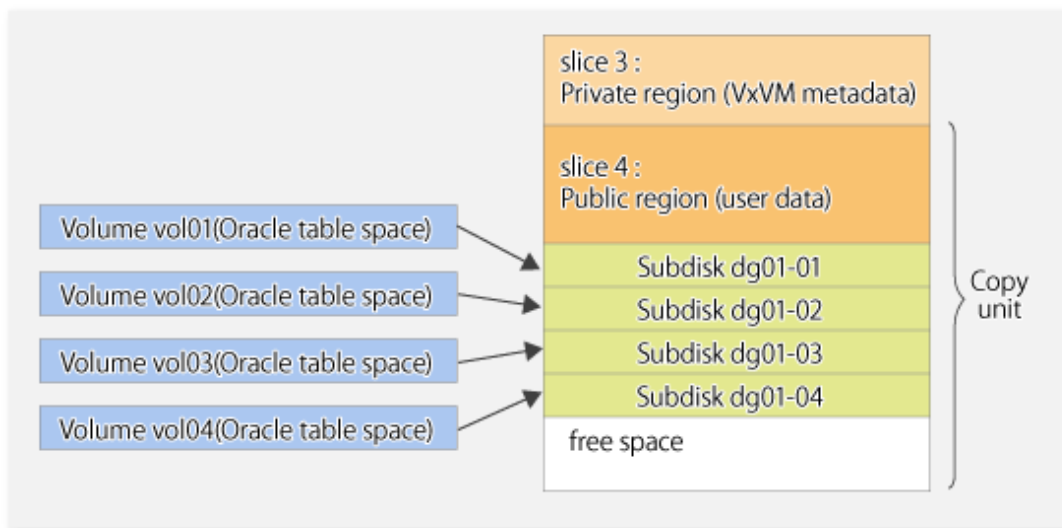


Figure 1.8 Relationship Between Volumes and Subdisks



AdvancedCopy Manager does not support logical volumes whose configuration is:

- Concatenation with spanning (Concatenation without spanning is supported.)
- Striping (RAID0)
- Mirroring (RAID1)
- RAID5
- Mirroring plus Striping (RAID1 + RAID0)

- Striping plus Mirroring (RAID0 + RAID1)

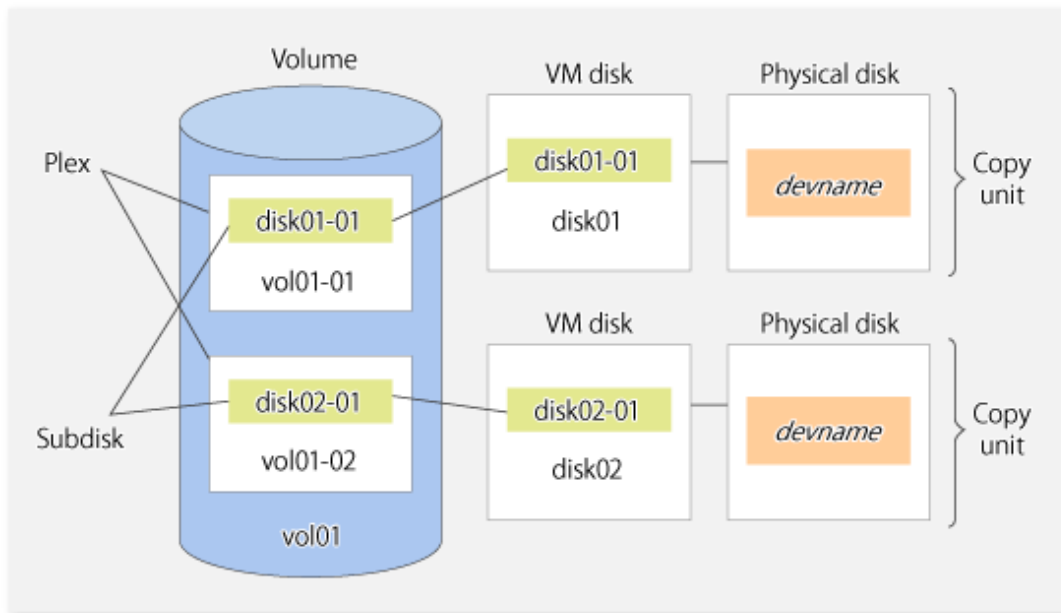
Note

- A disk group is created as a CDS disk group by default for VxVM 4.0. To obtain a disk group supported by AdvancedCopy Manager, create a non-CDS disk group, and specify "sliced" as the disk type.

1.4.3.2 Physical Disks That Constitute Disk Group

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

Figure 1.9 Example of VxVM Configuration



Since AdvancedCopy Manager copies an entire VM disk, not only the shared area that is a sub-disk but also the public region containing VxVM-internal configuration information is copied.

When backup or replication is performed, therefore, the integrity of VxVM configuration information must be maintained during copying.

Specify a device name as follows:

- /dev/vx/dmp/c#t#d#s#
- /dev/vx/dmp/c#t#d#
- /dev/vx/dmp/enclosure name (only for use under a SunCluster environment)

Note

Except when used under a SunCluster environment, operations are not supported under naming convention operations based on enclosures in VxVM.

Note

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

For operation in a SunCluster environment, the VxVM enclosure name must be the same at all servers that comprise the cluster, and the ETERNUS disk indicated by the enclosure name must be the same.



Note

It is only VM disk of the <nopriv> type that a physical slice becomes the unit of management.

1.5 Backup to Secondary Storage Media

Combining AdvancedCopy Manager with software having the backup feature to tape media (hereinafter referred to as "Tape backup software") can back up to secondary storage media the AdvancedCopy Manager copied data in an ETERNUS Disk storage system.

To back up data to secondary media, on completion of copy operation by AdvancedCopy Manager, back up the Advanced Copy destination area by the tape backup software.

In backup operation in combination of AdvancedCopy Manager with tape backup software, take note of the following matter to design:

- Copy operation by AdvancedCopy Manager

To ensure data consistency of backups, the Advanced Copy destination area should be kept out of access from other software than AdvancedCopy Manager. If copy operation is performed by AdvancedCopy Manager while the tape backup software is accessing or backing up the data in the Advanced Copy destination area, the copy operation fails.

Do not perform copy operation by AdvancedCopy Manager before keeping the Advance Copy destination area out of access from the tape backup software or making sure of that.

1.6 Operation with IPv6 Address

1.6.1 Software Requirements

The following are the basic software requirements to support the operation using IPv6 address in AdvancedCopy Manager:

Server Type	Basic Software
Management Server	Solaris 10 Solaris 11 (11/11 or later)
Managed Server	Solaris 10 Solaris 11 (11/11 or later)



See

Refer to "Operating Environment of ETERNUS SF Manager" and "Operating Environment of AdvancedCopy Manager's Agent" in the *ETERNUS SF Installation and Setup Guide* for information on support requirements of each basic software.

1.6.2 Support Target

For the IP address of the following servers, the operation using IPv6 address is supported:

- Management Server (manager)
- Managed Server (agent)



Note

In operation with IPv6 address, only a global address is supported.

1.6.3 Combination in IPv4/IPv6 Addresses Coexisting Environment

When using AdvancedCopy Manager's manager and agent in an IPv4/IPv6 addresses coexisting environment, the following combinations are possible:

Table 1.4 Combination in IPv4/IPv6 Addresses Coexisting Environment

Management Server's IP Address Environment	Managed Server's IP Address Environment		
	IPv4	IPv6	IPv4/IPv6
IPv4	A (Connect with IPv4)	N/A	A (Connect with IPv4)
IPv6	N/A	A (Connect with IPv6)	A (Connect with IPv6)
IPv4/IPv6	A (Connect with IPv4)	A (Connect with IPv6)	A (Connect with IPv4/IPv6(*1))

A: available; N/A: not available

*1: Connect with either IPv4 or IPv6 (IP address that the user registered on AdvancedCopy Manager).

Point


If both IPv4 and IPv6 format IP addresses are set to the operating system and AdvancedCopy Manager in the Management Server, a Managed Server with IPv4 address and a Managed Server with IPv6 address can be managed by one Management Server. Refer to "[stgxfwcmmodsv](#)" for information on setting method for IPv4 and IPv6 format IP addresses to AdvancedCopy Manager.

Note

- Since communication fails between a Managed Server with IPv4 and a Managed Server with IPv6, the following functions are not available in the IPv4/IPv6 addresses coexisting environment:
 - Inter-server replication between the Managed Server with IPv4 and the Managed Server with IPv6.
 - Management information batch backup with a Managed Server taken as a management information batch backup server.
- Even if both IPv4 and IPv6 format IP addresses are set to the operating system, the IP address of a Managed Server allowed to be registered on AdvancedCopy Manager is either IPv4 or IPv6.

If a Managed Server with IPv4 IP address and a Managed Server with IPv6 IP address are registered on AdvancedCopy Manager with IPv6 address, the range in which communication is allowed is as follows:

Figure 1.10 Communication Failure in IPv4/IPv6 Addresses Coexisting Environment



1.6.4 Combination in Environment Where Earlier Version Number and IPv6 Unsupported Operating System Coexist

The earlier version number and the platform that does not support IPv6 supports the connection with IPv4.

When using AdvancedCopy Manager's manager and agent at a different version number, the following combinations are possible:

Table 1.5 Combinations on Operating System in Which Earlier Version and IPv6 Are Not Supported

Manager Version No.	Agent Version No.		
	15.2 or Earlier	15.3 or Later (IPv6 Unsupported Operating System)	15.3 or Later (IPv6 Supported Operating System)
15.2 or earlier	A (Connect with IPv4)	N/A	N/A
15.3 or later (IPv6 unsupported operating system)	A (Connect with IPv4)	A (Connect with IPv4)	A (Connect with IPv4)
15.3 or later (IPv6 supported operating system)	A (Connect with IPv4)	A (Connect with IPv4)	A (Connect with IPv4/ IPv6(*1))

A: available; N/A: not available

*1: Connect with either IPv4 or IPv6 (IP address that the user registered on AdvancedCopy Manager).

1.6.5 IPv6 Address Notation

In AdvancedCopy Manager, IPv6 address can be input in full notation and abbreviated notation. However, the following notation is not available:

- IPv4-mapped address
- IPv4-compatible address
- Format in which an interface number or name is added as "%" to the end of an IP address (e.g. fe80::1234%fx0)



Note

AdvancedCopy Manager displays IPv6 addresses with reference to RFC5952. For this reason, the input IP address notation may be different from the IP address notation displayed by command output and on Web Console.

Chapter 2 Starting and Stopping Daemons

This chapter explains how to activate and stop daemons provided by AdvancedCopy Manager. For more information on how to activate and stop the daemons described in this chapter, refer to "9.4 Start/Stop Daemon Commands".

2.1 Daemons of AdvancedCopy Manager

The Manager is configured with the following daemon.

Table 2.1 Daemons of Manager

Daemon Type	Description
Communication daemon	Communicates with the Managed Server

The Agent is configured with the following daemon.

Table 2.2 Daemons of Agent

Daemon Type	Description
Communication daemon	Communicates with the Management Server and other Managed Servers

When the Management Server is in a clustered system, the following daemon composes Management Server transactions.

Table 2.3 Daemons Used for Management Server Transactions

Daemon Type	Description
Transaction-intended communication daemon	Communicates with the Managed Server

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

Table 2.4 Daemons Used for Managed Server Transactions

Daemon Type	Description
Transaction-intended communication daemon	Communicates with the Management Server and other Managed Servers



Note

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

- Communication daemon (local daemon)

2.2 Starting Daemons

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

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

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



Note

- Before executing the startacm command, confirm that the SWSTGNODE environment variable is not defined.
- Confirm that /var/tmp directory as a work directory exists.

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

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

To Start Daemon Separately

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

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

1. Execute the following commands depending on the communication daemon to be operated and check if it is activated. If it is activated, step 2 and subsequent steps are not necessary to be performed.
 - a. When the communication daemon is in the global zone and operation uses Oracle Solaris Zones Support Function.

```
# ps -z globale -f | grep stgxfws
```

- b. When the target communication daemon is other than "a" described above.

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

Confirm the command execution. If the strings are displayed as below, the daemon(s) has already been activated.

```
/opt/FJSVswstf/bin/stgxfws logicalNodeName
```

2. For the environment variable SWSTGNODE, specify the name of the logical node that the transaction is to be performed on.
3. Start the communication daemon.

For the startup command, refer to "[9.4.1 Starting and Stopping Communication Daemon](#)".

2.3 Stopping Daemons

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

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

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



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

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

To Stop Daemon Separately

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

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

1. Execute the following commands depending on the communication daemon to be operated and check if it is activated. If it is not activated, step 2 and subsequent steps are not necessary to be performed.
 - a. When the communication daemon is in the global zone and operation uses Oracle Solaris Zones Support Function.

```
# ps -z globale -f | grep stgxfws
```

b. When the target communication daemon is other than "a" described above.

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

Confirm the command execution. If the strings are displayed as below, the daemon(s) is activated.

```
/opt/FJSVswstf/bin/stgxfws logicalNodeName
```

2. For the environment variable SWSTGNODE, specify the name of the logical node that the transaction is to be performed on.
3. Stop the communication daemon.

For the stop command, refer to "9.4.1 Starting and Stopping Communication Daemon".

 **Point**

Terminating the communication daemon, may take up to 10 seconds to complete successfully.

Chapter 3 Backup Operation

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

For information on the relationships between the configurations of supported devices and the copy operation, refer to "1.4 Managing Device on AdvancedCopy Manager".

Note

For backup operations, the Advanced Copy function enabled for TFOVs is not available.

3.1 Overview

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

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

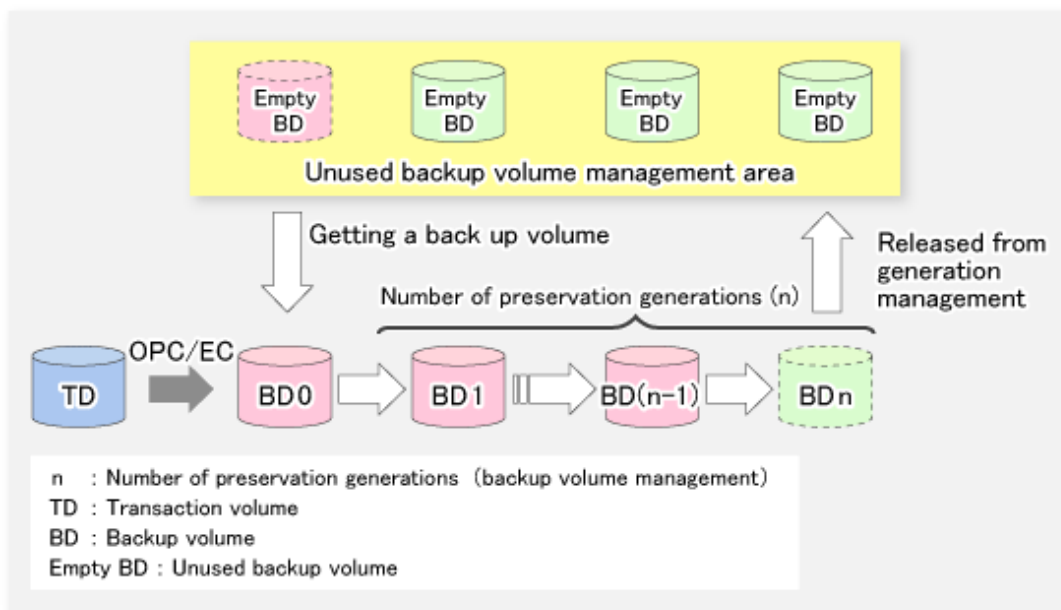
Automatic snapshot backup can be scheduled without first estimating the time required, because it does not use synchronous backup processing. However, when the save number is 1, a state may occur in which there is no backup data because a snapshot backup deletes the history prior to copy processing.

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

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

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

Figure 3.1 Managing Backup Volumes




3.1.1 Snapshot Backup Processing

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

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

1. When the `swstbackup` command is executed ((1) in the following figure), an unused backup volume with the same size as the relevant transaction volume is selected from the unused backup volume management area.
If the number of backup volumes exceeds the number of preservation generations by performing backup, the history of the oldest backup volume is deleted, and then the volume is selected as an unused backup volume.
2. Data on the transaction volume is copied to the selected unused backup volume. The logical copy is instantly performed and the physical copy is started on the hardware ((2) to (4) in the following figure).
3. Backup carrier information is set.

Figure 3.2 Snapshot Fast Backup Processing



The snapshot fast backup is completed as soon as you confirm the execution of the backup. The actual data copying is performed internally by the OPC function of ETERNUS Disk storage system.

The QuickOPC function of the snapshot high-speed backup creates a snapshot image by only copying the data updated since the start of previous backup. This reduces the time for physical copying significantly.

For information on the QuickOPC function, refer to "[Chapter 5 Backup Operation Using QuickOPC Function](#)".

Note

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

3.1.2 Synchronized High-speed Backup Processing

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

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


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

2. After the synchronous copy completes, the volumes enter equivalency maintain state which maintains the equivalency of the transaction volume and backup volume. ((3) in the following figure).
Note that the equivalency maintain state does not guarantee equivalency of the transaction volume and backup volume.
3. When the `swstbackup` command is executed ((4) in the following figure), the backup synchronization processing stops and the backup history information is registered.
4. When the generation number of the newly backed up data exceeds the number of generations of backup history being kept, the oldest backup volume is treated as an unused backup volume.

Note

The `swstbackup` command cannot be executed before the transaction and backup volumes have become equivalency maintain state.

Figure 3.3 Backup Using Synchronized High-speed Backup Function



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

For information on the Suspend/Resume function, refer to "[Chapter 4 Backup Operation Using Suspend/Resume Function](#)".

Note

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

3.1.3 Restoring Data

This section describes how data is restored in AdvancedCopy Manager.

The restoration operation refers to restoring data from a backup volume to a transaction volume using backup carrier management information.

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

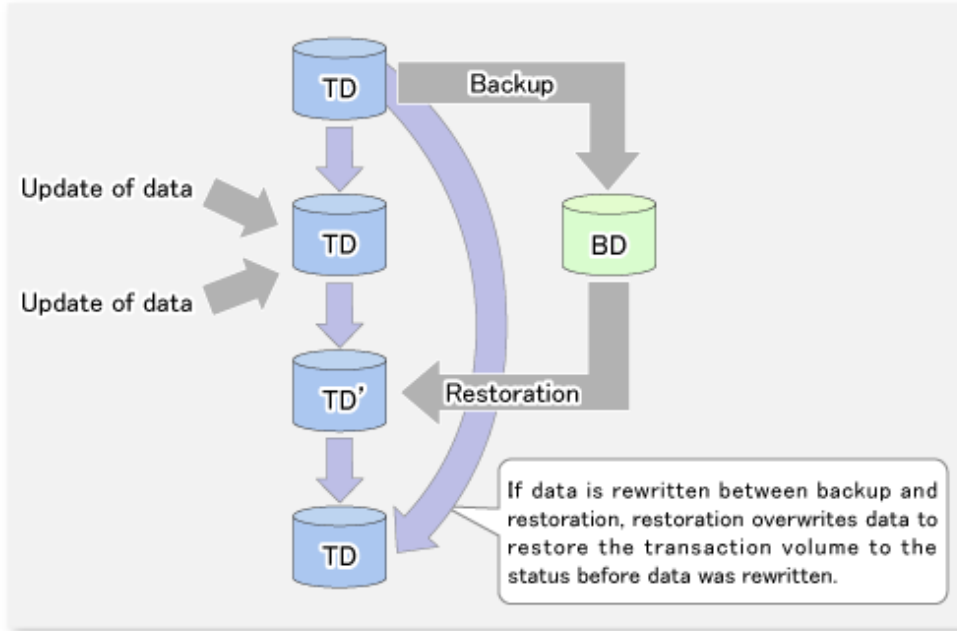
If the data on a backup volume under carrier management is restored to a transaction volume, the carrier management information does not change.

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

Note

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

Figure 3.4 Notes on Restore Operations in AdvancedCopy Manager



TD: Transaction volume
BD: Backup volume

3.2 Operation Design

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

- [3.2.1 Selecting Server on Which Backups Are Performed](#)
- [3.2.2 Selecting Backup Source](#)
- [3.2.3 Configuring Backup Policy](#)
- [3.2.4 Preparing Backup Volume](#)
- [3.2.5 Preparing Automatic Backup Operation](#)

3.2.1 Selecting Server on Which Backups Are Performed

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


Management Server

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

Managed Server

AdvancedCopy Manager operations can be performed on the Managed Server.

Figure 3.5 Configuration of Servers on Which Backup Operation Should Be Performed



3.2.2 Selecting Backup Source

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

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

For the supported units, refer to "1.4 Managing Device on AdvancedCopy Manager".

For information on setting up a transaction volume, refer to "3.4.6 Setting Operation Type for Device".

Note

- Do not back up the systems disk or slices on which AdvancedCopy Manager is installed.
- Slices must be allocated for transaction volumes in the following manner:
 - Create slices so that no VTOC is created.
 - No slice area may be duplicated on the same disk.
- For any slice that defines an entire volume, do not register it as a transaction volume.
- No file system of SynfinityFILE consisting of multiple slices and no PRIMECLUSTER GFS local file systems can be processed. When backing up or restoring a file system of SynfinityFILE consisting of multiple slices or a PRIMECLUSTER GFS local file system, unmount the multiple slices at the same time.

- Refer to "[10.1.1 General Notes](#)" for points to consider when deciding on a candidate volume for backup. Some points to consider include but may not be limited to:
 - [10.1.1.1 Backup and Replication Exclusions](#)
 - [10.1.1.6 Slice Containing VTOC](#)
 - [10.1.1.5 Maximum Number of Simultaneously Executing Operations](#)

3.2.3 Configuring Backup Policy

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

Number of Preservation Generations

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

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

When performing snapshot type high-speed backup using only one preservation generation, it is recommended that you also back up data to a tape.

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

Interval Days

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

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

For information on setting a backup policy, refer to "[3.4.7 Setting Backup Policy](#)".

3.2.4 Preparing Backup Volume

This section describes how to prepare the backup volume.

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

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

The backup volume size must be exactly the same number of bytes as that of the transaction volume.

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

Table 3.1 Number of Backup Volumes Required for Backup Operations

Backup Method	Number of Required Backup Volumes
Snapshot fast backup	Number of backup generations
Synchronized high-speed backup	Number of backup generations + 1

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

Table 3.2 Number of Backup Volumes Required for Backup Operations (If Backup Policy Has Been Set)

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

For information on configuring a backup volume, refer to "3.4.6 Setting Operation Type for Device".

Note

- Do not configure AdvancedCopy Manager to backup a system slice or a slice where AdvancedCopy Manager is installed.
- Slices must be allocated for backup volumes in the following manner:
 - Create slices without a VTOC.
 - No slice may be duplicated on the same disk.

In a backup volume, there is no need to create a slice for a whole volume, such as slice2.

Point

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

To do this, follow the procedure below:

1. Start backup synchronization.
2. Confirm the equivalency maintain status.
3. Execute backup.
4. Save the backup volumes to secondary media.
5. Delete history information.
6. Return to step 1.

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

3.2.5 Preparing Automatic Backup Operation


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

Automatic Snapshot Fast Backup

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

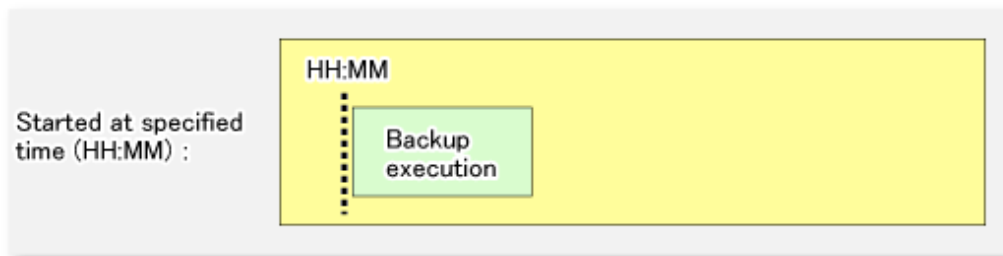
- Start the backup command as part of a batch job.

Figure 3.6 Start in Job Net



- Start the backup command at the specified time.

Figure 3.7 Start at Specified Time



Automatic Synchronized High-speed Backup

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

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

Starting Synchronized Backup Processing Start Command

Estimate the expected time required for the synchronized backup processing.

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


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

Starting Backup Execution Command

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

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

Figure 3.8 Start in Job Net




 Note

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

- Start the backup execution command with the time specified.


Figure 3.9 Start at Specified Time



3.3 Flow of Operation

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

Figure 3.10 Flow of Backup Operation on Transaction Volume



3.4 Preparing to Start Backup

3.4.1 Activating Daemons

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

For information on starting a daemon, refer to "[Chapter 2 Starting and Stopping Daemons](#)".

3.4.2 Starting Web Console

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

3.4.3 Registering Managed Server

With Web Console, register Managed Servers that are to be managed.



Note

- Change to the value of 800 characters or less before registering the Managed Server when the value of the environment variable PATH of the Management Server is more than 800 characters.
- Perform this task after confirming that it is possible to communicate from the Managed Server to the IP address of the Management Server. Check the Management Server IP address with the [stgxfwcmdisprv](#) command.

3.4.4 Fetching Device Information on Managed Server

This process can be carried out using either Web Console or the [stgxfwcmsetdev](#) command.

Refer to "Add/Reload Device" in the *ETERNUS SF Web Console Guide* for information on the operating procedures of Web Console.



Note

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

3.4.5 Configuring Environment for Backup Operation Server

This process can be executed using either Web Console or the [swstsvrset](#) command.



Point

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

3.4.6 Setting Operation Type for Device

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

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

Table 3.3 Number of Backup Volumes Required for Backup Operations

Backup Method	Number of Required Backup Volumes
Snapshot fast backup	Number of backup generations
Synchronized high-speed backup	Number of backup generations + 1

This process can be carried out using either Web Console or the [swstdevinfo](#) command.



Note

- For any slice that defines an entire volume, do not register it as a transaction volume.

- To change the configuration of a partition that has been registered as either a transaction volume or a backup volume, take the following steps:
 1. Delete the registered configuration, and then change the configuration.
 2. Attempt to "[3.4.4 Fetching Device Information on Managed Server](#)".
 3. Execute the `swstdevinfo` command to register the information.
- The backup volume must be the same size (have the same number of bytes) as the transaction volume.
- Backup and backup volumes are not supported for the system slice or the slice where AdvancedCopy Manager is installed.
- The mounted device cannot be set to backup volume. However, it is possible to set it to the transaction volume.
- Logical Units (disks) cannot be set for a transaction volume or a backup volume.

 **Point**

The registered volumes can be confirmed using either Web Console or the `swstdevdisp` command.

3.4.7 Setting Backup Policy

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

- Preservation generations

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

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

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

 **Note**

- Even if the number of interval days has been specified, AdvancedCopy Manager does not automatically invoke the backup when the date for the next backup is reached. This value is used only to prompt the user to perform backup when the number of backup interval days is exceeded.
- When you set a backup policy, there must be as many registered backup volumes as required to perform backup according to the specified backup policy. For information on the necessary number of backup volumes, refer to "[3.2.4 Preparing Backup Volume](#)".
- Even though the number of backup volumes required for the backup may not have been registered when backup policies for synchronized high-speed backup were set, this number can be set by default provided the number of backup volumes required for snapshot fast backup has been previously registered. Otherwise, it may not be possible to execute synchronized high-speed backup.

 **Point**

The registered backup policy can be confirmed using either Web Console or the `swstbkpoldisp` command.

3.4.8 Customizing Pre-processing and Post-processing

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

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

- For operational reasons, to avoid unmount and mount processing of transaction volumes that comprise a file system
- To add special pre-processing and post-processing to the corresponding scripts.

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

Point

When AdvancedCopy Manager Is Upgraded

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

3.4.9 Preparing Device Map File

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

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

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

Point

To manage multiple generations of backups using defined transaction volume to backup volume relationships, multiple device map files must be prepared.

The device map file used when backing up or starting synchronous processing should meet either of the following requirements:

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

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

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

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

3.4.9.1 Describing 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/c1t1d1s1          /dev/dsk/c1t1d20s1
/dev/dsk/c1t1d1s3          /dev/dsk/c1t1d20s3
/dev/dsk/c1t1d1s4          /dev/dsk/c1t1d20s4

# A comment can be described here also.
/dev/dsk/c1t1d2s1          /dev/dsk/c1t1d30s1
```



```
/dev/dsk/c1t1d2s3      /dev/dsk/c1t1d30s3
/dev/dsk/c1t1d2s4      /dev/dsk/c1t1d30s4

/dev/dsk/c1t1d5s4      /dev/dsk/c1t1d10s4
/dev/dsk/c1t1d5s5      /dev/dsk/c1t1d10s5
:
```

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

- Describe a transaction volume name and a corresponding target backup volume name in one line. Separate the transaction volume name and backup volume name with one or more space or tab characters. One or more space or tab characters may be used as delimiters between the beginning of a line and a transaction volume name or between the end of a target backup volume name and the end of a line (new line character).
- A blank line (Space or tab characters) may be included in a file.
- Anything after the symbol '#' is a comment. If the comment uses more than a single line, then each line must be preceded by a '#' character.
- You cannot specify more than one backup volume for the same transaction volume in one device map file. If you do, only the first occurrence is valid. While the device map file is still read, subsequent occurrences are ignored.
- A device map file may include descriptions of transaction volumes other than the processing target. Refer to the description example provided in the beginning of this chapter.

3.5 Operation

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

There are two kinds of backup in AdvancedCopy Manager:

- Snapshot type high-speed backup
- Synchronous high-speed backup

3.5.1 Performing Snapshot High-speed Backup

Snapshot high-speed backup can be carried out using either Web Console or the [swstbackup](#) command.

To check the backup status, use the [swstbackstat](#) command.



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

Enter the appropriate command for each operation.

- Execute the backup command with the -T option specified the [swstcanceltrk](#) command.

3.5.2 Performing Synchronized High-speed Backup

This process can be carried out using either Web Console or commands.

If operating with commands, perform the backup synchronous process, then carry out the backup for the transaction volume that you need to backup. For details on commands, refer to the following:

- Starting backup synchronous process: [swststartsync](#)
- Checking the status of backup synchronous process: [swstsyncstat](#)
- Executing backup: [swstbackup](#)

- Checking backup status: [swstbackstat](#)

Note

- The [swstbackup](#) command causes an error if it is executed before the transaction and backup volumes have entered equivalency maintain state.
- The device map cannot be specified while backup is processing.
- When the "[Chapter 4 Backup Operation Using Suspend/Resume Function](#)" is used for synchronous high-speed backup, the operations listed below cannot be performed from Web Console. Enter the appropriate command for operation.
 - If one transaction volume has multiple synchronous processing (ie, EC sessions), specify the backup volume and cancel only the specified EC session ([swstcancelsync](#) command).

3.5.3 Restoring Data

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

3.5.3.1 Restoring Entire Volume

All of the data from a volume that has been saved with the snapshot high-speed backup can be restored using either Web Console or the [swstrestore](#) command.

To check the restore status, use the [swstreststat](#) command.

Note

If restoring for a general file system, stop operations.

3.5.3.2 Restoring Individual File

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

1. Mount the backup volume. Check the backup volume using Web Console or the [swsthistdisp](#) command.
2. Copy files to restore using the cp command.
3. Unmount the backup volume.

3.5.4 Managing Backup History

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

Point

- Backup history information is only used to manage information stored on a backup volume.
- Any backup history information exceeding the number of preservation generations defined in a backup policy is automatically deleted.

3.5.4.1 Displaying Backup History

Backup history can be displayed using either Web Console or the [swsthistdisp](#) command.

3.5.4.2 Deleting Backup History

Backup history can be deleted using either Web Console or the [swsthistdel](#) command.

3.6 Changing Setting Information

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

3.6.1 Adding Device

After a device is added to any Managed Server, incorporate it into the backup as described below. This process can be carried out using either Web Console or commands.

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

1. Execute the `stgxfwcmsetdev` command to acquire information for the newly added device, and accept management information.
2. Execute the `swstdevinfofet` command to set up the operation type of the added device.
3. If the added device was defined as a transaction volume, execute the `swstbkpolset` command to set up the backup policy.



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

Refer to "Required Resources for Operation" in "Operating Environment of AdvancedCopy Manager's Agent" in the *ETERNUS SF Installation and Setup Guide* for details.

3.6.2 Deleting Device

This section describes how to delete different types of devices.

3.6.2.1 Deleting Transaction Volume

This process can be carried out using either Web Console or commands.

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

1. Execute the `swsthistdel` command to delete the backup history.
2. Execute the `swstbkpoldel` command to delete the backup policy.
3. Execute the `swstdevinfofet` command to delete the device information for the deleted transaction volume.



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

3.6.2.2 Deleting Backup Volume

This process can be carried out using either Web Console or commands.

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

1. If deleting the backup volume, ensure that you first execute the `swstbkpoldisp` command to check for a previously defined backup policy.



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

2. Execute the `swsthistdel` command to delete the backup history, removing the relevant volume from the backup history. If there is no backup history, proceed to the next step.
3. Execute the `swstdevinforeset` command to delete the device information for the deleted backup volume.

 **Point**

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

3.6.3 Adding Managed Server

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

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

1. Register a Managed Server. For more information, refer to "[3.4.3 Registering Managed Server](#)".
2. Fetch all the information on the devices connected to the newly added Managed Server. For more information on this, refer to "[3.4.4 Fetching Device Information on Managed Server](#)".
3. Set the environment information for the Managed Server. For more information on this, refer to "[3.4.5 Configuring Environment for Backup Operation Server](#)".
4. Set the operation type for the device for which information has been fetched in step 2. For more information on this, refer to "[3.4.6 Setting Operation Type for Device](#)".
5. Configure a backup policy for all the transaction volumes registered in step 4. For more information on this, refer to "[3.4.7 Setting Backup Policy](#)".

 **Note**

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

Refer to "Required Resources for Operation" in "Operating Environment of ETERNUS SF Manager" in the *ETERNUS SF Installation and Setup Guide* for details.
.....

3.6.4 Deleting Managed Server

The process of deleting a Managed Server can be carried out using either Web Console or the `stgxfwcmdelsrv` command.

 **Point**

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

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

1. Delete the backup history of all transaction volumes on the Managed Server to be deleted. For more information, refer to "[3.5.4.2 Deleting Backup History](#)".
2. Delete the backup policy of all transaction volumes on the Managed Server to be deleted.
The process of deleting a backup policy can be carried out using either Web Console or the `swstbkpdel` command.
3. Delete the volume type that was registered using "[3.4.6 Setting Operation Type for Device](#)" for the device connected to the Managed Server.
4. Delete the Managed Server.

5. If you have deleted the Managed Server with the [stgxfwcmdelsrv](#) command, take the following actions to reflect to Web Console depending on how the agent installed in the deleted Managed Server is used:
 - If both of AdvancedCopy Manager's agent and Storage Cruiser's agent are used:
With Web Console, execute the **Reload Conf.** operation to update to the latest information.
 - If only AdvancedCopy Manager's agent is used:
With Web Console, delete the Managed Server that you have deleted.

3.7 Stopping Operation

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

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



- When all daemons on the Managed Server have stopped, all functions of AdvancedCopy Manager running on the Managed Server stop.
- Stop daemons on a Management Server only after ensuring that all the Managed Servers that it manages have also been stopped.

3.8 SDX Object Operations

This section describes the backup and restore operations of SDX objects.

A good understanding of the basic operation of normal volumes is required for these operations.

For details on PRIMECLUSTER GD, refer to the *PRIMECLUSTER Global Disk Services User's Guide (Solaris(TM) Operating Environment)*.

3.8.1 Units of SDX Object Operations

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

1. Backup operation in units of logical volumes (method with PRIMECLUSTER GD Snapshot)
2. Backup operation in units of slices (method without PRIMECLUSTER GD Snapshot)

For details on how to confirm and select the type of operation, refer to "[3.8.1.1 Configuring Unit of SDX Object Operations](#)".



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

3.8.1.1 Configuring Unit of SDX Object Operations

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

Units of Logical Volumes

When PRIMECLUSTER GD Snapshot has been included PRIMECLUSTER GD

Units of Slices


In cases other than the above

For operation in units of slices in an environment with PRIMECLUSTER GD installed, use the command described in the [stgxfwcmsetsdx](#) command, and then set operation in units of slices by taking action as described in "[3.4.4 Fetching Device Information on Managed Server](#)".

3.8.2 Backup Operation in Units of Logical Volumes (Method with PRIMECLUSTER GD Snapshot)

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

Figure 3.11 Backup Operation in Units of Logical Volumes



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

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



Note

Configuration and Conditions of Unsupported SDX Objects

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

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

Copying cannot be processed in the following circumstances:

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

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

Point

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

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

Volume Object Type Attributes

- single: Single volume (supported)
 - mirror: Mirror volume (supported)
 - stripe: Stripe volume (not supported)
 - concat: Volume in a concatenation group (not supported)
 - switch: Switch volume (not supported)
-

3.8.2.1 Designing Backup Operation

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

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

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

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

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

When configuring SDX objects of PRIMECLUSTER GD for use by AdvancedCopy Manager, pay careful attention to the relevant structure attributes of "[3.8.2.2 Transaction Volume](#)" and "[3.8.2.3 Backup Volumes](#)" described in the next sections.

For details on PRIMECLUSTER GD Snapshot and premises of the proxy structure, refer to the *PRIMECLUSTER Global Disk Services User's Guide (Solaris(TM) Operating Environment)*.

3.8.2.2 Transaction Volume

To use an SDX object (logical volume) for transactions, register it as a transaction volume.

The transaction volume operates as a master object of PRIMECLUSTER GD Snapshot.

Therefore, note the following about configuring the backup operation:

- An SDX object that belongs to a shadow class cannot be registered as the transaction volume.
- The only types of logical volumes that can be registered as the transaction volume are single volumes and mirror volumes.
- A stripe volume, a volume in a concatenation group, or a switch volume cannot be registered as the transaction volume. Stripe groups and concatenation groups, however, can be used as subordinate groups of a mirror group.
- A volume in a disk group to which a backup volume also belongs cannot be registered as the transaction volume.
- A logical volume that belongs to the single disk group with a backup volume cannot be registered as the transaction volume.
- If any of the slices that make up the transaction volume are separated temporarily or are being copied, backup cannot be performed.
- If the transaction volume (logical volume) is in the invalid state, backup cannot be performed.
- A proxy object of PRIMECLUSTER GD Snapshot cannot be used as the transaction volume.

3.8.2.3 Backup Volumes

Register an SDX object (logical volume) as a backup volume.

The backup volume operates as a proxy object of PRIMECLUSTER GD Snapshot.

Therefore, note the following about configuring the backup operation:

- An SDX object that belongs to a shadow class cannot be registered as a backup volume.
- The types of logical volumes that can be registered as backup volumes are single volumes and mirror volumes. A stripe volume, a volume in a concatenation group, or a switch volume cannot be registered as a backup volume. Stripe groups and concatenation groups, however, can be used as subordinate groups of a mirror group.
- A volume in a disk group to which a transaction volume also belongs cannot be registered as a backup volume.
- A logical volume that belongs to the single disk with the transaction volume cannot be registered as a backup volume.
- An object whose class is different from that of the transaction volume cannot be used as a backup volume.
- A master object of PRIMECLUSTER GD Snapshot cannot be used as a backup volume.
- A backup volume for backup by using EC or OPC must be a non-mirroring volume. If the backup volume is in a mirroring structure, only software-copy can be used.
- If a backup volume is in use or is accessible by any application, backup cannot be performed.
- If any slices of the slices that make up a backup volume is separated temporarily or being copied, backup cannot be performed.
- As a backup volume, a logical volume must be the same size as the transaction volume.
- For the backup operation, the total number of the slices that make up the transaction volume and slices that make up all backup volumes associated with the transaction volume must not exceed 32.

3.8.2.4 Copy Functions Available When Used with PRIMECLUSTER GD Snapshot

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

- Advanced copy: Copying with the hardware function of ETERNUS Disk storage system
- Software-copy: Copying with the disk driver of PRIMECLUSTER GD

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

Table 3.4 Copy Functions Available in Linkage with PRIMECLUSTER GD Snapshot (Backup)

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

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

Table 3.5 Copy Functions Available in Linkage with PRIMECLUSTER GD Snapshot (Restore)

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

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

3.8.2.5 Configuration Options

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

3.8.2.6 Fetching Device Information from Managed Server

To perform the backup operation with a Managed Server and obtain information of a subordinate device related to the Managed Server, you must register the server.

For details on how to fetch the device information, refer to "[3.4.4 Fetching Device Information on Managed Server](#)".

3.8.2.7 Configuring Operation Type for Device

Register a PRIMECLUSTER GD logical volume as a transaction volume or backup volume.

```
# /opt/FJSSVswsts/bin/swstdevinfofet -t /dev/sfdsk/CLS1/dsk/voll
swstdevinfofet completed
# /opt/FJSSVswsts/bin/swstdevinfofet -b /dev/sfdsk/CLS1/dsk/bvoll
swstdevinfofet completed
```



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

- Object of a shadow class
- Stripe-type object
- Concatenation-type object
- Switch-type object

3.8.2.8 Configuring Backup Policy

Checking Number of Backup Generations

Before configuring a backup policy, the required number of backup volumes for operation must be prepared according to the backup policy. (For details, refer to "[3.4.7 Setting Backup Policy](#)".)

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

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

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



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

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

Limit on Number of Slices

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

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

3.8.2.9 Backup

Backup in Units of Logical Volumes

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

AdvancedCopy Manager automatically selects it from the registered backup volumes.

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

Available Copy Processing Methods

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

EC/OPC cannot be used in the following cases:

- The backup volume is an SDX object in a mirroring structure
- A subordinate group is connected to the disk group to which the transaction volume and backup volume belong

Table 3.6 Hardware and Available Copy Processing Methods

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

3.8.2.10 Restore

Restoration to Transaction Volume in Units of Logical Volumes

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

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

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

Restoration with Restore Destination Specified

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

Available Copy Processing Methods

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

OPC cannot be used in the following case:

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

Table 3.7 Hardware and Available Copy Processing Methods

Hardware	Available Copy Processing Method
OPC can be used.	OPC, software-copy
OPC cannot be used.	Software-copy

3.8.2.11 Operational Notes

Do not use the "sdxproxy" command directly to change the state of a volume registered for backup management.

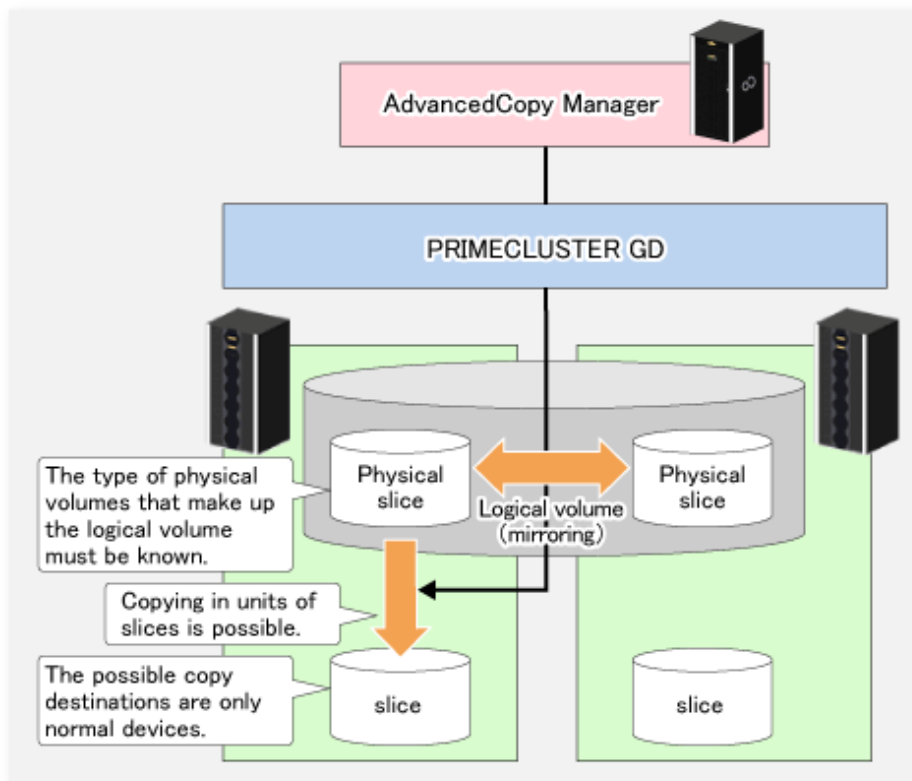
Use of the "sdxproxy" command may create inconsistencies between the backup management information and the actual volume state. In such cases, perform recovery by executing the [swstsrsemtch](#) command.

3.8.3 Backup Operation in Units of Slices (Method without PRIMECLUSTER GD Snapshot)

This backup operation uses only the AdvancedCopy Manager function without PRIMECLUSTER GD Snapshot.

For reliable configuration and operation, you need to know the type of physical volumes that make up the logical volume.

Figure 3.12 Backup Operation in Units of Slices



When specifying a name in an AdvancedCopy Manager command, combine the logical volume name and the AdvancedCopy Manager device name using the format shown below. Refer to "1.4.2 SDX Objects of PRIMECLUSTER GD/SynfinityDISK" for details.

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



Note

- For notes on SDX object operations, refer to "[10.1.5.1 Notes on SDX Object Operations \(Backup Management\)](#)", and "[10.1.5.2 Notes on SDX Object Operations \(Replication Management\)](#)".
- If mirroring between boxes is enabled and restoration by OPC is required in the event of a box failure, both mirrored systems must be backed up. In such cases, the required capacity of the backup volume is not the size of the logical volume but the size of the physical volume.
If a Symfoware DBSP is being mirrored within the ETERNUS Disk storage system, only one of the paired devices (mirrored device or mirroring device) can be registered as a transaction volume. Therefore, if a hardware error occurs on the device that carries out backup, backup and recovery can only be implemented after restoring the device from the hardware error.
- The SDX objects that can be used for operations in slice units are as follows:
 - Single volumes
 - Non-layered mirror volumes



Point

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

Volume Object Type Attribute

- single: Single volume (operation possible)
- mirror: Mirror volume (operation possible)
- stripe: Stripe volume (operation not possible)
- concat: Volume in a concatenation group (operation not possible)
- switch: Switch volume (operation not possible)

3.8.3.1 Configuring Backup Operation

For notes on configuring the backup operation, refer to "[10.1.5.1 Notes on SDX Object Operations \(Backup Management\)](#)".

3.8.3.2 Fetching Device Information from Managed Server

To perform the backup operation with a Managed Server and obtain information of a subordinate device of the Managed Server, you must register the server.

For details on how to fetch the device information, refer to "[3.4.4 Fetching Device Information on Managed Server](#)".

3.8.3.3 Configuring Operation Type for Device

Transaction Volume

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

```
#/opt/FJSVswsts/bin/swstdevinfoaset -t /dev/sfdsk/CLS01/dsk/VOL01:c1t0d1
swstdevinfoaset completed
#
```

Backup Volume

Slices of an SDX object cannot be registered as a backup volume.

General slices must be used for a backup volume.

```
#/opt/FJSVswsts/bin/swstdevinfo -b /dev/dsk/c1t0d2s6
swstdevinfo completed
#
```

3.8.3.4 Backup

Backup Example

```
#/opt/FJSVswsts/bin/swstbackup /dev/sfdsk/CLS01/dsk/VOL01:c1t0d1
/dev/sfdsk/CLS01/dsk/VOL01:c1t0d1 swstbackup completed
#
```

Backup States

Backup can only be performed when the SDX object that makes up the applicable logical volume is in one of the states listed below. If the object is in any other state, backup cannot be performed.

AdvancedCopy Manager checks the state of the SDX object by using the "sdxinfo" command of SafeDISK/PRIMECLUSTER GD.

- The state of the volume is "ACTIVE" or "STOP".
- The state of the physical disk is "ENABLE".
- The state of the slice is "ACTIVE" or "TEMP".

Backup Pre-processing and Post-processing

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

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

3.8.3.5 Restore

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

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

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

For details on restoration of an SDX object that is not a single volume, refer to "[10.1.5.1 Notes on SDX Object Operations \(Backup Management\)](#)".

Restoration Example

```
#/opt/FJSVswsts/bin/swstrestore /dev/sfdsk/CLS01/dsk/VOL01:c1t0d1
/dev/sfdsk/CLS01/dsk/VOL01:c1t0d1 swstrestore completed
#
```

3.9 VxVM Volume Operation

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

3.9.1 Backup Operation in Units of Logical Volumes

When the VxVM volume configuration satisfies the relevant conditions, backup operation can be performed in units of logical volumes.



Note

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

3.9.1.1 Configuring Backup Operation

3.9.1.1.1 Supported VxVM Volume Structures

For the supported VxVM volume structures, refer to "1.4 Managing Device on AdvancedCopy Manager".

3.9.1.1.2 Preparing Backup Volume

The following types of volumes only can be used as backup volumes of VxVM logical volumes:

- Logical volume of VxVM
- Normal device (slice)

For this backup volume, however, use a VxVM logical volume with the same structure as that of the transaction volume, in anticipation of mounting of the backup volume or its use by other applications.

The same volume structure means the following conditions are satisfied:

- The size of the physical disk (LUN) is the same.
- The size of the private region is the same.
- The size of the public region is the same.
- Each sub-disk in the public region has a corresponding start offset and size.
- The logical volume and sub-disks have a corresponding association.

Figure 3.13 Example of Correspondence Between Transaction Volume and Backup Volume (1)

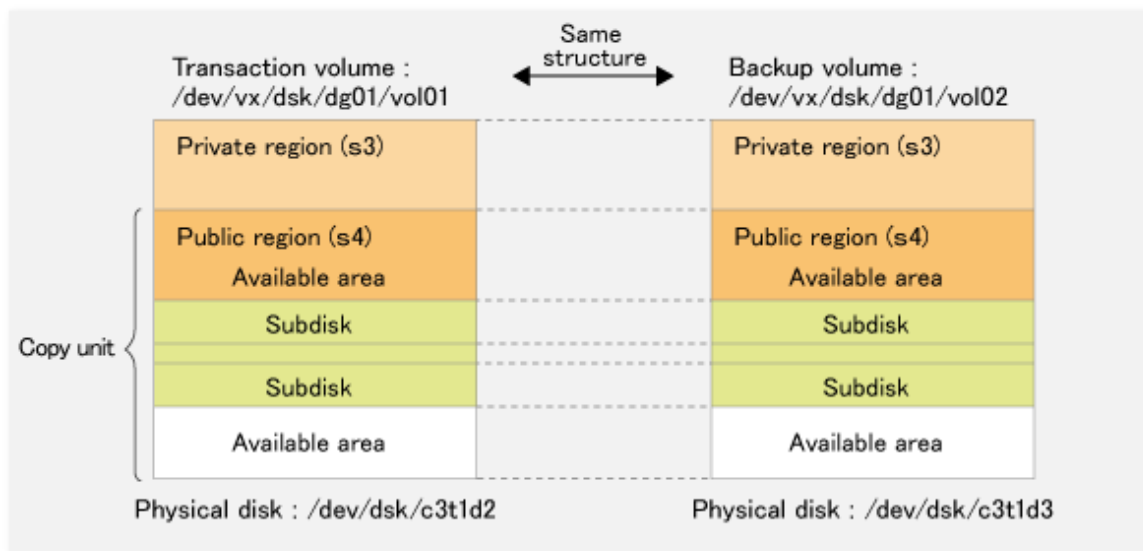



Figure 3.14 Example of Correspondence Between Transaction Volume and Backup Volume (2)



 **Note**

.....

If the backup volume is the same size as the public region, a backup volume of a VxVM volume can be made into a normal device (slice). However, it is generally not possible to mount the backup volume and have other applications use it. Similarly, a backup volume of a normal device (slice) can be made into a VxVM volume. However, it is generally not possible to mount the backup volume and have other applications use it.

.....


3.9.1.2 Preparations

3.9.1.2.1 Configuring Operation Type of Device

If the target VxVM logical volume satisfies the condition of logical volumes:VM disk = N:1 (N>1), use one of the logical volume names included with each VM disk to register a transaction volume or backup volume.

As an example, assume the transaction volume and backup volume in the following volume structure.

Figure 3.15 Volume Structure Example



- VM disk disk01 (transaction volume) has logical volumes named tranvol01, tranvol02, and tranvol03, and VM disk disk02 (backup volume) has logical volumes named backvol01, backvol02, and backvol03.
- VM disks disk01 and disk02 have the same slice size. The sub-disk structures of VM disks disk01 and disk02 are the same. (The sub-disk structure can be checked by executing "vxprint -AGTs".)

In the above example, use the logical volume names tranvol01 and backvol01 when configuring the device information.

```
# /opt/FJSVswsts/bin/swstdevinfofet -t /dev/vx/dsk/AGT01/tranvol01
swstdevinfofet completed
# /opt/FJSVswsts/bin/swstdevinfofet -b /dev/vx/dsk/AGT01/backvol01
swstdevinfofet completed
```

The logical volume specified in the device information configuration command is used to specify the VM disk (slice) for backup processing. This means that tranvol02 or tranvol03 instead of tranvol01 can be set in the above example (and this is also true for backvol01). However, if tranvol01 and backvol01 are already registered, the other volumes (tranvol02, tranvol03, backvol02, and backvol03) cannot be registered.

3.9.1.2.2 Customizing Pre-processing and Post-processing Scripts

If the target VxVM logical volume satisfies the condition where logical volumes:VM disk = N:1 (N>1) and unmount/mount processing of the transaction volume with a constructed file system is to be avoided for operational reasons, customize the backup pre-processing script (OpcBackup.pre) so that the transaction volume cannot be unmounted(*).

(*) In such cases, add processing before unmount processing of the backup pre-processing script (OpcBackup.pre) to clear the mount point of the VxVM logical volume, as shown below.

Customization Example

```
79: exit 0
80: fi
81: post_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".pre"
82: fstype_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".fstype"
83: bd_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".bd"
84:
    if [ "$device" = "/dev/vx/dsk/AGT01/tranvol01" ]
```



```

        then
            mount_point=""
        fi
85: if [ "$mount_point" != "" ]
86: then
87:
88: # When device cannot be unmounted
89: #
90: # if [ "$device" = "/dev/dsk/c*t*d*s*" ]

```

3.9.1.2.3 Preparing Device Map File

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

Example of Device Map File

```

# Transaction volume           Output destination backup volume
/dev/vx/dsk/AGT01/tranvol01   /dev/vx/dsk/AGT01/backvol01

```

For details on the device map file, refer to "[3.4.9 Preparing Device Map File](#)".

3.9.1.3 Operations

3.9.1.3.1 Executing Backup

If the target VxVM logical volume satisfies the condition where logical volumes:VM disk = N:1 (N>1), the transaction volume cannot be unmounted/mounted using the pre-processing or post-processing script.

For this type of structure, unmount the file system before executing backup, and mount the file system after backup.

If the target VxVM logical volume satisfies the condition where logical volumes:VM disk = 1:1, the transaction volume is unmounted and mounted in the same way as with normal slices, and the following procedure is not required.

Snapshot Backup Procedure

```

# umount /global/AGT01/tranvol01 <-- unmount tranvol01
# umount /global/AGT01/tranvol02 <-- unmount tranvol02
# umount /global/AGT01/tranvol03 <-- unmount tranvol03
# /opt/FJJSVswsts/bin/swstbackup /dev/vx/dsk/AGT01/tranvol01 -Xdevmap /acm/devmap.txt
/dev/vx/dsk/AGT01/tranvol01 swstbackup completed
# mount -F ufs /dev/vx/dsk/AGT01/tranvol01 /global/AGT01/tranvol01 <-- mount tranvol01
# mount -F ufs /dev/vx/dsk/AGT01/tranvol02 /global/AGT01/tranvol02 <-- mount tranvol02
# mount -F ufs /dev/vx/dsk/AGT01/tranvol03 /global/AGT01/tranvol03 <-- mount tranvol03

```

To prevent unmount/mount processing of the transaction volume for operational reasons, take action as described in "[3.9.1.2.2 Customizing Pre-processing and Post-processing Scripts](#)" and execute backup by using the procedure described below.

Snapshot Backup Procedure (While Preventing Unmount/Mount Processing of Transaction Volume)

```

# lockfs -w /global/AGT01/tranvol01 <-- lock file system tranvol01
# lockfs -w /global/AGT01/tranvol02 <-- lock file system tranvol02
# lockfs -w /global/AGT01/tranvol03 <-- lock file system tranvol03
# /opt/FJJSVswsts/bin/swstbackup /dev/vx/dsk/AGT01/tranvol01 -Xdevmap /acm/devmap.txt
/dev/vx/dsk/AGT01/tranvol01 swstbackup completed
# lockfs -u /global/AGT01/tranvol01 <-- unlock file system tranvol01
# lockfs -u /global/AGT01/tranvol02 <-- unlock file system tranvol02
# lockfs -u /global/AGT01/tranvol03 <-- unlock file system tranvol03
# fsck -n -F ufs /dev/vx/dsk/AGT01/backvol01 <-- check file system backvol01
# fsck -n -F ufs /dev/vx/dsk/AGT01/backvol02 <-- check file system backvol02
# fsck -n -F ufs /dev/vx/dsk/AGT01/backvol03 <-- check file system backvol03

```



Note

Even if the command "lockfs -w" is used to flush the file system cache and prevent further update, file contents are not guaranteed whenever the file is in Write Open (excluding the database space of Oracle used to support on-line backup).

If renewal of metadata is currently occurring, consistency of file system data itself cannot be guaranteed because "lockfs" has been specified.

To manage this problem, perform "fsck" after backup to check the consistency of the backup data.

Even if the following message is output in fsck processing, there is no problem with the consistency of the file system. Execute "fsck -y" to restore.

```
FILE SYSTEM STATE IN SUPERBLOCK IS WRONG; FIX?
```

3.9.1.3.2 Executing Restoration

If the target VxVM logical volume satisfies the condition where logical volumes:VM disk = N:1 (N>1), the transaction volume cannot be unmounted/mounted using the pre-processing or post-processing script.

For this type of structure, unmount the file system before executing restoration, and mount the file system after restoration.

If the target VxVM logical volume satisfies the condition where logical volumes:VM disk = 1:1, the transaction volume is unmounted and mounted in the same way as in operation with normal slices, and the following procedure is not required.

Restoration Procedure

```
# umount /global/AGT01/tranvol01 <-- umount tranvol01
# umount /global/AGT01/tranvol02 <-- umount tranvol02
# umount /global/AGT01/tranvol03 <-- umount tranvol03
# /opt/FJVSswsts/bin/swstrestore /dev/vx/dsk/AGT01/tranvol01
/dev/vx/dsk/AGT01/tranvol01 swstrestore completed
# mount -F ufs /dev/vx/dsk/AGT01/tranvol01 /global/AGT01/tranvol01 <-- mount tranvol01
# mount -F ufs /dev/vx/dsk/AGT01/tranvol01 /global/AGT01/tranvol02 <-- mount tranvol02
# mount -F ufs /dev/vx/dsk/AGT01/tranvol01 /global/AGT01/tranvol03 <-- mount tranvol03
```

3.9.1.4 Changing Configuration Information

To change the volume structure during the backup operation, follow the procedure given below.

1. Specify the -z option for the [swsthistdel](#) command, to delete all history information.
2. Execute the [swstbkpdel](#) command to delete relevant backup policies.
3. Delete the registration of the transaction volume by specifying the -o option in the [swstdevinfo](#) command. If the backup volume is a VxVM volume, also delete the registration of the backup volume by specifying the -o option in the [swstdevinfo](#) command.
4. Change the transaction volume. If the backup volume is a VxVM volume, also change the volume structure of the backup volume so that the backup volume and transaction volume can be in the same structure.
5. With Web Console, execute the discovering all devices and reflecting information (Discover/Reload).
6. Register the transaction volume and backup volume again by executing the [swstdevinfo](#) command.
7. Execute the [swstbkpolset](#) command to set a backup policy.
8. Resume the backup operation.

3.9.2 Backup Operation in Units of Physical Disks

When the VxVM volume configuration does not satisfy the conditions for backup operation using units in logical volumes, backup operation can be performed using units of physical disks.

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

Point

For the physical volume which should be synchronized, confirm it by using either Web Console or the `stgxfwcmdispdev` command.

Note

- It is only VM disk of the "nopriv" type that a physical slice becomes the unit of management.
- Before starting this operation, be sure to understand the basic operation of ordinary volumes.
- Snapshot backup is recommended for backing up in units of physical disks. In the case of synchronized backup, commands that access the destination volume, such as VxVM commands cannot be used during full copy or differential copy.
- For operation in a cluster configuration, the device name (`/dev/(r)dsk/c#t#d#`) of the physical disks that comprise the disk group must be the same at all servers that comprise the cluster, and the ETERNUS Disk storage system's disk indicated by the device name must be the same.

For operation in a SunCluster environment, the VxVM enclosure name must be the same at all servers that comprise the cluster, and the ETERNUS Disk storage system's disk indicated by the enclosure name must be the same.


3.9.2.1 Operational Configuration

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

Observe the following conditions when configuring the disk groups:

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

Figure 3.16 Transaction Volume and Backup Volume



3.9.2.2 Preparations

3.9.2.2.1 Confirming Disk Group Configuration Information File

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

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

3.9.2.2.2 Setting Transaction Volume and Backup Volume

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

Example:

```
# /opt/FJSVswsts/bin/swstdevinfoaset -t /dev/vx/dmp/c1t0d10s2
swstdevinfoaset completed
# /opt/FJSVswsts/bin/swstdevinfoaset -t /dev/vx/dmp/c1t0d11s2
swstdevinfoaset completed
# /opt/FJSVswsts/bin/swstdevinfoaset -b /dev/vx/dmp/c1t0d20s2
swstdevinfoaset completed
# /opt/FJSVswsts/bin/swstdevinfoaset -b /dev/vx/dmp/c1t0d21s2
swstdevinfoaset completed
#
```

3.9.2.2.3 Preparing Device Map File

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

Example of Device Map File

```
# Transaction volume          Backup volume
/dev/vx/dmp/clt0d10s2        /dev/vx/dmp/clt0d20s2
/dev/vx/dmp/clt0d11s2        /dev/vx/dmp/clt0d21s2
```

For details on the device map file, refer to "3.4.9 Preparing Device Map File".

3.9.2.3 Backup

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

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

Example of Snapshot Backup

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

# /opt/FJSVswsts/bin/swstbackup /dev/vx/dmp/clt0d10s2 -Xdevmap /acm/devmap.txt
/dev/vx/dmp/clt0d10s2 swstbackup completed
# /opt/FJSVswsts/bin/swstbackup /dev/vx/dmp/clt0d11s2 -Xdevmap /acm/devmap.txt
/dev/vx/dmp/clt0d11s2 swstbackup completed
#

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

Example of Synchronous Backup

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

# /opt/FJSVswsts/bin/swststartsync /dev/vx/dmp/clt0d10s2 -Xdevmap /acm/devmap.txt
/dev/vx/dmp/clt0d10s2 swststartsync completed
# /opt/FJSVswsts/bin/swststartsync /dev/vx/dmp/clt0d11s2 -Xdevmap /acm/devmap.txt
/dev/vx/dmp/clt0d11s2 swststartsync completed
#

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

# /opt/FJSVswsts/bin/swstbackup /dev/vx/dmp/clt0d10s2
/dev/vx/dmp/clt0d10s2 swstbackup completed
# /opt/FJSVswsts/bin/swstbackup /dev/vx/dmp/clt0d11s2
/dev/vx/dmp/clt0d11s2 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.

Pre-processing and Post-processing for Restoration

Table 3.8 In Case of Normal Operation (Non-clustered System)

Volume Type	Pre-processing	Post-processing
Transaction volume	<ol style="list-style-type: none"> Secure data integrity by stopping access to all logical volumes in the disk group. If file systems are included, unmount all file systems in the disk group. 	If file systems are included, mount the volumes that were unmounted during preprocessing.

Volume Type	Pre-processing	Post-processing
	3. Import the disk group, when the disk group is not imported.	
Backup volume	<ol style="list-style-type: none"> 1. Stop access to all logical volumes in the disk group. 2. If file systems are included, unmount all file systems in the disk group. 3. Deport the disk group. 4. Set the physical disks in the disk groups to offline. 	<ol style="list-style-type: none"> 1. Online the physical disks that were set to offline during preprocessing. 2. Reconfigure the disk group. 3. If file systems are included, remount the volumes that were unmounted during preprocessing.

Table 3.9 In Case of Operation in Clustered System

Volume type	Pre-processing	Post-processing
Transaction volume	<ol style="list-style-type: none"> 1. Secure data integrity by stopping access to all logical volumes in the disk group. 2. If file systems are included, unmount all file systems in the disk group. 3. Implement the following tasks depending on the disk group registration status in the clustered system disk group resource. <ul style="list-style-type: none"> - When the Disk Group is not Registered: If the disk group is not import, import the disk group. - When the Disk Group is Registered: If the disk group resource is not online, bring the disk group resource online. 	If file systems are included, mount the volumes that were unmounted during preprocessing.
Backup volume	<p>- When not using the VERITAS Cluster Server, or when using a version earlier than the VERITAS Cluster Server 6.0</p> <ol style="list-style-type: none"> 1. Stop access to all logical volumes in the disk group. 2. If file systems are included, unmount all file systems in the disk group. 3. Implement the following tasks depending on the disk group registration status in the clustered system disk group resource. <ul style="list-style-type: none"> - When the Disk Group is not Registered: Deport the disk group. - When the Disk Group is Registered: Offline the disk group resource. 4. Set the physical disks in the disk groups to offline. 	<p>- When not using the VERITAS Cluster Server, or when using a version earlier than the VERITAS Cluster Server 6.0</p> <ol style="list-style-type: none"> 1. Online the physical disks that were set to offline during preprocessing. 2. Reconfigure the disk group. 3. If file systems are included, remount the volumes that were unmounted during preprocessing.
	<p>- When using a version later than the VERITAS Cluster Server 6.0</p> <ol style="list-style-type: none"> 1. Stop access to all logical volumes in the disk group. 	<p>- When using a version later than the VERITAS Cluster Server 6.0</p> <ol style="list-style-type: none"> 1. Online the physical disks that were set to offline during preprocessing.

Volume type	Pre-processing	Post-processing
	<ol style="list-style-type: none"> 2. Implement the following tasks to all the file systems in the disk group depending on the registration status of the mounting point as a resource of the clustered system. <ul style="list-style-type: none"> - When the Mount Point is not Registered: Unmount the file system. - When the Mount Point is Registered: Instead of unmounting the file system, offline the mounting resource, and then configure the mounting resource Enabled attribute to off. 3. Implement the following tasks depending on the volume registration status in the clustered system volume resource. <ul style="list-style-type: none"> - When the Volume is not Registered: Suspend the volume. - When the Volume is Registered: Offline the volume, and then configure the volume resource Enabled attribute to off. 4. Implement the following tasks depending on the disk group registration status in the clustered system disk group resource. <ul style="list-style-type: none"> - When the Disk Group is not Registered: Deport the disk group. - When the Disk Group is Registered: Instead of deporting the disk group, offline the disk group resource, and then configure the mounting resource Enabled attribute to off. 5. Set the physical disks in the disk groups to offline. 	<ol style="list-style-type: none"> 2. Reconfigure the disk group. 3. Deport the disk group. 4. Implement the following tasks depending on the disk group registration status in the clustered system disk group resource. <ul style="list-style-type: none"> - When the Disk Group is not Registered: Import the disk group. - When the Disk Group is Registered: Instead of importing the disk group, after configuring the Enabled attribute of the disk group resource to on, and then online the disk group resource. 5. Implement the following tasks depending on the volume registration status in the clustered system volume resource. <ul style="list-style-type: none"> - When the Volume is not Registered: Start the volume. - When the Volume is Registered: Configure the volume resource Enabled attribute to on, and then online the volume resource. 6. Implement the following tasks depending on the registration status of the mounting point as a resource of the clustered system. <ul style="list-style-type: none"> - When the Mount Point is not Registered: Mount file systems that were unmounted during preprocessing. - When the Mount Point is Registered: Instead of mounting the file system, after configuring the mounting resource Enabled attribute to on, and then online the mounting resource.

Reconfiguring Disk Group

Reconfigure the disk group as follows:

1. Pre-commit analysis for restoration

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

Installing volume manager disk header for clt0d20s2 ...
Installing volume manager disk header for clt0d21s2 ...
-
dstdg's diskgroup configuration is restored (in precommit state).
Diskgroup can be accessed in read only and can be examined using
vxprint in this state.
```

Run:

```
vxconfigrestore -c dstdg ==> to commit the restoration.
```

```
vxconfigrestore -d dstdg ==> to abort the restoration.
#
```

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

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

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

Note

- When performing a system disk exchange, there are cases when conflicting backup configuration information may exist. In such cases, the disk group ID needs to be reset after executing the above command.
- After this operation, where the volumes within a disk group are required to be run in synchronous mode in background, synchronous processing may take some time depending on the volume configuration. It is also possible to use the volumes during this time.

3.9.2.4 Restoration

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

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

Example of Restoration

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

```
# /opt/FJSVswsts/bin/swstrestore /dev/vx/dmp/clt0d10s2
/dev/vx/dmp/clt0d10s2 swstrestore completed
# /opt/FJSVswsts/bin/swstrestore /dev/vx/dmp/clt0d11s2
/dev/vx/dmp/clt0d11s2 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.

Pre-processing and Post-processing for Restoration

Table 3.10 In Case of Normal Operation (Non-clustered System)

Volume Type	Pre-processing	Post-processing
Backup volume	<ol style="list-style-type: none"> 1. Secure data integrity by stopping access to all logical volumes in the disk group. 2. Import the disk group, when the disk group is not imported. 	Post-processing is not required.
Restoration destination volume	<ol style="list-style-type: none"> 1. Stop access to all logical volumes in the disk group. 2. If file systems are included, unmount all file systems in the disk group. 3. Deport the disk group. 4. Set the physical disks in the disk groups to offline. 	<ol style="list-style-type: none"> 1. Online the physical disks that were set to offline during preprocessing. 2. Reconfigure the disk group. 3. If file systems are included, remount the volumes that were unmounted by preprocessing.

Table 3.11 In Case of Operation in Clustered System

Volume Type	Pre-processing	Post-processing
Backup volume	<ol style="list-style-type: none"> 1. After suspending access to all logical volumes within the disk group, secure consistency of the data. 2. Implement the following tasks depending on the disk group registration status in the clustered system disk group resource. <ul style="list-style-type: none"> - When the Disk Group is not Registered: If the disk group is not import, import the disk group. - When the Disk Group is Registered: If the disk group resource is not online, bring the disk group resource online. 	<p>Post-processing is not required.</p>
Restoration destination volume	<ul style="list-style-type: none"> - When not using the VERITAS Cluster Server, or when using a version earlier than the VERITAS Cluster Server 6.0 <ol style="list-style-type: none"> 1. Stop access to all logical volumes in the disk group. 2. If file systems are included, unmount all file systems in the disk group. 3. Implement the following tasks depending on the disk group registration status in the clustered system disk group resource. <ul style="list-style-type: none"> - When the Disk Group is not Registered: Deport the disk group. - When the Disk Group is Registered: Offline the disk group resource 4. Set the physical disks in the disk groups to offline. 	<ul style="list-style-type: none"> - When not using the VERITAS Cluster Server, or when using a version earlier than the VERITAS Cluster Server 6.0 <ol style="list-style-type: none"> 1. Online the physical disks that were set to offline during preprocessing. 2. Reconfigure the disk group. 3. If file systems are included, remount the volumes that were unmounted by preprocessing.
	<ul style="list-style-type: none"> - When using a version later than the VERITAS Cluster Server 6.0 <ol style="list-style-type: none"> 1. Stop access to all logical volumes in the disk group. 2. Implement the following tasks to all the file systems in the disk group depending on the registration status of the mounting point as a resource of the clustered system. <ul style="list-style-type: none"> - When the Mount Point is not Registered: Unmount the file system. - When the Mount Point is Registered: Instead of unmounting the file system, offline the mounting resource, and then configure the mounting resource Enabled attribute to off. 3. Implement the following tasks depending on the volume registration status in the clustered system volume resource. 	<ul style="list-style-type: none"> - When using a version later than the VERITAS Cluster Server 6.0 <ol style="list-style-type: none"> 1. Online the physical disks that were set to offline during preprocessing. 2. Reconfigure the disk group. 3. Deport the disk group. 4. Implement the following tasks depending on the disk group registration status in the clustered system disk group resource. <ul style="list-style-type: none"> - When the Disk Group is not Registered: Import the disk group. - When the Disk Group is Registered: Instead of importing the disk group, after configuring the Enabled attribute of the disk group resource to on, and then online the disk group resource.

Volume Type	Pre-processing	Post-processing
	<ul style="list-style-type: none"> - When the Volume is not Registered: Stop the volume. - When the Volume is Registered: Offline the volume, and then configure the volume resource Enabled attribute to off. <p>4. Implement the following tasks depending on the disk group registration status in the clustered system disk group resource.</p> <ul style="list-style-type: none"> - When the Disk Group is not Registered: Deport the disk group. - When the Disk Group is Registered: Instead of deporting the disk group, offline the disk group resource, and then configure the mounting resource Enabled attribute to off. <p>5. Set the physical disks in the disk groups to offline.</p>	<p>5. Implement the following tasks depending on the volume registration status in the clustered system volume resource.</p> <ul style="list-style-type: none"> - When the Volume is not Registered: Start the volume. - When the Volume is Registered: Configure the volume resource Enabled attribute to on, and then online the volume resource. <p>6. Implement the following tasks depending on the registration status of the mount point as a resource of the clustered system.</p> <ul style="list-style-type: none"> - When the Mount Point is not Registered: Mount file systems that were unmounted during preprocessing. - When the Mount Point is Registered: Instead of mounting the file system, after configure the mounting resource Enabled attribute to on, and then online the mounting resource.

Reconfiguring Disk Group

Reconfigure the disk group as follows:

1. Restoration pre-commit analysis

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

Installing volume manager disk header for clt0d10s2 ...
Installing volume manager disk header for clt0d11s2 ...
-
srcdg's diskgroup configuration is restored (in precommit state).
Diskgroup can be accessed in read only and can be examined using
vxprint in this state.

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

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

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

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



Note

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

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

Chapter 4 Backup Operation Using Suspend/Resume Function

This chapter provides details of the backup operation using the Suspend/Resume function.



For backup operations with the Suspend/Resume function, the Advanced Copy function enabled for TFOVs is not available.

4.1 Overview

The Suspend/Resume function sets Suspend/Resume for the equivalency maintenance status using EC. Using this function for differential copying from the Suspend state enables higher-speed synchronous backup.


AdvancedCopy Manager provides two backup functions: snapshot high-speed backup and synchronous high-speed backup. In snapshot high-speed backup, OPC is used to create a backup. In synchronous high-speed backup, EC is used to create a backup. The Suspend/Resume function can be thought of as a kind of synchronous high-speed backup that uses EC.

With ordinary synchronous high-speed backup, the synchronous processing start command starts copying using the EC function, then keeps the equivalency after the copying is completed. The backup execution start command is then executed. This cancels EC and creates a backup.

In synchronous high-speed backup using the Suspend/Resume function, when the backup execution start command is executed (with -suspend specified), EC is suspended, the backup process is completed, and the EC session status is saved. The next time that the synchronous processing start command is executed, EC resumes and only data that has been updated since the suspend state was last set is copied. As a result of this, the preparation time of the backup is shortened.

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

Figure 4.1 Comparison Between Ordinary Synchronous High-speed Backup and Synchronous High-speed Backup Using Suspend/Resume Function



4.2 Backup Operation Design

Notes about synchronous high-speed backup using the Suspend/Resume function and backup operation design are provided below, as well as information related to backup management operation design (refer to "3.2 Operation Design").

Transaction Volume

Register the backup source device, used for storing transaction data, as the transaction volume.

Backup Volume

Register the backup destination device in which transaction volume data is saved as the backup volume.

When the backup command with `-suspend` option specified is executed, the EC session between the transaction volume and backup volume is maintained in a Suspended state. In this Suspend state, the backup volume cannot be used as the backup destination of another transaction volume.


Point

In the backup process that uses the Suspend/Resume function, each transaction volume requires as many backup volumes as the saved generation number + 1.

4.3 Flow of Operations

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

Figure 4.2 Flow of Backup Operations Using Suspend/Resume Function



4.4 Preparations

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

4.4.1 Activating Daemons

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

For information on starting a daemon, refer to "[Chapter 2 Starting and Stopping Daemons](#)".

4.4.2 Starting Web Console

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

4.4.3 Registering Managed Server

With Web Console, register Managed Servers that are to be managed.



- Change to the value of 800 characters or less before registering the Managed Server when the value of the environment variable PATH of the Management Server is more than 800 characters.
 - Perform this task after confirming that it is possible to communicate from the Managed Server to the IP address of the Management Server. Check the Management Server IP address with the [stgxfwcmdisprv](#) command.
-

4.4.4 Fetching Device Information on Managed Server

This process can be carried out using either Web Console or the [stgxfwcmsetdev](#) command.

Refer to "Add/Reload Device" in the *ETERNUS SF Web Console Guide* for information on the operating procedures of Web Console.



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

4.4.5 Configuring Environment for Backup Operation Server

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

This process can be carried out using either Web Console or the [swstsvrset](#) command.

In addition, once it is set, you can display the Managed Server configuration information by executing the [swstsvrdisp](#) command.

4.4.6 Setting Operation Type for Device

Register a transaction volume as the backup source and a backup volume to be used as the backup destination.

This process can be carried out using either Web Console or the [swstdevinfo](#) command.

In addition, to display the registered volume information, execute the [swstdevdisp](#) command.



Note

- A suspended transaction volume cannot be changed to be used for a different purpose.
- A suspended backup volume cannot be changed to be used for a different purpose.



Point

When the device information of a suspended backup volume is displayed, "Backup (used)" is displayed in the **Device-Mode** column.

4.4.7 Setting Backup Policy

Before a backup policy is set, you need to register as many usable backup volumes as are required to perform the operation in accordance with the specified backup policy.

This process can be carried out using either Web Console or the following commands.

- Backup Policy Setup: [swstbkpolset](#)
- Check Backup Policy that has been set up: [swstbkpoldisp](#)
- Backup Policy Deletion: [swstbkpoldel](#)

4.4.8 Customizing Pre-processing and Post-processing

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

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

- For operational reasons, to avoid unmount and mount processing of transaction volumes that comprise a file system
- To add special pre-processing and post-processing to the corresponding scripts.

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



Point

When AdvancedCopy Manager Is Upgraded

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

4.4.9 Preparing Device Map File

In its backup operations, AdvancedCopy Manager automatically selects, from a group of devices registered as backup volumes, a device that has the same capacity as the transaction volume. It then uses the selected device as the backup destination.

If there is a suspended backup volume at the time of backup destination selection and the start of the synchronous processing, its corresponding device is selected first, and is used as the backup destination.

To conveniently monitor backup destination devices, create a [device map file](#) that maps transaction volumes to backup volumes in advance. Specifying this file at the start of the synchronous processing or during the backup process enables operation with the backup destination specified.

P Point

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

4.5 Backup Operations

For synchronous high-speed backup using the Suspend/Resume function, execute the `swstbackup` command after the synchronous processing has started and the equivalency maintenance status has been set.

Then suspend the synchronous processing.


When history information is deleted because the number of generations has been exceeded or when the `swsthistdel` command is executed, the backup volume becomes the Suspend state without any history information.

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


Note that synchronous high-speed backup using the Suspend/Resume function can also be operated from Web Console.

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


Figure 4.3 Synchronous High-speed Backup Operations Using Suspend/Resume Function




4) Execution of the backup command with -suspend specified




5) Synchronous processing started



6) Equivalency maintain status




7) Backing up with the suspend specification




History information deleted because number of generations exceeded (*1) at backup command execution

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

8) Transaction volume data updated





4.5.1 Starting Synchronous Processing

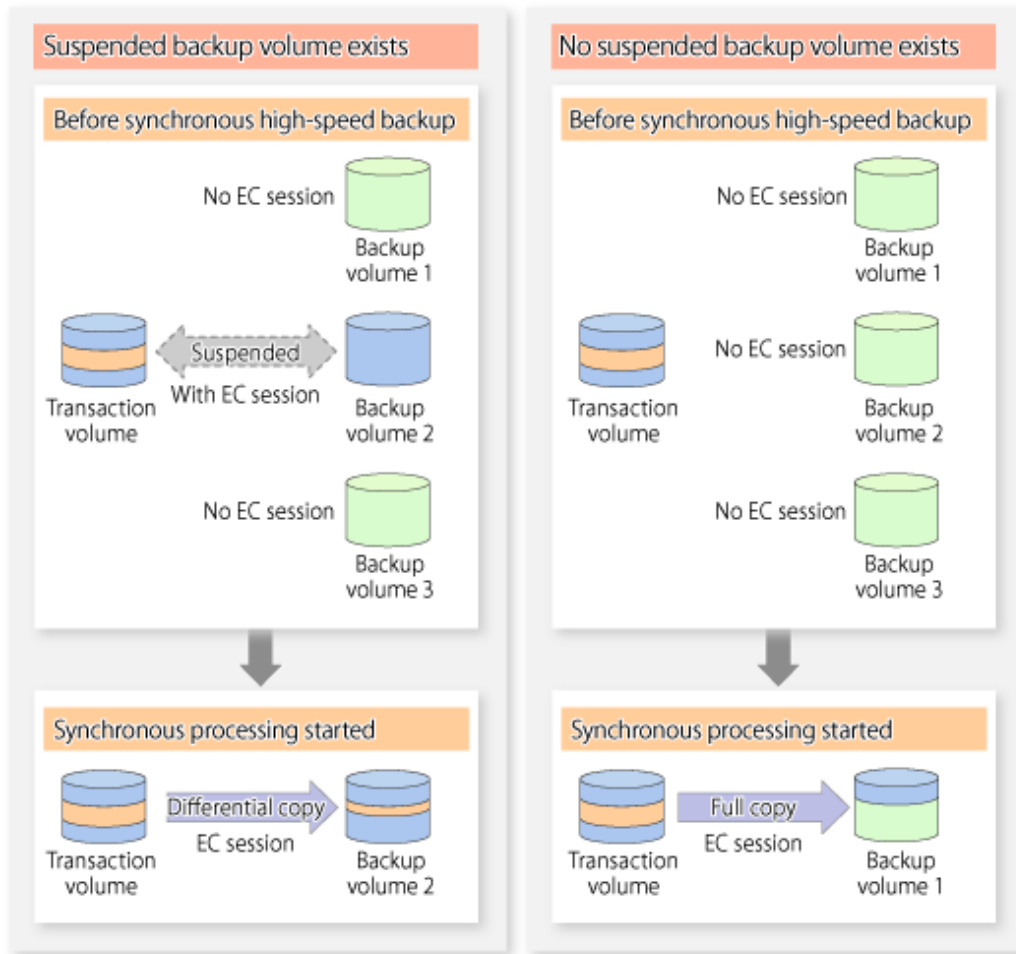
For synchronous high-speed backup, execute the `swststartsync` command to start the synchronous processing. If there is a suspended backup volume at the start of the synchronous processing, then the following devices are selected first, and EC resumes (i.e., the Resume function causes differential copying to start).

- The backup volume from which history information is deleted because the number of generations has been exceeded; or
- The suspended backup volume without history information.

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

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

Figure 4.4 Backup Volume Selection at 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 `swstsyncstat` command.


4.5.2 Backing Up

For synchronous high-speed backup using the Suspend/Resume function, start synchronous processing. Then, when the copy status of both the transaction volume and backup volume attain equivalency maintenance status, execute the `swstbackup` command with the suspend specification (using `-suspend` option).

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

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

Figure 4.5 Ordinary Backup Operations and Backup Operations with Suspend Specification




4.5.3 Displaying Synchronous Processing Execution Status

To check the status of synchronous high-speed backup that uses the Suspend/Resume function, execute the `swstsyncstat` command. If one transaction volume has multiple EC sessions, then this command displays the status of all EC sessions. The statuses displayed for synchronous processing are explained below.

- executing : Synchronous processing (i.e., copying) is in progress
- equivalent : Equivalency maintain status
- suspend : Suspended processing

The command displays the synchronous processing execution status as shown below:

Figure 4.6 Ordinary Backup Operations and Backup Operations with Suspend Specification



4.5.4 Cancelling Synchronous Processing

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

- executing : Synchronous processing (i.e., copying) is in progress
- equivalent : Equivalency maintain status
- suspend : Suspended processing

After the backup command with `-suspend` option specified is executed, the Suspend state is maintained even if history information is deleted. A suspended EC session cannot be released unless the synchronous processing is cancelled.

If the synchronous processing is cancelled with no options specified, the EC session in the synchronous processing status or in the equivalency maintenance status is cancelled.

If one transaction volume has multiple synchronous processing in progress (i.e., running multiple EC sessions), then only the EC sessions mapped to specific backup volumes can be cancelled. All of the synchronous processing of the transaction volume can be cancelled by specifying the synchronous processing.

Note

The backup volume and all synchronous processing cannot both be specified simultaneously.

Point

- Cancelling the suspended synchronous processing that has history information does not delete the history information.
- Cancelling synchronous processing (during copying, in the equivalency maintenance status, or in the Suspend state) that does not have any history information sets the backup volume to the unused state.
- Cancelling synchronous processing (in the Suspend state) that has history information leaves the backup volume in the used state.

A range of operational states when synchronous processing is cancelled are shown below:

Figure 4.7 Operational States in Synchronous Processing Caused by Cancellation



4.5.5 History Management


Backup history information can be referenced using the `swsthistdisp` command.
Backup history information can be deleted using the `swsthistdel` command.

Point

- Even when the history information created during backup operations with the suspend specification is deleted, an associated EC session is not released.
- Deleting history information during ordinary synchronous high-speed backup releases the backup volume, and "Backup(free)" is displayed in the **Device-Mode** column by the `swstdevdisp` command.
- Deleting history information during backup operations with the suspend specification does NOT release the backup volume after the information is deleted, and "Backup(used)" is displayed in the **Device-Mode** column by the `swstdevdisp` command.

The process of deleting ordinary backup history information and backup history information with the suspend specification is shown below.

Figure 4.8 Deleting Ordinary Backup History Information and Backup History Information with Suspend Specification



4.5.6 Restoring

Before restoration, the `swstcancelsync` command must be executed to release all EC sessions that are set on the transaction volume. To check EC sessions set on the transaction volume, execute the `swstsyncstat` command.

To restore data from a backup volume, execute the `swstrestore` command.

To cancel restoration or to delete management information associated with restoration that ended abnormally, execute the `swstcancelrest` command.

Note

- Restoration is not possible if the transaction volume has a backup volume being copied by EC, a backup volume under the equivalency maintenance status, or a suspended backup volume.
- A suspended backup volume cannot be used as the restore destination volume.

Chapter 5 Backup Operation Using QuickOPC Function

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



For backup operations with the QuickOPC function, the Advanced Copy function enabled for TFOVs is not available.

5.1 Overview

For a snapshot high-speed backup using OPC, the entire transaction volume must be copied to the backup volume every time the backup is started. However, as long as the same transaction volume and backup volume are being used, the entire transaction volume need not be physically copied every time OPC is run. Instead, physically copying only the data that has been modified since the start of the previous backup can create a complete snapshot image. This can greatly reduce the time taken for physical copying.

The QuickOPC function achieves this by copying only the data that has been changed from the point of the previous physical copy.

The snapshot high-speed backup using the QuickOPC function is referred to as the differential snapshot high-speed backup.

The conventional snapshot high-speed backup and the differential snapshot high-speed backup are compared below:




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



When the QuickOPC function is used, the hardware records the changes that are made to the copy source and destination after completing the OPC logical copy. The state in which the hardware keeps a record of changes is called "tracking state".

Figure 5.1 Comparison Between Conventional Snapshot High-speed Backup and Differential Snapshot High-speed Backup



5.2 Configuring Backup Operation

Configure the differential snapshot high-speed backup using the QuickOPC function using information provided in the section relating to "3.2 Operation Design" for backup management with consideration to the following notes:

Note

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


- In a backup operation using the QuickOPC function, an OPC session remains active even after the physical copy has been completed. Operators should specify the number of preservation generations so that the maximum number of OPC sessions that can be defined for the same logical unit cannot be exceeded.

- Operators must prepare a backup volume in the same box as the transaction volume to be backed up.

5.3 Operation Flow

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

Figure 5.2 Flow of Backup Operations Using QuickOPC Function



5.4 Preparation

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

5.4.1 Activating Daemons

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

For information on starting a daemon, refer to "[Chapter 2 Starting and Stopping Daemons](#)".

5.4.2 Starting Web Console

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

5.4.3 Registering Managed Server

With Web Console, register Managed Servers that are to be managed.



- Change to the value of 800 characters or less before registering the Managed Server when the value of the environment variable PATH of the Management Server is more than 800 characters.
 - Perform this task after confirming that it is possible to communicate from the Managed Server to the IP address of the Management Server. Check the Management Server IP address with the [stgxfwcmdisprv](#) command.
-

5.4.4 Fetching Device Information on Managed Server

This process can be carried out using either Web Console or the [stgxfwcmsetdev](#) command. Refer to "Add/Reload Device" in the *ETERNUS SF Web Console Guide* for information on the operating procedures of Web Console.



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

5.4.5 Configuring Environment for Backup Operation Server

Set the configuration information of the Managed Server processing the backup operations.

This process can be carried out using either Web Console or the [swstsvrset](#) command.

In addition, once it is set, you can display the Managed Server configuration information by executing the [swstsvrdisp](#) command.

5.4.6 Setting Operation Type for Device

Register a transaction volume as the backup source and a backup volume to be used as the backup destination.

This process can be carried out using either Web Console or the [swstdevinfo](#) command.

In addition, to display the registered volume information, execute the [swstdevdisp](#) command. The transaction volume in the tracking state can be checked with the [swstrkstat](#) command.



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

Point

The both `-b` and `-u` options can be specified in the `swstdevdisp` command to check for backup volumes, where relevant histories have been deleted, in the suspend or tracking state.

5.4.7 Setting Backup Policy

Before a backup policy is set, you need to register as many usable backup volumes as are required to perform the operation in accordance with the specified backup policy.

This process can be carried out using either Web Console or the following commands.

- Backup Policy Setup: `swstbkpolset`
- Check Backup Policy that has been set up: `swstbkpoldisp`
- Backup Policy Deletion: `swstbkpoldel`

5.4.8 Customizing Pre-processing and Post-processing

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

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

- For operational reasons, to avoid unmount and mount processing of transaction volumes that comprise a file system
- To add special pre-processing and post-processing to the corresponding scripts.

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

Point

When AdvancedCopy Manager Is Upgraded

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

5.4.9 Preparing Device Map File

In the backup, AdvancedCopy Manager automatically selects a device which has the same capacity as the target transaction volume as a backup volume from a group of registered devices.

If a backup volume in the tracking state exists when AdvancedCopy Manager looks for a backup volume for the snapshot high-speed backup, it selects the backup volume in the tracking state in preference to any other and uses it as the backup destination.

When it is necessary to recognize a specific backup destination device for any reason a "device map file" that defines the relationship between a transaction volume and backup volume must be created in advance. By specifying the "device map file" when starting the backup, the backup can always be performed because the backup destination has already been specified.

Point

If multiple generations are to be managed, multiple device map files must be prepared.

5.5 Operation

5.5.1 Invoking Backup


Invoke the differential snapshot high-speed backup by specifying the -T option in the `swstbackup` command.

When the QuickOPC starts, the physical copy and tracking processing begins. When the physical copy is complete, only tracking processing is active. When this backup command is executed by specifying the -T option, only new data or data which has changed from the point in time that OPC was started previously is actually copied.

Note that differential snapshot high-speed backup can also be operated from Web Console.

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


Figure 5.3 Conventional Snapshot High-speed Backup and Differential Snapshot High-speed Backup



After history information is deleted in a differential snapshot high-speed backup, the backup volume is put into the tracking state without history information. If a new differential snapshot high-speed backup is started under that condition, the backup volume in the tracking state without history information is selected to start the differential data copy. The physical copy can thus be completed in a short period because only the difference in data from where OPC was started previously is copied.

The figure below shows a flow of a differential snapshot high-speed backup in which only one generation is preserved.

Figure 5.4 Flow of Differential Snapshot High-speed Backup Operation



5.5.2 History Management

Backup history information can be referenced using the [swsthistdisp](#) command.

Backup history information can be deleted using the [swsthistdel](#) command.

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


In the conventional snapshot high-speed backup, after the history information is deleted, the backup volume is released and "Backup (free)" is displayed in the **Device-Mode** column by the [swstdevdisp](#) command. In the differential snapshot high-speed backup,

however, the backup volume is not released even after the history information is deleted and "Backup (used)" is displayed in the **Device-Mode** column by the `swstdevdisp` 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 5.5 Comparison of History Deletion Between Snapshot High-speed Backup and Differential Snapshot High-speed Backup




5.5.3 Restoration

The section explains how to use the `swstore` command for restoration.

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


Figure 5.6 Tracking of Differential Snapshot High-speed Backup Operations



When restoration is executed, normal OPC from the backup volume to the transaction volume is started while the tracking state from the transaction volume to the backup volume is maintained.

Because OPC for restoration physically copies only the data that has been updated after the backup, the time taken for the physical copy can be reduced for the restoration, as well as for the backup.


Figure 5.7 Restoration from Backup Volume to Transaction Volume



While tracking processing is performed between the transaction volume and backup volume, restoration to other volumes cannot be performed. If restoration to other volumes is required, the following operational steps must be performed in advance:

1. If the OPC physical copy from the transaction volume to the backup volume is in progress, wait for the completion of the copy.
2. Stop tracking processing using the `swstcanceltrk` command.

Figure 5.8 Restoration to Another Volume



 Note

If tracking processing is terminated, all data in the transaction volume is physically copied by the next differential snapshot high-speed backup.

5.5.4 Cancelling Restoration

Execute the `swstcancelrest` command to cancel the restoration.

 **Point**

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

5.5.5 Displaying Execution Status of Tracking Processing

The execution status of tracking processing by the QuickOPC function can be checked with the [swstrkstat](#) command.

Information on an OPC session that is not in a tracking state is excluded from the display. For an OPC session in tracking, the amount of data that has been updated after execution of the QuickOPC (ie, the logical copy) is displayed as a percentage in the **Update** column.

5.5.6 Stopping Tracking Processing

Execute the [swstcanceltrk](#) command to stop tracking processing.

 **Point**

.....
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 6 Replication

This chapter explains AdvancedCopy Manager replication in Solaris.

A slice or Logical Unit (disk) is used as the replication unit.

In this manual, the term "volume" is used instead of "slice" or "Logical Unit (disk)".

6.1 Overview

Using the Advanced Copy function of ETERNUS Disk storage system, the replication function performs high-speed replication between volumes, regardless of the volume capacities.

Snapshot replication is the term used for the replication function that uses the OPC function of ETERNUS Disk storage system.

Synchronized Replication is where either the EC or REC functions are used.

The replication function creates copies so that the information at a certain point in time can be used for different purposes. It can be used either independently of or in combination with the backup function.

Unlike in the case of backup management, Symfoware database is not supported by replication management.


6.1.1 Snapshot Replication Processing

Snapshot replication copies from a source volume to a destination volume using the function of ETERNUS Disk storage system.

The ordinary snapshot replication using OPC performs the following two types of processing:

1. When a copy is created, the snapshot processing (OPC) from the source volume to the replica volume is executed with the [swsrpmake](#) command.
The [swsrpmake](#) command dismounts or mounts the original volume.
This process determines the scope of original data (for details, refer to "[Appendix C Pre-processing and Post-processing of Replication](#)"). Refer to steps (1). and (2). in the figure below.
2. If the user wants to recreate a replica, only the [swsrpmake](#) command is required. If the previous snapshot processing is in progress, it is stopped and a new snapshot process is started.

Figure 6.1 Ordinary Snapshot Replication Processing



Snapshot replication is completed when the replication creation command is issued. During snapshot processing performed internally by the ETERNUS Disk storage system, if there is an attempt to access an area where copying is not yet complete at the replication destination, priority is given to copying the accessed area and, after it is copied, access is permitted ((a) in the figure). If there is an attempt to update the replication source data, the pre-update data at the replication source is first copied to the replication destination, and then the replication source data is updated ((b) in the figure). These operations make it seem that creation of the replica is completed almost instantly.

6.1.1.1 Types of Snapshot Replication


The following types of snapshot replication are available, depending on the used Advanced Copy functions:

- [Ordinary Snapshot Replication](#)
- [QuickOPC Type Replication](#)
- [SnapOPC Type Replication](#)
- [SnapOPC+ Type Replication](#)

Ordinary Snapshot Replication

Ordinary snapshot replication copies all the data at a certain point in time (logical copy) to the copy destination disk area.

Figure 6.2 Ordinary Snapshot Replication




P Point

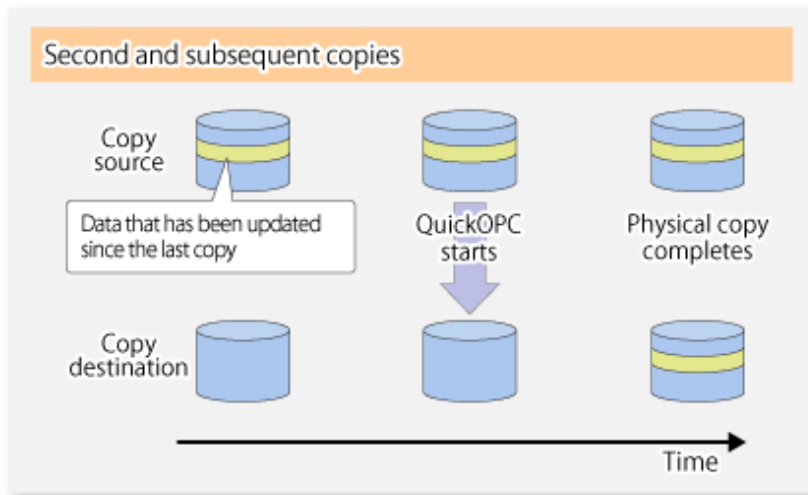
The copy destination disk must be the same size or larger than the copy source disk.
The copying time is the time taken to copy all the data.

QuickOPC Type Replication

The initial replication copies all the data at a certain point in time (logical copy) to the copy destination disk area. Subsequent replications copy only the data that has been updated since the previous replication. This enables large reductions in the time taken for physical copying.

Figure 6.3 QuickOPC Replication





P Point

- The copy destination disk must be the same size or larger than the copy source disk. The copying time for subsequent replications is the time taken to copy the differential data.
- 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 "tracking status".


Note

The QuickOPC function cannot be used for replicating SDX objects in the units of logical volumes.

SnapOPC Type Replication

SnapOPC is a function that enables only data that has been updated in the copy source disk after a certain point in time (a logical copy) to be copied to the destination disk.


Figure 6.4 SnapOPC Replication



P Point

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

Figure 6.5 SnapOPC Mechanism



In some ways, SnapOPC is superior to conventional OPC, since copy times are shorter and copy destination volumes can have a smaller capacity. However, there are problems in terms of access performance and the reliability of copy data, as explained in the following:

Access Performance

Since the data copy from the source to the destination occurs inside SnapOPC, the access performance of the copy source is sometimes reduced, as well as the access performance of the copy destination.

Reliability of Copy Data

Copy data is made up of two parts: the "copy source" and "updated part of the copy source". Therefore, a hardware fault with the copy source can cause copy data to be lost.

Additionally, performing a restoration is equivalent to losing the "copy source" data, and so even if there are multiple "updated parts of the copy source", these become invalid from the point when the restoration is performed.


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

SnapOPC+ Type Replication

When data is changed, the SnapOPC+ function copies only the original data of the changed data set, and saves it in snap generation units.

The following figure shows the processing flow of the SnapOPC+ function.

Figure 6.6 SnapOPC+ Function Flow

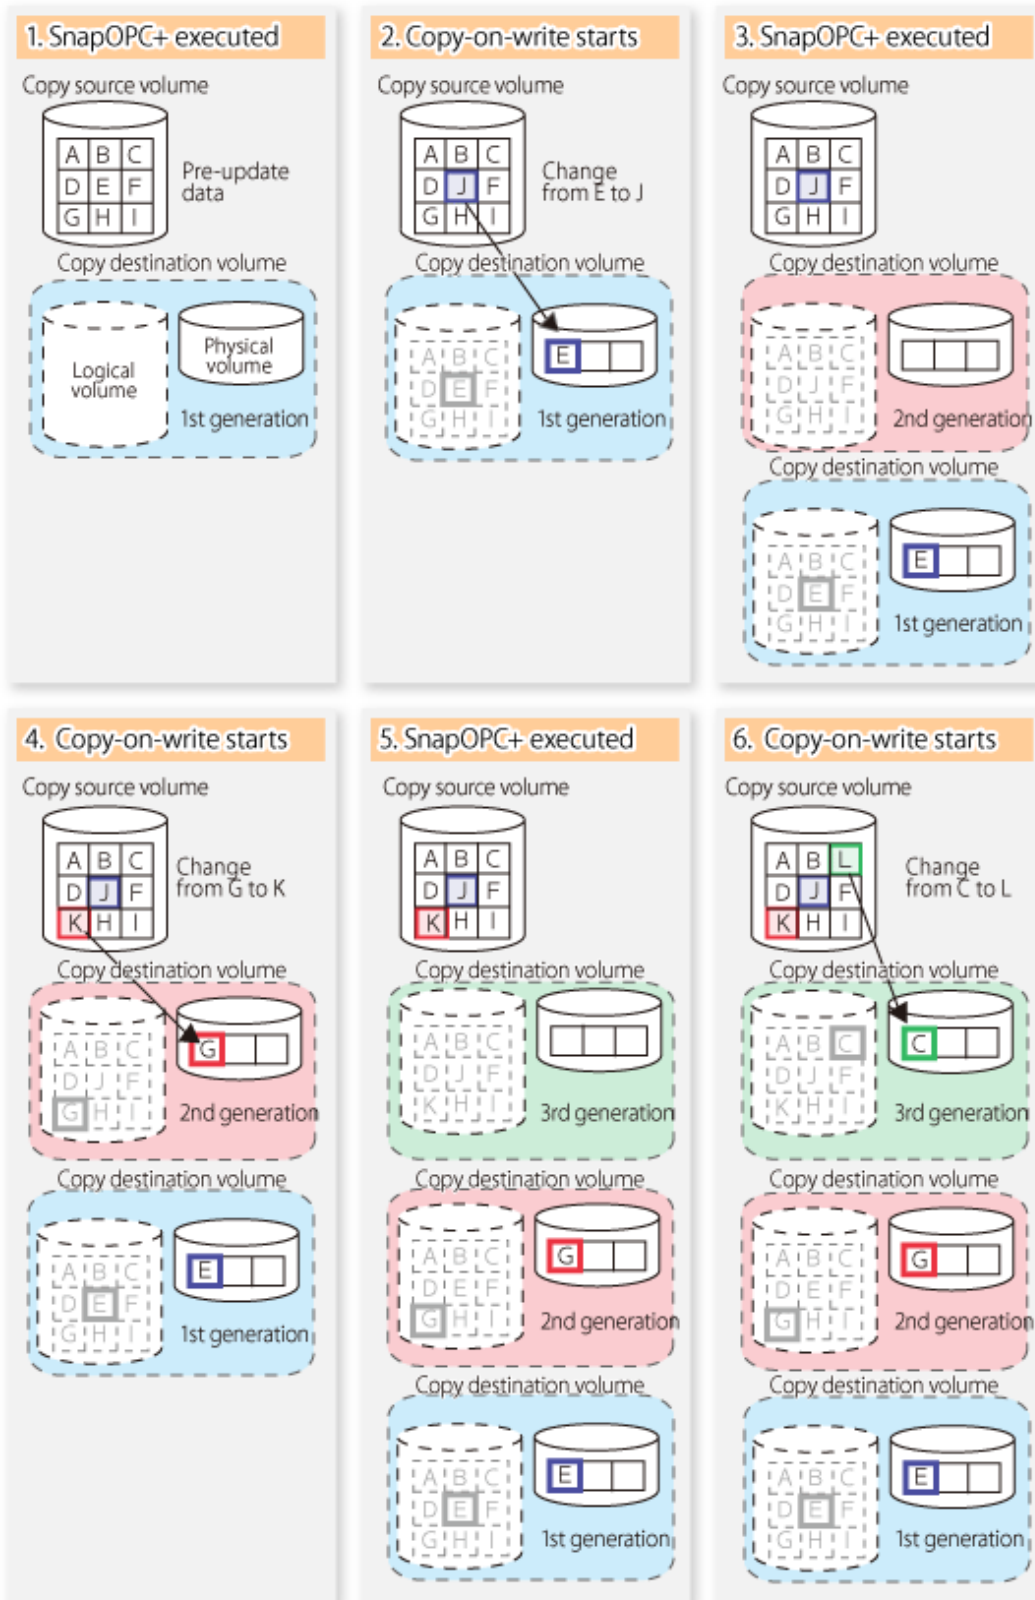


1. Logical copy is performed from the copy source to the copy destination.
2. After the logical copy, a SnapOPC+ session is set between the copy source and the copy destination.
3. After the SnapOPC+ session is set, only the original data of the changed data at the copy source is copied to the copy destination (copy-on-write processing).
4. When the next copy destination volume is created, the copy-on-write processing is stopped.
5. Then, a SnapOPC+ session is set between the copy source volume and the new copy destination volume, and copy-on-write processing is performed.
6. Each time a copy destination volume is created, a snap generation number is assigned to that volume. (Snap generation numbers are assigned in sequence from 1, starting with the oldest.)

Since SnapOPC+ copy-on-write processing is performed only between the copy source volume and the latest copy destination volume, there is no effect on access performance to the copy source volume even when the number of copy destination volumes increases. In addition, space at the copy destination volume is required only for the amount of original data (for data updated immediately prior) and for hardware management. Thus, use of SnapOPC+ enables the saved data area to be kept to the minimum area required. Snapshot replication that uses SnapOPC+ is called "SnapOPC+ replication".

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

Figure 6.7 1:3 Replication Operations Using SnapOPC+



Due to the mechanism and features of SnapOPC+, it is assumed that it is 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, it is recommended 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 6.1 List of Function Differences Between SnapOPC and SnapOPC+

Function	SnapOPC	SnapOPC+
Copy method	Copy-on-write method	Copy-on-write method
Copy time (logical copy duration)	Instant	Instant
Copy destination volume type	Snap Data Volume (SDV), Thin Provisioning Volume (TPV), Flexible Tier Volume (FTV)	Snap Data Volume (SDV), Thin Provisioning Volume (TPV), Flexible Tier Volume (FTV)
1:N operation	Available	Available (recommended): Access performance to the copy source volume is better than with SnapOPC. Also, compared with SnapOPC, less physical space is required at the copy destination volume.
Snap generation numbers	No	Yes
Restoration using the replication function	Yes	Yes (*1)
Usefulness	<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

*1: Only available when using ETERNUS Disk storage systems that support the restoration using replication.

Information

However, for the following ETERNUS Disk storage systems, SDV is the only copy destination volume type that can be used:

- ETERNUS DX60
- ETERNUS DX400 series
- ETERNUS DX8000 series
- ETERNUS DX S2 series
- ETERNUS DX S3 series (for earlier than V10L60 firmware)

6.1.1.2 SnapOPC/SnapOPC+ Copy Destination Volume

Regarding the SnapOPC/SnapOPC+ copy destination volume, the capacity that can be seen in the host is called "logical capacity" and the capacity that is physically configured is called "physical capacity".

Configure the logical capacity so that it is greater than or equal to the copy source capacity.

Prepare the total physical capacity of the backup data stored capacity and the control information (conversion table) stored capacity for managing the backup data. Estimate the capacity required to store the backup data from the update amount in the copy source volume using the `swstestupdate` command. The size of the control information is 0.1 % of the logical capacity.

If the physical capacity of the copy destination volume is insufficient, you cannot access the copy destination volume. For this reason, you must monitor the use of the copy destination volume so that a capacity shortage of the physical space in the copy destination volume does not occur during operation. If SDV is used as the SnapOPC+ copy destination volume, it is recommended that the Snap Data Pool (SDP) is set.



See

- Refer to "[6.2.3.3 Monitoring Usage of Copy Destination Volume](#)" for information on how to monitor the usage of the copy destination volume.
- If the physical capacity of the copy destination volume is insufficient, refer to "[8.4.2.3 Troubleshooting When Lack of Free Physical Space Has Occurred in Copy Destination Volume](#)" and take appropriate action.
- If SDV is used as the copy destination volume and SDP is set, the capacity is automatically added from SDP when the physical capacity of SDV is insufficient. In addition, if the SDP capacity is insufficient, you cannot access the copy destination volume contrary when the physical capacity of SDV is insufficient.
If the SDP capacity is insufficient, refer to "[8.4.2.3 Troubleshooting When Lack of Free Physical Space Has Occurred in Copy Destination Volume](#)" and take appropriate action. Refer to the storage device manuals for details of SDV and SDP.

6.1.2 Synchronized Replication Processing

Replicas are created by copying from the source volume to the replica volume using the EC or REC function of ETERNUS Disk storage system.


Synchronized replication uses two methods to create a copy: full copying and incremental copying. In the case of full copying, all of the source volume is completely copied. In the case of incremental copying, only the data that has been updated since the last replication is copied.

- Creating replicas with full copying
Use this method to create the first copy in a replication.
- Creating copies with incremental copying
Use this method to create copies in the future once a previous replica exists.
Copies are created by reflecting the data that has been updated after an initial copy has been made (full copying and incremental copying).

The procedure for synchronized replication incorporates the following steps:

1. Full copying starts when synchronized processing that has been invoked using the [swsrpstartsync](#) command starts copying from the source volume to the replica volume. Refer to step (1) in the figure below.
2. The overall copy by synchronous processing is completed, updates in the source volume are subsequently reflected in the replica volume, so that data equivalence is always maintained (equivalency maintenance state). Refer to the states in steps (2) and (3) in the figure below.
Note that, this state is not a state that guarantees the equivalence of the replication source volume and the replication destination volume.
3. Synchronized processing is temporarily stopped with the [swsrpmake](#) command to create a copy of the source volume.
The [swsrpmake](#) command dismounts or mounts the original volume.
This processing determines the scope of original data (for details of this, refer to "[Appendix C Pre-processing and Post-processing of Replication](#)"). After execution of the replication creation command, the replica volume can be accessed (copy established status). Refer to the states in steps (3) and (4) in the figure below.
4. To create another copy, the synchronized processing that has been temporarily stopped is restarted with the [swsrpstartsync](#) command. In the case of this copy, the only data copied to the replica volume is the data on the source volume that has been updated since the previous copy was made. Refer to step (4) in the figure below.
When the replica volume has been updated after the previous copy is made, the contents of the update of the replica volume are cleared.
5. When the incremental data has been copied, the status returns to the equivalency maintenance state. Refer to the states in steps (5) and (6) in the figure below.
6. When this status has been set, the copy is re-created with the [swsrpmake](#) command. Refer to step (6) in the figure below.

Figure 6.8 Synchronized Replication Processing



Note

- When using the EC or REC function in synchronized replication, ETERNUS Disk storage system must support the EC or REC function.
- The `srsrpmake` command cannot be executed unless both the source volume and replica volume are in the equivalency maintenance state.
- If the replication destination volume is being synchronized or in the equivalency maintenance state, it is not accessible.

6.2 Configuration

This section explains the operation of replication.

Configure replication using the following procedure:

1. [6.2.1 Determining Server That Performs Replication](#)

2. 6.2.2 Determining Replication Target

Refer to "6.2.3 Notes on Configuration of SnapOPC/SnapOPC+ Replication", when SnapOPC/SnapOPC+ snapshot replication is used. To backup or restore a system volume, refer to "Appendix D System Volume Backup Operation".

6.2.1 Determining Server That Performs Replication

Determine the server that performs replication.

The following servers are used to perform replication:

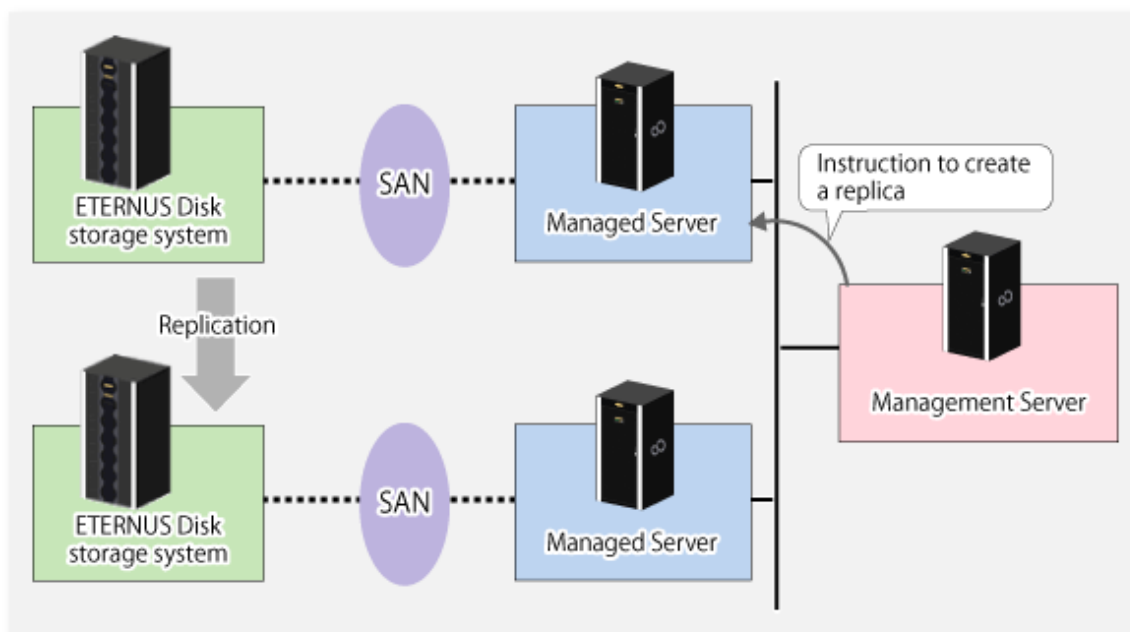
Management Server

Multiple Managed Servers are unified and operated centrally. The Management Server can also be used as a Managed Server.

Managed Server

AdvancedCopy Manager operations are performed.

Figure 6.9 Configuration of Servers That Perform Replication



Note

- The following requirements apply when an inter-box replication is performed using the REC function of ETERNUS Disk storage system:
 - The REC function must be supported in both ETERNUS Disk storage systems.
 - Both ETERNUS Disk storage systems must already be connected to each other using an FC Remote Adapter (FCRA). This is a connection that uses an FCRA, in which data flows only from the Initiator side to the Target side. For bi-directional copying, at least two FCRA connections are required.

6.2.2 Determining 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 source and replica volumes:

- Copy direction

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

Bi-directional

Data can be copied in either direction between source volume and replica volume.

Uni-directional

Data can only be copied from the source volume to the replica volume. During this operation, copying from the replica to the original volume is disabled.

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

In the case of server-to-server replication, you can specify whether to allow replication to be performed on either the source server or destination server, or on both servers.

Original Server

Sets the Managed Server with the source volume that was connected as the operational server.

Replica Server

Sets the Managed Server with the replica volume that was connected as the operational server.

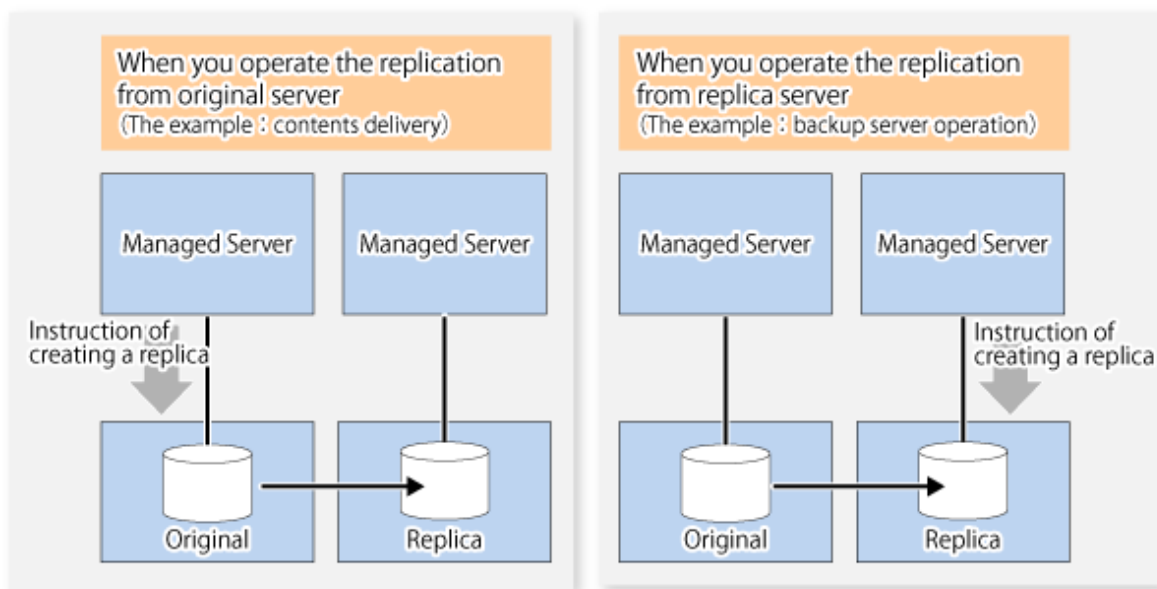
Both Servers

Sets both the source and destination servers as operational servers.

It is possible to create a group by arranging multiple source volumes and destination volumes.

In addition, the operation for multiple source volumes and destination volumes can be performed in a group. For details, refer to "6.4.7 Creating Groups".

Figure 6.10 Servers Used in Performing Replication



In this manual, a Managed Server used to perform replication is called the operational server. An original server or replica server is specified as the operational server.

An operational server can perform all of the replication functions. In contrast, any other type of server can use only the information display function and operational release function.



Note

In order to set up a volume as a source volume, there are some points which should be taken into consideration. For details, refer to "10.1.1 General Notes", for notes on determining a replication target.

- [10.1.1.1 Backup and Replication Exclusions](#)
- [10.1.1.6 Slice Containing VTOC](#)
- [10.1.1.5 Maximum Number of Simultaneously Executing Operations](#)

6.2.2.1 Logical Unit (Disk) Operation

Replication operations can copy data in Logical Units (disk units).

Logical Unit (disk) unit copying supports the following disk formats:

- EFI labeled disks
- VTOC labeled disks

To execute Logical Unit (disk unit) copy, refer to "10.1.1.11 Notes on Executing Logical Unit (Disk) Unit Copy".



Note


To copy a VTOC labeled disk, the second slice, indicating the entire disk, and the -Xvtoc option must be specified in setting of the original volume and replica volume.

6.2.3 Notes on Configuration of SnapOPC/SnapOPC+ Replication

6.2.3.1 Notes on Session Configurations


Only a single session can be set for the copy destination disk of SnapOPC/SnapOPC+.

Figure 6.11 When a Single Session Is Set for Copy Destination Disk




Multiple sessions cannot be set for a single copy destination disk, as shown in the following figure.

Figure 6.12 When Multiple Sessions Are Set for a Single Copy Destination Disk



For a SnapOPC copy source disk, you can set the session mixed with an SDV copy destination disk and a TPV/FTV copy destination disk. However, for a SnapOPC+ copy source disk, that session cannot be set.

Figure 6.13 When Multiple Sessions Mixed with SDV and TPV Are Set for Copy Destination Disk



6.2.3.2 Notes on Copy Destination Disk

If the physical capacity of the copy destination volume is insufficient, you cannot access the copy destination volume. Because of that, after estimating the physical capacity required for the copy destination volume according to the following procedure, create the copy destination disk.

1. [Estimated Required Capacity for Copy Destination Volumes](#)
2. [Creating of Copy Destination Disk](#)

Estimated Required Capacity for Copy Destination Volumes

To estimate the physical capacity required for the copy destination volume, you need to measure the update amount of the copy source volume. The measurement methods vary depending on the operation.

- In the case of one-to-one replication operation (SnapOPC/SnapOPC+)
Start a session in relation to the pair, and measure the update amount that is generated in relation to the replication source volume up until the session is stopped or a new session is started.
- In the case of one-to-many replication operation (SnapOPC)
Start a session in relation to all of the pairs, measure the update amount that is generated in relation to the replication source volume up until the session is stopped or a new session is started, and total all of the update amounts.
- In the case of one-to-many replication operation (SnapOPC+)
Start a session in relation to all of the pairs, measure the update amount that is generated in relation to the replication source volume up until SnapOPC+ is performed to the next generation, and total all of the update amounts.

For the update amount of the copy source volume, measure the following procedure using the [swstestupdate](#) command:

1. Start measuring the update amount by configuring a monitor session on the copy source volume.

```
# /opt/FJSSVswsts/bin/swstestupdate start /dev/dsk/c1t0d0s1
/dev/dsk/c1t0d0s1 swstestupdate completed
#
```

2. Start transactions.
Blocks that are updated by transactions are recorded on the hardware.
3. After the measurement period has elapsed, check the number of blocks that have been updated.

```
# /opt/FJSSVswsts/bin/swstestupdate status /dev/dsk/c1t0d0s1
Volume-Name      Update
/dev/dsk/c1t0d0s1 644333
#
```

4. After the measurements are complete, cancel the monitor session.

```
# /opt/FJSSVswsts/bin/swstestupdate stop /dev/dsk/c1t0d0s1
/dev/dsk/c1t0d0s1 swstestupdate completed
#
```

The formula for estimating the capacity is as follows.

$(\text{Updated volume of replication source volume}) * (\text{safety coefficient}) < (\text{Physical capacity of the copy destination volume})$
--

Point

.....
If an advance estimation is difficult, it is recommended that the physical capacity of the copy destination volume is 30 to 50 % of the total capacity of the copy source. As this is a recommended value, you must adjust according to the operational status.
.....

Creating of Copy Destination Disk

To create a copy destination disk, use Storage Cruiser or ETERNUS Web GUI.

Point

- Make logical space on a copy destination volume equal to that on a copy source volume.
- Consider a disk expansion from the required physical capacity in the copy destination volume.
 - If TPV is used as the copy destination volume

Use Storage Cruiser or ETERNUS Web GUI to check the status of the Thin Provisioning Pool and expand the capacity of the Thin Provisioning Pool.

For the operation procedure when using Storage Cruiser, refer to "Display Thin Provisioning Pool" and "Expand Capacity of/Format/Change Threshold Value of/Delete Thin Provisioning Pool" in the *ETERNUS SF Web Console Guide*.

- If FTV is used as the copy destination volume

Use Storage Cruiser or ETERNUS Web GUI to check the status of the Tier pool and expand the sub-pool capacity of the Tier pool. For the operation procedure when using Storage Cruiser, refer to "Display Tier Pool" and "Expand Capacity of Sub-Pool in Tier Pool" in the *ETERNUS SF Web Console Guide*.

- If SDV is used as the copy destination volume

Use the `swstsdv` command or ETERNUS Web GUI to check the status of the SDP and expand the SDP capacity.

Creating an exclusive volume named Snap Data Pool Volume (SDPV) enables the SDP and the created SDPV is automatically incorporated in the SDP. Creating an SDPV of the physical capacity that is assigned to a copy destination volume expands the SDP capacity.

The operation procedure when using the `swstsdv` command is as follows:

1. Execute the command with the "poolstat" subcommand and check the status of the SDP.
2. Create the SDPV with ETERNUS Web GUI.

- If you use an SDV as the copy destination volume, initialize the SDV. To avoid wasting the physical space of the SDV, limit the creation of file systems to the minimum requirement for the update after initializing the SDV.

See

Refer to "Create Volume" in the *ETERNUS SF Web Console Guide* for information on how to create TPV/SDV/SDPV using Storage Cruiser. Refer to "Create FTV" in the *ETERNUS SF Web Console Guide* for information on how to create FTV.

6.2.3.3 Monitoring Usage of Copy Destination Volume

If the physical capacity of the copy destination volume is insufficient, you cannot access the copy destination volume. This status makes it impossible to read data from the copy destination volume and write data into the copy destination volume. As for SnapOPC+, read/write is not allowed not only on the copy destination volume, but also on earlier generation's copy destination volumes.

For this reason, you must monitor so that a capacity shortage of the physical space in the copy destination volume does not occur during operation.

Monitoring When Using TPV/FTV as Copy Destination Volume

The physical space on the TPV/FTV is assigned per block. This may require more physical space than estimated. To prevent a capacity shortage of the physical space in the copy destination volume, use Storage Cruiser to change the settings of capacity threshold as required and monitor the usage.

Refer to "Thin Provisioning Management" in the *ETERNUS SF Storage Cruiser Operation Guide* for details of TPV monitoring.

Refer to "Automated Storage Tiering Management" in the *ETERNUS SF Storage Cruiser Operation Guide* for details of FTV monitoring.

Point

When you use the TPV/FTV as a copy destination volume, if the free space of the pool to which the TPV/FTV of the copy destination is created (Thin Provisioning Pool for TPV, Tier pool for FTV) is insufficient, you cannot access the copy destination volume. For this reason, when you use TPV/FTV as a copy destination volume, monitor so that a capacity shortage of the free space in the TPV/FTV and the pool to which the TPV/FTV of the copy destination is created does not occur.

Monitoring When Using SDV as Copy Destination Volume

- If only the SDV is used (the SDP is not configured)
Execute the [swstsdv](#) command with the "stat" subcommand to periodically monitor whether there is a shortage in the SDV capacity.
- If the SDP is used (the SDP is configured)
Execute the [swstsdv](#) command with the "poolstat" subcommand to periodically monitor whether there is a shortage in the SDP capacity.

In the SDP area, if there is an area that exceeds 50 % of usage ratio regardless of the encryption, consider increasing SDPVs. In addition, increase the monitoring frequency of the SDP area.

In the SDP area, if there is an area that exceeds 70 % of usage ratio regardless of the encryption, increase SDPVs immediately.

Point

If the physical capacity of the copy destination volume is insufficient, refer to "[8.4.2.3 Troubleshooting When Lack of Free Physical Space Has Occurred in Copy Destination Volume](#)" and take appropriate action.

6.2.3.4 Notes When Using SDV as Copy Destination Volume

Creating Access Volume


If the capacity of the SDV or the SDP is insufficient, you cannot access the copy destination volume. In the following cases, you cannot access the copy destination volume either:

- Access to the SDV after the SDV initialization
- Access to the SDV after the disk management information initialization

For this reason, it is recommended that you reserve a volume (Access volume) that is separate from the SDV. Create the volume that is used as the Access volume in the ETERNUS Disk storage system that the SDV has been created. Since the Access volume is not used in the storage of data, there is no problem even if there is little capacity. Any size can be used for the creation of the Access volume.

The following operations can be executed by creating the Access volume:

- [Snap Data Volume Initialization](#)
- [Snap Data Volume Status Display](#)
- [Snap Data Pool Status Display](#)



6.2.3.5 Notes on Cluster Operations

Do not use the copy destination volume of SnapOPC/SnapOPC+ as the shared disks in a clustered system. This is to avoid the clustered system failing over when the physical capacity is insufficient.


In cluster operations, use the following methods to prevent the copy destination volume of SnapOPC/SnapOPC+ from being used as a shared disk:

- Permit the copy destination volume of SnapOPC/SnapOPC+ to be referenced from all the nodes in the clustered system.
- Use inter-server replication between the clustered system and a non-clustered system.

6.3 Flow of Operations

The following figure shows the flow of the replication.

Figure 6.14 Flow of Replication Operations



6.4 Preparations

This section describes the preparations that are required before performing replication.

6.4.1 Starting Daemon

Before replication is started, the AdvancedCopy Manager daemon must be started on the Management Server and Managed Server. Normally, the daemon is configured to be automatically invoked at system startup.

However, if a daemon fails to start for some reason or has stopped, both servers must restart their respective daemons. For information about starting daemons, refer to "[Chapter 2 Starting and Stopping Daemons](#)".

6.4.2 Starting Web Console

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

6.4.3 Registering Managed Server

The Management Server registers the Managed Server to be managed. This process can be executed using Web Console.

Note

- Change to the value of 800 characters or less before registering the Managed Server when the value of the environment variable PATH of the Management Server is more than 800 characters.
 - Perform this task after confirming that it is possible to communicate from the Managed Server to the IP address of the Management Server. Check the Management Server IP address with the [stgxfwcmdisprv](#) command.
-

6.4.4 Fetching Device Information on Managed Server

This process can be executed using either Web Console or the [stgxfwcmsetdev](#) command. Refer to "Add/Reload Device" in the *ETERNUS SF Web Console Guide* for information on the operating procedures of Web Console.

Note

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

6.4.5 Configuring Source Volume and Replica Volume

Execute the [swsrpsetvol](#) command to set the source volume and replica volume that are used in the replication. The information on the registered source volume and replica volume can be checked with the [swsrpvolinfo](#) command.

When specifying the volume of another Managed Server in server-to-server replication, specify the volume in the format "*Volume-Name@Managed-Server-name*".

6.4.6 Customizing Pre-processing and Post-processing

In AdvancedCopy Manager, replication processing must be performed while volumes are unmounted. Therefore, when replication is executed, processing to mount and unmount the volume is performed by the pre-processing and post-processing scripts respectively.

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

For details of the pre-processing and post-processing, refer to "[Appendix C Pre-processing and Post-processing of Replication](#)".

Replication pre-processing and post-processing need to be customized when operational requirements change.

For example:

- For specific operational requirements, you want to avoid the unmount/mount processing of the volume where a file system is constructed.
- If you want to add special pre-processing or post-processing.

Refer to "[Appendix C Pre-processing and Post-processing of Replication](#)" for information on the customization.

Note

- The pre-processing and post-processing are not performed in case of replicating a group. Unmount all the volumes in the group before replication, and mount all the volumes after replication.

- Pre-processing and post-processing are not executed for the slices included in a Logical Unit (disk) when Logical Unit (disk) unit copying is executed. Therefore, the slices included in a Logical Unit (disk) must be unmounted before a replica is created, and the unmounted slices must be mounted after the replica is created.

P Point

When AdvancedCopy Manager Is Upgraded


The script may need to be updated after an upgrade of AdvancedCopy Manager. Therefore, you should check the script after any upgrade to ensure it is compatible with newer versions.

6.4.7 Creating Groups


6.4.7.1 Conditions for Replication Volumes That Make Up Group

The replication volumes that make up the group (the source volume and the destination volume pairs) must meet the following conditions:


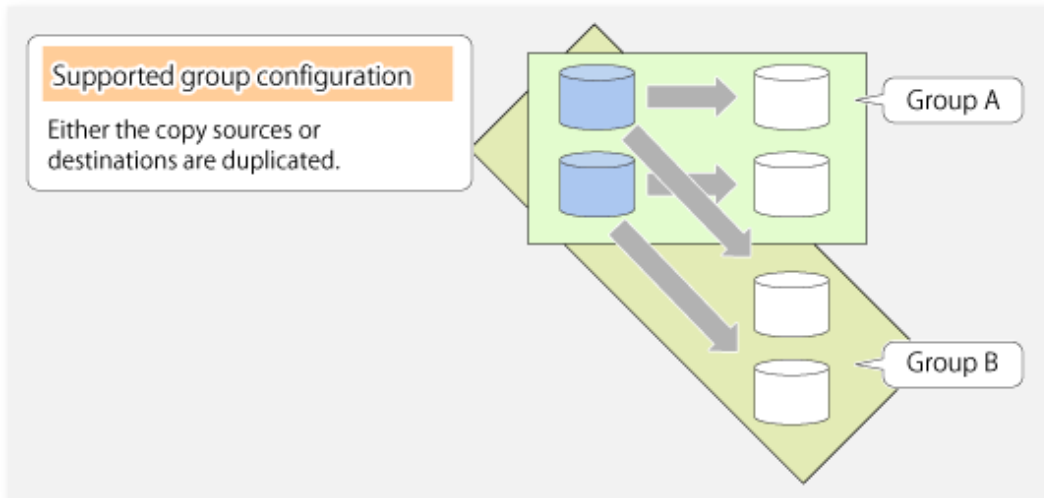
- The copy source servers and destination servers must have matching attributes such as size and device configuration.





- The operational servers and the copy directions must all match. The values for the operational servers and the copy directions can be checked using the **Op-Server** and **Copy** columns displayed by executing the `swsrpvolinfo` command.




- The replication volume information being registered must not be registered in any other group.



- Source volumes and destination volumes must not be duplicated within the group.



- A replication group cannot have both PRIMECLUSTER GD logical volumes and other types of volumes.



- If a Logical Unit (disk) is specified, the information of the replication volume in a partition included in the Logical Unit (disk) must not be registered to another group.
- If a Logical Unit (disk) is specified, the information of the replication volume in a partition included in the Logical Unit (disk) must not be registered within the same group.
- If the Storage Cluster Continuous Copy function is used, all of the replication volume information in the group must belong to the same TFO group.

6.4.7.2 How to Create Groups

Groups are created using the [swsrpsetvol](#) command.

Information for created groups can be displayed using the [swsrpvolinfo](#) command.

Example

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

```
# /opt/FJVSwsrp/bin/swsrpsetvol -Xgroup GRP1 /dev/dsk/c1t1d1s4@SRC /dev/dsk/
c1t1d1s4@TARG-1
swsrpsetvol completed
# /opt/FJVSwsrp/bin/swsrpsetvol -Xgroup GRP1 /dev/dsk/c1t1d2s4@SRC /dev/dsk/
c1t1d1s4@TARG-1
swsrpsetvol completed
```

```
# /opt/FJSVswsrp/bin/swsrpvoinfo -L
Server Original-Volume      Size      Replica-Volume      Size
Copy      Op-Server Group
SRC      /dev/dsk/c1t1d1s4@SRC 4.0 Gbyte /dev/dsk/c1t1d1s4@TARG-1 4.0 Gbyte bi-
direction both      GRP1
SRC      /dev/dsk/c1t1d2s4@SRC 4.0 Gbyte /dev/dsk/c1t1d2s4@TARG-1 4.0 Gbyte bi-
direction both      GRP1
#
```

6.5 Operations

Before performing replication, refer to "6.4 Preparations" to set up the environment required for replication.

This section describes the operation by command.

6.5.1 Executing Snapshot Replication

Use the `swsrpmake` command to perform snapshot replication.

Refer to "6.1.1 Snapshot Replication Processing" for an explanation of snapshot replication.

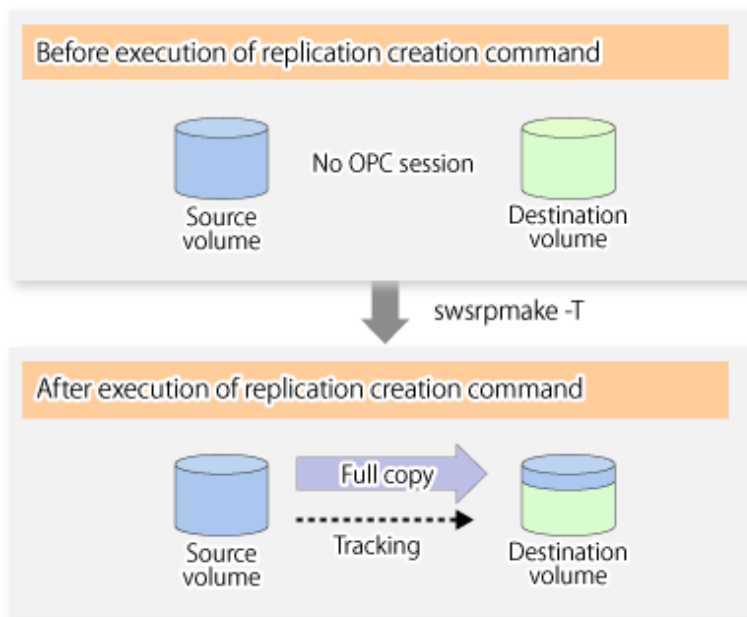
The operation status of a physical copy can be checked by executing the `swsrpstat` command.

QuickOPC Type Replication

Execute QuickOPC replication by specifying the `-T` option in the `swsrpmake` command.

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

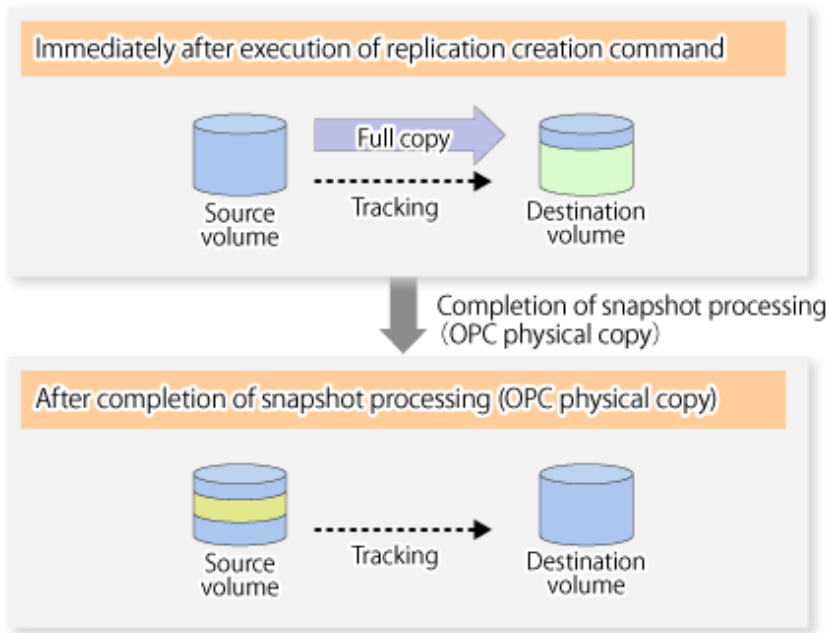
Figure 6.15 When Replication Creation Command Is Executed (1st Time)



To check the execution status of physical copying, use the `swsrpstat` command in the same way as for an ordinary snapshot replication. After snapshot processing (OPC physical copy) is complete, only tracking processing is active.


To check the tracking status, execute the `swsrpstat` command with the `-L` option.

Figure 6.16 When Snapshot Processing Is Completed



Entering the `swsrpmake` command with the `-T` option specified during tracking processing performs the physical copying of only the data that has been generated since the previous snapshot processing. This means that physical copying can be accomplished in a shorter period of time.

Figure 6.17 When Replication Creation Command Is Executed (2nd Time)



When you want to perform a restoration while tracking processing is being executed, you need to perform a restoration by OPC (to achieve this, you need to execute the `swsrpmake` command without the `-T` option). QuickOPC cannot be executed in the reverse direction while tracking processing is being executed. The replication using QuickOPC is done as follows:

```
[backup]
swsrpmake -T <original volume name> <replica volume name>

[restore]
swsrpmake <replica volume name> <original volume name>
```

Although a restoration is executed with OPC, only the data that has been updated since the previous replication (it can be obtained from the **Update** 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. The restore execution status can be checked by executing the `swsrpstat` command with the `-E` option specified.

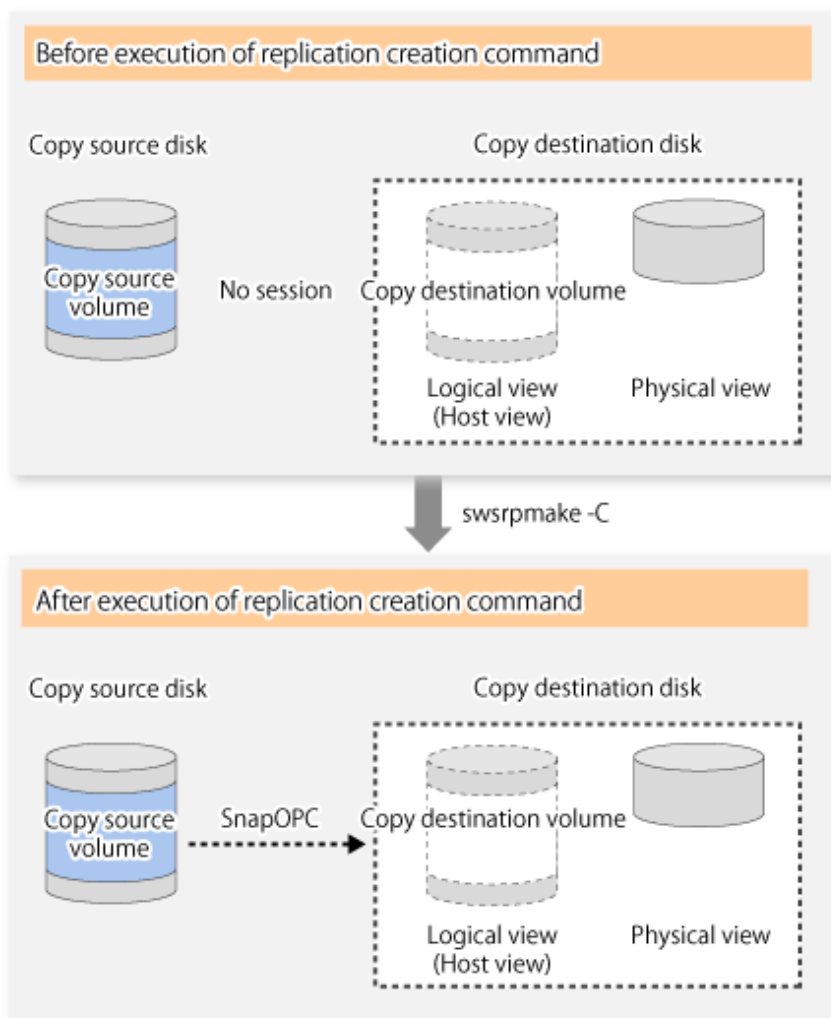
SnapOPC Type Replication

Execute SnapOPC type replications with the `-C` option specified in the `swsrpmake` command. When the `swsrpmake` command is executed, a SnapOPC session is set up between the copy source volume and the copy destination volume.

Example

```
# /opt/FJSVswsrp/bin/swsrpmake -C /dev/dsk/c1t0d0s1 /dev/dsk/c1t0d11s1
FROM=/dev/dsk/c1t0d0s1@SV1,TO=/dev/dsk/c1t0d11s1@SV1 swsrpmake completed
#
```

Figure 6.18 When Replication Creation Command Is Executed



Unlike normal OPCs and QuickOPCs, SnapOPCs do not copy all of the data from the source volume, but instead copy only the data that has been updated on the source or destination since SnapOPC started. This kind of copy processing is referred to as "Copy-on-Write".

Figure 6.19 When Copy Source Volume Is Updated





Figure 6.20 When Copy Destination Volume Is Updated



Note: The units for host I/O and storage device copies are different, and therefore data copies also occur when the copy destination is updated.

The status of SnapOPC sessions can be checked using the `swwrstat` command.

The following example shows the execution of the `swwrstat` command immediately after a SnapOPC snapshot has started. While SnapOPC is being performed, "copy-on-write" is displayed in the **Status** column, and the amount of data updated since the last copy was created is displayed in the **Update** column as a percentage.

Example

```
# /opt/FJSVswsrp/bin/swwrstat -L /dev/dsk/clt0d0s1
Server Original-Volume      Replica-Volume      Direction Status      Execute Trk  Update Rcv
Split Xfer Snap-Gen
SV1 /dev/dsk/clt0d0s1@SV1 /dev/dsk/clt0d1s1@SV1 regular copy-on-write ---- off 0% ----
-----
#
```

If the `swwrmake` command is executed again during SnapOPC processing, the SnapOPC session that has already been set up is cancelled, and a new session is set up.

Note

If the physical capacity of the copy destination volume is insufficient, the SnapOPC execution status changes to error suspend status ("failed"), and replication volume cannot be used.

The SnapOPC execution status can be checked in the `swwsrpstat` command output result Status column.

Example

```
# /opt/FJSVswsrp/bin/swwsrpstat -L /dev/dsk/c1t0d0s1
Server Original-Volume      Replica-Volume      Direction Status Execute Trk  Update Rcv  Split
Xfer Snap-Gen
SV1    /dev/dsk/c1t0d0s1@SV1 /dev/dsk/c1t0d11s1@SV1 regular   failed ----  off  ----  ----  ----
-----
#
```

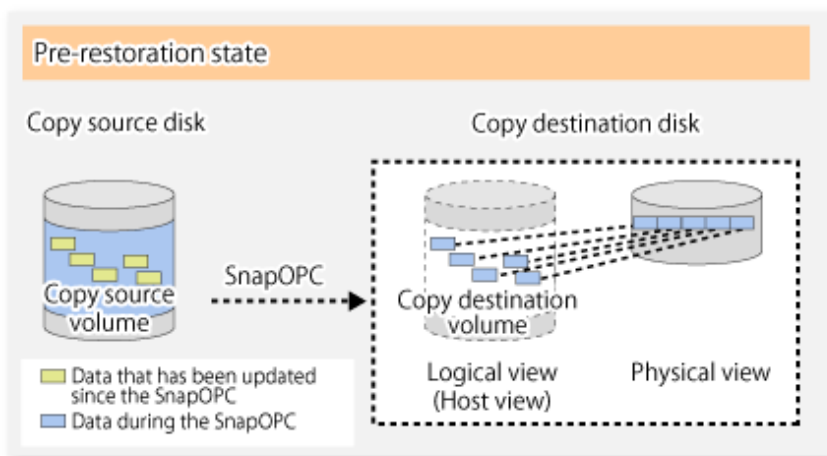
When the SnapOPC execution status is error suspend status ("failed"), refer to ["8.4.2.3 Troubleshooting When Lack of Free Physical Space Has Occurred in Copy Destination Volume"](#) and take appropriate action.

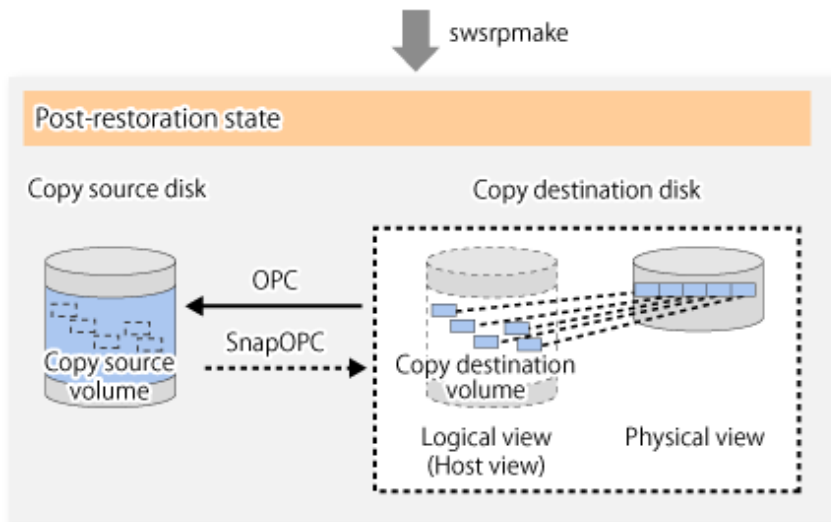
Perform restorations from the copy destination volume by running an OPC using the `swwsrpmake` command.

```
# /opt/FJSVswsrp/bin/swwsrpmake /dev/dsk/c1t0d11s1 /dev/dsk/c1t0d0s1
FROM=/dev/dsk/c1t0d11s1@SV1,TO=/dev/dsk/c1t0d0s1@SV1 swwsrpmake completed
#
```

When restorations are executed, the SnapOPC session from the source volume to the destination volume is maintained as is, and a normal OPC from the replication destination volume to the replication source volume is started. At this point, the time taken to restore the physical copy is reduced, because only data that has been updated since the last copy is restored.

Figure 6.21 When Restoration Is Executed






The execution status of restorations can be checked by specifying the -E option with the [swsrpstat](#) command.

```
# /opt/FJVSwsrp/bin/swsrpstat -E /dev/dsk/clt0d0s1
Server Original-Volume      Replica-Volume      Direction Status Execute
SV1    /dev/dsk/clt0d0s1@SV1    /dev/dsk/clt0d11s1@SV1 reverse  snap   80%
#
```

Note


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

Figure 6.22 When Restoring to Volume Other Than Copy Source Volume




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

Figure 6.23 When SnapOPC Is Performed for Multiple Copy Destination Volumes



In this case, restoration using an OPC can be performed by cancelling the other SnapOPCs. However, the backup data on the copy destination volumes whose SnapOPC sessions were cancelled is lost.


Figure 6.24 When SnapOPC Session Cancelled to Perform Restoration



To perform a restoration while still maintaining all SnapOPC sessions, operating system copy functions (such as the cp command or the copy command) must be used for the restoration.

However, if restoration is performed using operating system functions, the amount of updated data on the source volume increases, and there is a risk that the capacity of the SnapOPC volume is insufficient.

Figure 6.25 When Performing Restoration without Cancelling SnapOPC Session



SnapOPC+ Type Replication

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

An execution example of the `swsrpmake` command with the `-P` option is shown below.

Execution Example

```
# /opt/FJJSVswsrp/bin/swsrpmake -P /dev/dsk/c1t1d0s3@SV1 /dev/dsk/c1t1d1s3@SV1
FROM=/dev/dsk/c1t1d0s3@SV1,TO=/dev/dsk/c1t1d1s3@SV1 swsrpmake completed
#
```

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

The next time this command is executed with a different copy destination volume for the same copy source volume, the copy-on-write processing being executed between the copy source volume and the previous generation of the copy destination volume is stopped. Then, a SnapOPC+ session is set between the copy source volume and the newly specified copy destination volume, and copy-on-write is performed.

An execution example of the `swsrpmake` command with the `-P` option for the newly specified copy destination volume is shown below.

Execution Example

```
# /opt/FJJSVswsrp/bin/swsrpmake -P /dev/dsk/c1t1d0s3@SV1 /dev/dsk/c1t1d2s3@SV1
FROM=/dev/dsk/c1t1d0s3@SV1,TO=/dev/dsk/c1t1d2s3@SV1 swsrpmake completed
#
```

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


Similarly, each time there is a new copy destination volume, a snap generation number is assigned.



Note

If an earlier snap generation (other than the oldest snap generation) is specified as the copy destination volume when the `swsrpmake` command is executed, the command terminates with an error. If the oldest snap generation is specified as the copy destination volume, that snap generation is automatically discarded and a replica is created as the newest snap generation. In this case, subsequent snap generations (second, third) are assigned a snap generation number that is one generation prior (second generation => first generation, and third generation => second generation).

Figure 6.26 When Oldest Snap Generation Number Is Specified as Replication Volume



The operation status of SnapOPC+ replication can be checked by executing the `swsrpstat` command with the `-L` option. 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 /dev/dsk/clt1d0s3
Server Original-Volume      Replica-Volume      Direction Status      Execute Trk
Update Rcv  Split Xfer Snap-Gen
SV1    /dev/dsk/clt1d0s3@SV1    /dev/dsk/clt1d1s3@SV1 regular  copy-on-write(inactive) ---- off
0%     ---- ---- ---- 1
SV1    /dev/dsk/clt1d0s3@SV1    /dev/dsk/clt1d2s3@SV1 regular  copy-on-write(active) ---- off
5%     ---- ---- ---- 2
#
```

Note

If the physical capacity of the copy destination volume is insufficient, the SnapOPC+ execution status changes to error suspend status ("failed"), and the execution status of SnapOPC+ that was executed before it also changes to error suspend status ("failed"). Replication volume of error suspend status ("failed") cannot be used.

The SnapOPC+ execution status can be checked in the `swsrpstat` command output result **Status** column.

Execution Example

```
# /opt/FJSVswsrp/bin/swsrpstat -L /dev/dsk/clt1d0s3
Server Original-Volume      Replica-Volume      Direction Status Execute Trk Update Rcv  Split
Xfer Snap-Gen
SV1    /dev/dsk/clt1d0s3@SV1    /dev/dsk/clt1d1s3@SV1 regular  failed ---- off ---- ---- ----
---- ----
SV1    /dev/dsk/clt1d0s3@SV1    /dev/dsk/clt1d2s3@SV1 regular  failed ---- off ---- ---- ----
---- ----
#
```

When the SnapOPC+ execution status is error suspend status ("failed"), refer to "[8.4.2.3 Troubleshooting When Lack of Free Physical Space Has Occurred in Copy Destination Volume](#)" and take appropriate action.


To restore from the copy destination volume, execute the `swsrpmake` command to start OPC.

```
# /opt/FJSVswsrp/bin/swsrpmake /dev/dsk/c1t1d2s3@SV1 /dev/dsk/c1t1d0s3@SV1
FROM=/dev/dsk/c1t1d2s3@SV1,TO=/dev/dsk/c1t1d0s3@SV1 swsrpmake completed
#
```

The SnapOPC+ session from the replication source volume to the replication destination volume is maintained even if the replication creation command is executed.

Execution of restoration while maintaining the SnapOPC+ session reduces the physical copying time, because physical copying is performed only for data updated after the replica creation.

Figure 6.27 Restoration with SnapOPC+ Session Maintained



To check the restoration execution status, execute the `swsrpstat` command with the `-E` option.

```
# /opt/FJSVswsrp/bin/swsrpstat -E /dev/dsk/clt1d0s3
Server Original-Volume      Replica-Volume      Direction Status Execute
SV1    /dev/dsk/clt1d0s3@SV1    /dev/dsk/clt1d1s3@SV1    ----    ----    ----
SV1    /dev/dsk/clt1d0s3@SV1    /dev/dsk/clt1d2s3@SV1    reverse  snap    80%
#
```

Note

Restoration may cause a capacity shortage of the copy destination volume, due to updates issued to the most recent snap data generation where the copy-on-write status is active. Make sure that there is enough free space in the copy destination volume usage area before performing restoration.

The most recent snap generation is the data written to the replication source volume by the restoration, updated by the previously existing data. The update amount to the most recent snap generation generated by the restoration is the total of the Copy usage amount for the restoration target snap generation and subsequent snap generations except for the most recent snap generation.

An example of how to check when restoring from snap generation (**Snap-Gen**) 2 is displayed below.

Use the procedure below to check that there is enough free space in the copy destination volume usage area:

1. Execute the `swsrpstat` command to check the device name of the restoration target and subsequent snap generations, except for the most recent snap generation (**Snap-Gen** 4 data in the example below).

```
# /opt/FJSVswsrp/bin/swsrpstat -L /dev/dsk/c1t1d0s3
Server Original-Volume      Replica-Volume      Direction Status          Execute Trk
Update Rcv  Split Xfer Snap-Gen
SV1 /dev/dsk/c1t1d0s3@SV1 /dev/dsk/c1t1d1s3@SV1 regular copy-on-write(inactive) ---- off
8% ---- ---- ---- 1
SV1 /dev/dsk/c1t1d0s3@SV1 /dev/dsk/c1t1d2s3@SV1 regular copy-on-write(inactive) ---- off
12% ---- ---- ---- 2
SV1 /dev/dsk/c1t1d0s3@SV1 /dev/dsk/c1t1d3s3@SV1 regular copy-on-write(inactive) ---- off
0% ---- ---- ---- 3
SV1 /dev/dsk/c1t1d0s3@SV1 /dev/dsk/c1t1d4s3@SV1 regular copy-on-write(active) ---- off
3% ---- ---- ---- 4
```

In this example, `/dev/dsk/c1t1d2s3` and `/dev/dsk/c1t1d3s3` are targeted.

2. Calculate the total usage capacity of the devices that was checked in step 1.

- If the copy destination volume is TPV

Use Storage Cruiser or ETERNUS Web GUI to calculate the total usage capacity of the devices in step 1.

If you use Storage Cruiser, the used capacity can be checked in the **Used Capacity** column in the **Volume** tab of the Thin Provisioning Details screen on Web Console. Refer to "Display Thin Provisioning Pool" in the *ETERNUS SF Web Console Guide* for information on how to check the used capacity.

- If the copy destination volume is FPV

Use Storage Cruiser or ETERNUS Web GUI to calculate the total usage capacity of the devices in step 1.

If you use Storage Cruiser, the used capacity can be checked in the **Used Capacity** column of the FTV screen on Web Console. Refer to "Display FTV" in the *ETERNUS SF Web Console Guide* for information on how to check the used capacity.

- If the copy destination volume is SDV

Execute the `swstsdv` command with the "stat" subcommand to calculate the total used capacity of the devices in step 1.

If SDP is used, also add the SDP usage capacity.

[/dev/dsk/c1t1d2s3 disk usage]

```
# /opt/FJSVswsts/bin/swstsdv stat /dev/dsk/c1t1d2s3
BoxID = 00E4000M3#####E450S20A####KD4030639004##
LUN = 110 (0x6E)
Rate Logical(sector) Physical(sector) Used(sector) Copy(sector) Host(sector) Pool(sector)
100% 8388608          1048576          1048576          1048384          192             640
```

[/dev/dsk/c1t1d3s3 disk usage]

```
# /opt/FJSVswsts/bin/swstsdv stat /dev/dsk/c1t1d3s3
BoxID = 00E4000M3#####E450S20A####KD4030639004##
LUN = 111 (0x6F)
Rate Logical(sector) Physical(sector) Used(sector) Copy(sector) Host(sector) Pool(sector)
4% 8388608          1048576          46928            16              46912           0
```

In this example, the quantity updated by the restoration is **1049040** (1048384+640+16) sectors.

3. Check the free capacity of the copy destination volume.

- If the copy destination volume is TPV

Use Storage Cruiser or ETERNUS Web GUI to check the free capacity of the Thin Provisioning Pool.

If you use Storage Cruiser, the free capacity can be checked in the **Total Capacity** column and the **Used Capacity** column of the Thin Provisioning Overview screen on Web Console. Refer to "Display Thin Provisioning Pool" in the *ETERNUS SF Web Console Guide* for information on how to check the capacity.

- If the copy destination volume is FTV

Use Storage Cruiser or ETERNUS Web GUI to check the free capacity of the Tier pool.

If you use Storage Cruiser, the free capacity can be checked in the **Total Capacity** column and the **Used Capacity** column of the Tier pool detail screen that is displayed by selecting the target Tier pool in the Tier pool Overview screen on Web Console. Refer to "Display Tier Pool" in the *ETERNUS SF Web Console Guide* for information how to check the capacity.

- If the copy destination volume is SDV

Execute the `swstdv` command with the "poolstat" subcommand to check the total capacity and the used capacity of SDP.

If SDV is not encrypted, check the total capacity and the used capacity where [Pool-Type] is "Normal".

If SDV is encrypted, check the total capacity and the used capacity where [Pool-Type] is "Encrypted".

```
# /opt/FJSVswstdv/bin/swstdv poolstat -G /dev/dsk/c1t1d0s3
BoxID = 00E4000M3#####E450S20A####KD4030639004##
Pool-Type Rate Total(sector) Used(sector) Copy(sector) Host(sector) Free(sector)
Normal 10% 20971520 2097152 0 2097152 18874368
Encrypted 0% 20971520 0 0 0 20971520
```


The disk usage in this example is 15% = $\sim (2097152 + 1049040) / 20971520 \times 100$

If the update amount with a restore is less than the free capacity of the copy destination volume, a restoration is possible. However, in order to safely perform a restoration, it is recommended that you increase the free capacity with a disk expansion if the disk usage after the restoration is predicted to exceed 70 %.

Since the required physical capacity is increased depending on the restoration, the capacity of the copy destination volume may be insufficient. Therefore, in order to prevent the physical capacity of the copy destination volume from becoming insufficient, refer to "6.2.3.3 Monitoring Usage of Copy Destination Volume" to review how to monitor the used capacity of the copy destination volume. In addition, consider increasing the number of disks as required.

 Note

If SnapOPC+ is being performed between the replication source volume and the replication destination volume, restoration cannot be performed to a volume other than the replication source volume.

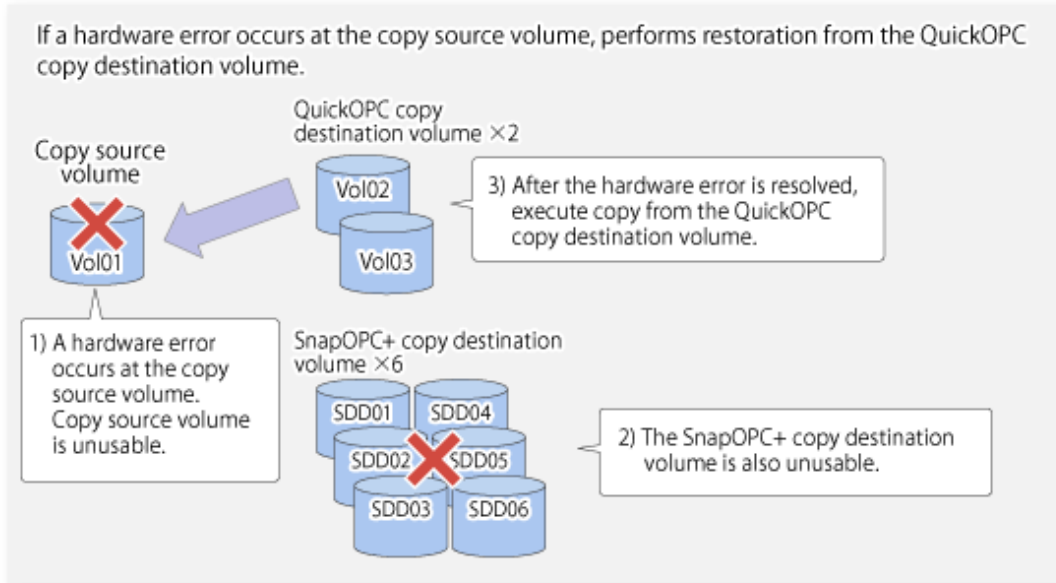



Point

As a precaution against hardware malfunctions with SnapOPC+, it is recommended to operate it in conjunction with making full copies using OPC/QuickOPC/EC(REC).

An example of performing QuickOPC on Sundays and SnapOPC+ on Mondays to Saturdays is displayed below.

Figure 6.28 Example of Operation Using SnapOPC+ and QuickOPC




6.5.1.1 Concurrent OPC Functions

Concurrent OPC function is one of ETERNUS Disk Storage system's function which creates snapshots of the multiple logical volumes simultaneously. This function allows backup the replication of the database consisting of multiple volumes in a consistent state. This function is available to use in one of these replications; OPC, QuickOPC, SnapOPC, or SnapOPC+.

The following diagram shows the operation in ETERNUS Disk Storage system.

Figure 6.29 Operations in ETERNUS Disk Storage System



Note

- If you force-quit the `swsrpmake` command with task manager or Ctrl-C, the session may remain defined Concurrent OPC. In this case, execute the command again, or stop the session with the `swsrpcancel` command.
- The maximum number of pairs that can simultaneously implement Concurrent OPC depends on the specifications of the ETERNUS Disk storage system.

6.5.2 Executing Synchronized Replication

To perform synchronized replication, use the following procedure:

1. Start synchronized processing using the [swsrpstartsync](#) command. To cancel synchronized processing that has already started, execute the [swsrpcancel](#) command.
2. After making sure that equivalency maintenance state has been established with the [swsrpstat](#) command, temporarily stop synchronized processing with the [swsrpmake](#) command to create a replica of the source volume.
3. To copy the updated (ie, incremental) data, restart synchronized processing with the [swsrpstartsync](#) command.

6.5.2.1 Intra-box Synchronous Replication

Intra-box synchronous replication creates a replication from a source volume to a destination volume by using the EC function of ETERNUS Disk storage system.

The EC function operates in a mode in which a copy is made to a destination volume in synchronization through a write to a source volume (ie, in synchronous mode).

It is also possible to reverse the copy direction from suspend status using the EC functionality.

6.5.2.2 Inter-box Synchronous Replication

Inter-box synchronous replication creates a replication from a source volume to a destination volume by using the REC function of ETERNUS Disk storage system.

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

- Transfer mode
- Recovery mode
- Split mode

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

Transfer Mode

The transfer mode provides the REC data transfer modes described below.

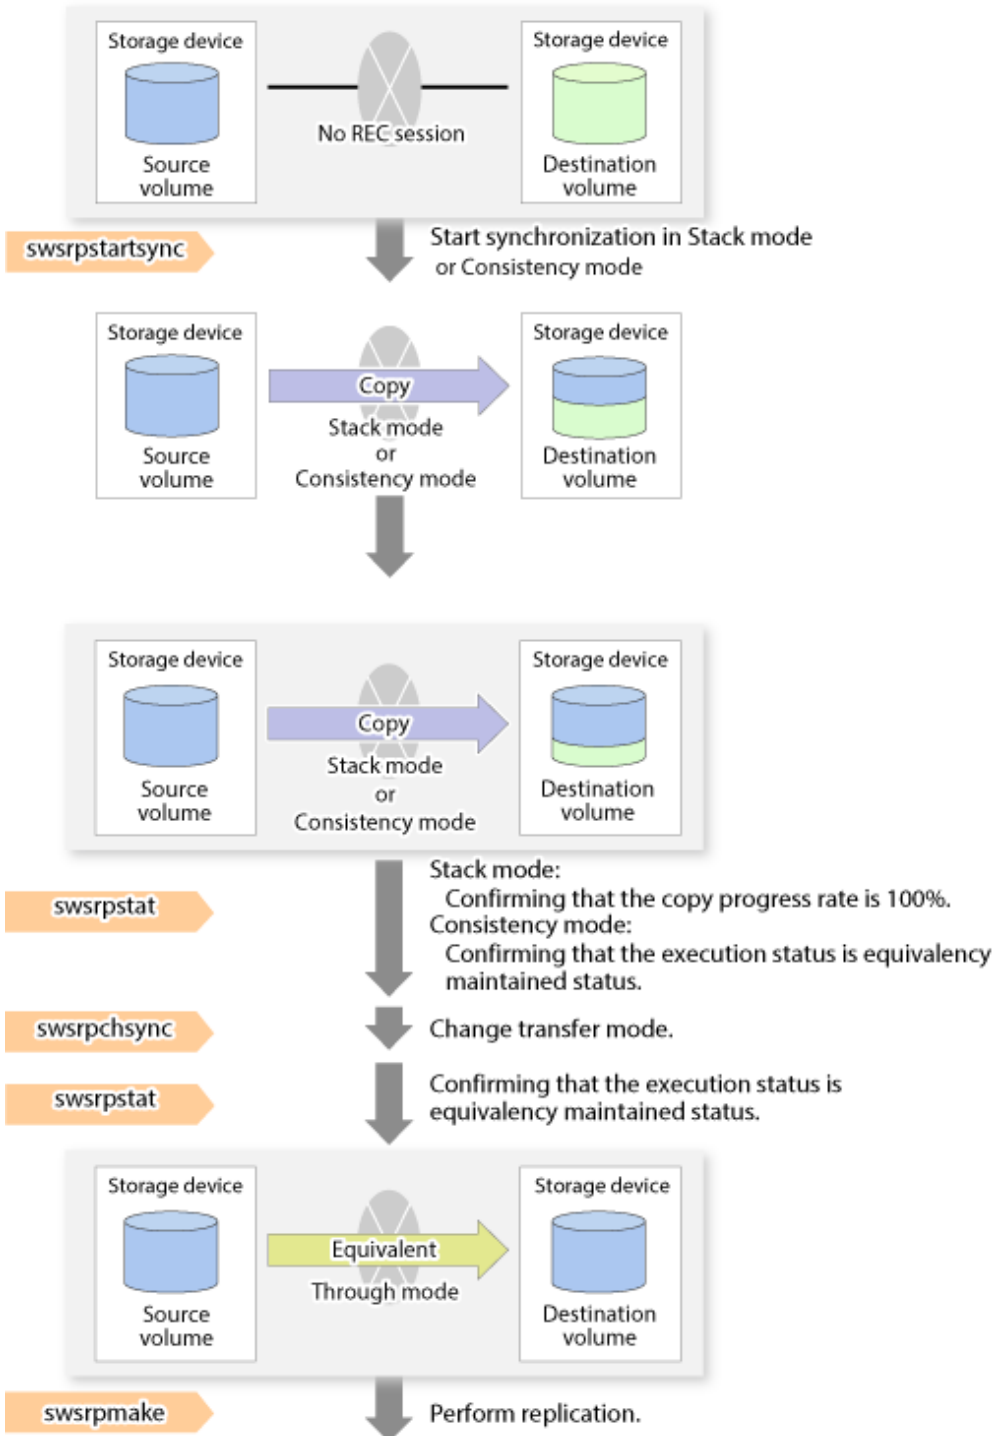
Table 6.2 Transfer Mode Types


Transfer Method	Transfer Mode	Explanation
Synchronous transfer method	Synchronous	<p>In this mode, responses to write requests from the server do not return as "complete" until data has been written to the copy source volume and copied to the destination volume.</p> <p>Synchronizing data copies with writes to the copy source guarantees the integrity of the data on both the copy source volume and the copy destination volume when the copy completes.</p> <p>This mode is suitable to RECs within sites with short delay times, because it has a large impact on the performance of write accesses from the server.</p>
Asynchronous transfer methods	Stack	<p>This mode has only a small impact on the time taken to respond to the server because it only records the position of the block that has been updated before returning "complete" to the server. The data for this block is transferred using an independent transfer engine.</p> <p>Copies can be executed even when there is only limited bandwidth, but in this case the amount of data that has not been transferred may build up.</p>
	Consistency	<p>This mode guarantees the sequence of transfers between copy sessions for the copy destination ETERNUS Disk storage system, and is suitable for operations, such as on databases, that perform mirroring using copies that consist of multiple areas.</p> <p>This mode uses part of cache memory for transfer buffers. Data is transferred to the copy destination by first "storing" a block of data to be copied in a transmission transfer buffer and then using a reception transfer buffer to "expand" the data that has been transferred.</p>

Transfer Method	Transfer Mode	Explanation
	Through	This mode is for transferring data that has not been transferred when copies in Stack mode or Consistency mode are stopped or suspended. If any data is updated on a replication volume, copy processing to transfer untransferred data may be given priority, delaying the response to Write request. Therefore, it is not recommended to use this mode during normal operations.

To perform a synchronous replication in Stack mode or Consistency mode, use the [swsrpstartsync](#), [swsrpmake](#) and [swsrpchsinc](#) commands. To check the execution status, use the [swsrpstat](#) command. The figures below show how synchronous replication is performed in Stack mode or Consistency mode.

Figure 6.30 Flow of Synchronous Replication (for Stack Mode / Consistency Mode)





Point

- When changing the transfer mode from the Consistency mode to the Through mode, implement the operations for changing the mode after confirming that the execution status is equivalency maintenance state. There may be cases when the execution state does not enter the equivalency maintenance state immediately after changing the mode. When the transfer mode has been changed, create a replication after confirming that the execution status is equivalency maintenance state.
- When changing the transfer mode from the Stack mode to the Through mode, implement the operations for changing the mode after confirming that the copy progress rate is 100%. Although the mode may be changed even when not in a state where the copy progress rate is 100% (a state where data that has not yet been transferred is remaining), the period of the Through mode is longer as the data that has not yet been transferred is transferred after the transfer mode has been changed.

Recovery Mode

Recovery mode consists of two methods of restarting the copying after recovery is performed from an inter-box path error (e.g., halt status).

Table 6.3 Recovery Mode Types

Mode	Description
Automatic Recovery	In this mode, the REC session automatically switches from halt status to regular status, and copy processing resumes when the inter-box FCRA path is recovered.
Manual Recovery	In this mode, the REC session remains in halt status and copy processing does not resume even if the inter-box FCRA path is recovered. Manually resume copying. This mode is used, for example, when operating a standby database.

Point

The REC restart method varies, depending on the halt status. Refer to "8.4.2.4 Error (halt) on Remote Copy Processing" for details.

Split Mode

Split mode includes two modes for performing a write operation to the source volume when REC is used for synchronous transfer mode, and for recovery from an inter-box path error (halt status).

Table 6.4 Split Mode Types

Mode	Description
Automatic Split	This split mode forcibly executes successful write operations to source volumes even if the inter-box FCRA path is fully blocked and halt status occurs. This mode enables write operations to source volumes even if the inter-box FCRA path is fully blocked, thus this mode does not affect transactions. When the inter-box FCRA path is recovered, copy processing resumes according to the recovery mode configurations.
Manual Split	This split mode rejects write operations to source volumes (it returns an error) if the inter-box FCRA path is fully blocked and halt status occurs. This mode enables source volumes and destination volumes to be fully synchronized even if the FCRA path is fully blocked. When the inter-box FCRA path is recovered, copy processing resumes according to the recovery mode configurations.


Copy Direction Reverser

The copy direction reverser allows you to smoothly switch between different sites.

The following figures show an example of how to switch the copy direction:

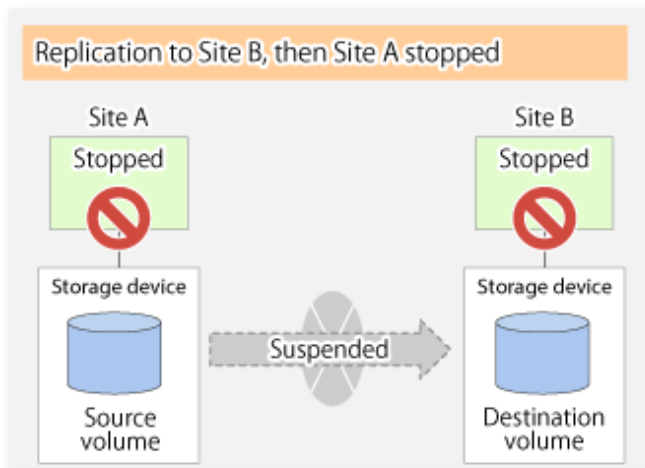
1. Assume that Site A is operating and REC is operating from Site A to Site B.

Figure 6.31 When REC Performed from Site A to Site B




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

Figure 6.32 When Replication Creation Command Used to Create Replica at Site B




- Execute the `swsrprevsync` command to reverse the copy direction.

Figure 6.33 When Synchronization Reverse Command Is Executed




- Bring Site B into operation. At this stage, synchronization is still suspended, thus any update to the volume at Site B is not reflected at Site A.

Figure 6.34 When Operation Is Started at Site B (Synchronization Suspend State)



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

Figure 6.35 When Operation Is Started at Site B (Synchronization Resumed)




6.5.2.2.1 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.


- Suppose that operations at Site A have stopped.

Figure 6.36 If Operations at Site A Have Stopped




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

Figure 6.37 When Synchronization Is Started by Initial Copy Skip Function




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

Figure 6.38 When Copy Source Data Is Backed Up to Tape




- The tape media is sent to Site B and jobs at Site A restart.

Figure 6.39 When Site A Operations Restart



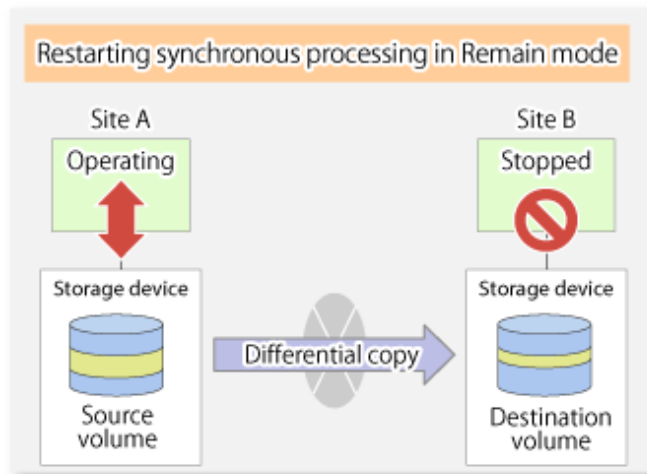
- The data on the tape media is restored to the copy destination volume. At this point, the data on the copy destination volume is the same as the data on the copy source volume that existed before operations restarted.

Figure 6.40 When Data Is Restored from Tape to Copy Destination



- Synchronous processing restarts in Remain mode. Restarting synchronous processing in Remain mode means that only data that has been updated on the copy source volume is reflected on the copy destination volume. If Remain mode is not used, all data on the copy source volume is copied.

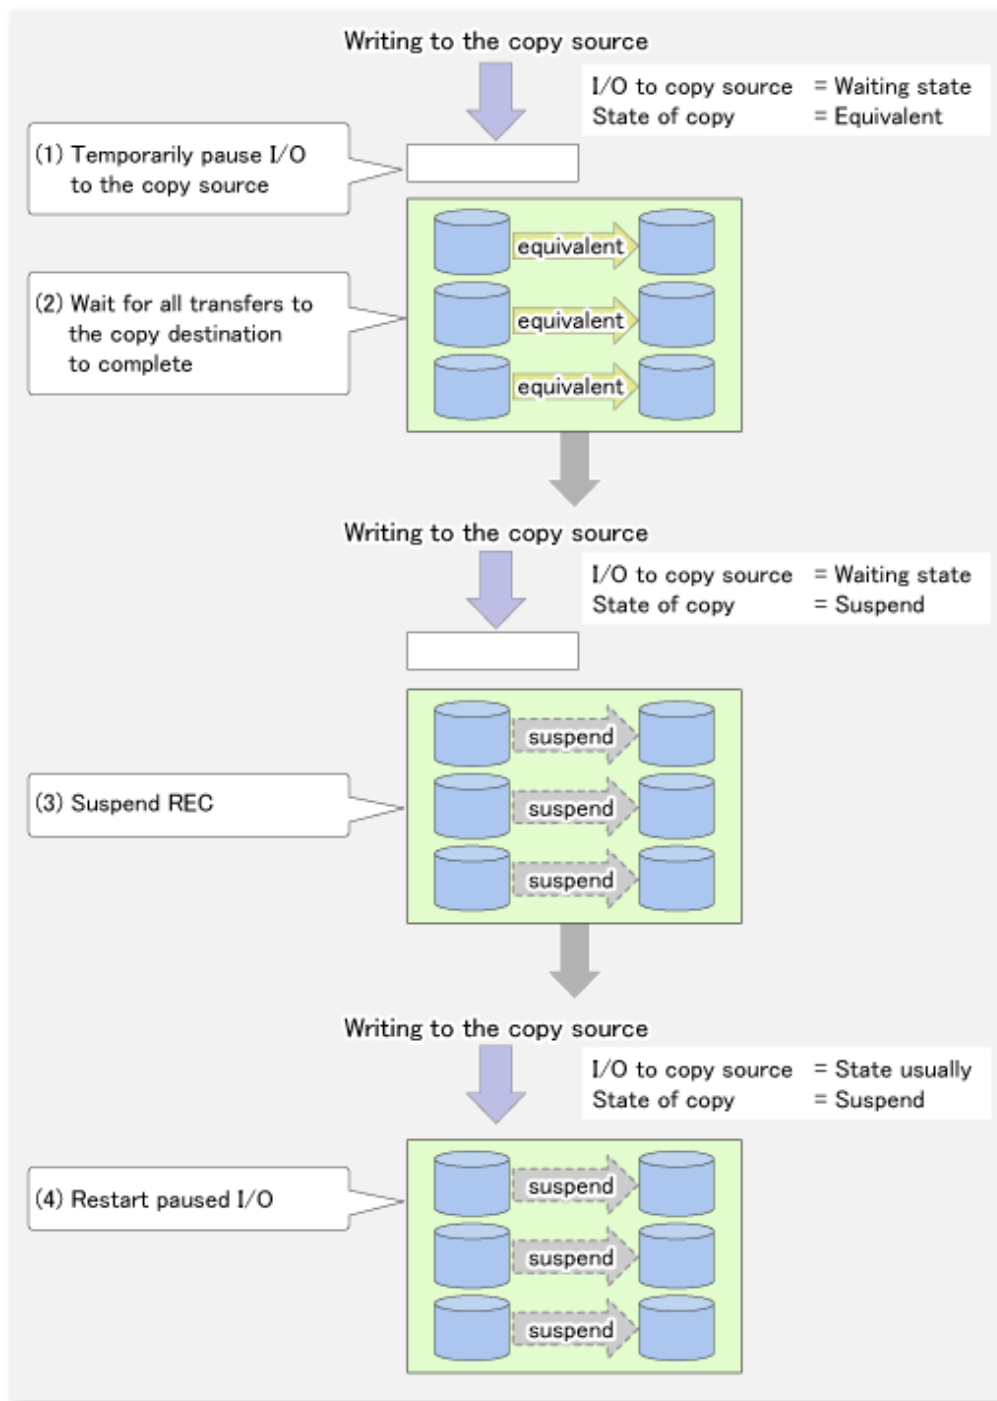
Figure 6.41 When Synchronization Is Restarted in Remain Mode



6.5.2.3 Concurrent Suspend Function


The Concurrent Suspend function simultaneously suspends multiple EC/REC sessions for ETERNUS Disk storage system. By using this function, a copy with consistency can be easily taken. For example, this is useful for a database composed of multiple volumes. The behavior that takes place within the ETERNUS Disk storage system is shown below.

Figure 6.42 Internal Operation of ETERNUS Disk Storage System



Replication using the Concurrent Suspend function is performed by specifying the `-Xconcur` option for the `swwrmake` command. Additionally, if Concurrent Suspend is performed using the Consistency transfer mode, it is no longer necessary to make temporary mode changes during multiple creation processes. (Refer to the figure below). Accordingly, the operating procedure used when Concurrent Suspend is executed in Consistency mode is the same as the procedure for Through mode and Synchronous mode.

Figure 6.43 When Consistency Mode Used



 Note

The maximum number of pairs that can simultaneously implement Concurrent Suspend depends on the specifications of the ETERNUS Disk storage system.

6.5.2.3.1 Changing Timeout Value for Concurrent Suspend Function

A timeout may occur before suspend is completed if there is a large quantity of data, if the communication link is poor, or if the REC Disk buffer is used for inter-box synchronous replication (REC function) in Consistency mode.

If the REC Disk buffer is used, refer to "[6.5.2.3.2 Concurrent Suspend When REC Disk buffer Is Used](#)".

The following message is shown if it times out.

```
swsrp2709 Concurrent suspend failed. Session ID=[-], Reason=[timeout]
```

If this message is shown, change the value of timeout by creating a file which sets the timeout value. Create a file which sets the timeout value for the Concurrent Suspend function under the following name. The file must be created on the same server that contains the source volume.

In Case of Normal Operation (Non-clustered Environment)

`/etc/opt/FJSVswsrp/data/DEFAULT/check.ini`

In Case of Operation in Clustered Environment

`/etc/opt/FJSVswsrp/<logical node>/data/DEFAULT/check.ini`

Table 6.5 File Format of check.ini

Section Name	Key Name	Value
[check]	ConcurCheckTimeout	Set the timeout value for the Concurrent Suspend function in units of seconds. Setting range: From 1 to 2,147,483,647 seconds. If you set a value beyond this range, the default value (60 seconds) is applied.

Example

Setting the timeout value for the Concurrent Suspend function to 120 seconds.

```
[check]
ConcurCheckTimeout=120
```

If the setting file for the Concurrent Suspend function is not created, the default value (60 seconds) is applied.

6.5.2.3.2 Concurrent Suspend When REC Disk buffer Is Used

If the REC Disk buffer is used for inter-box synchronous replication (REC function) in Consistency mode, check that no transferred data exists in the REC Disk buffer before executing the Concurrent Suspend.

If you execute the Concurrent Suspend when any data to be transferred exists in the REC Disk buffer, timeout may occur.

You can check the data volume of the REC Disk buffer in the **DiskRate** column when you execute the `swsrprebuffstat` command with specifying -L option.


6.5.2.4 Destination Access Permission Function

The Destination Access Permission function serves to inhibit user access to replication target volumes when copying is in progress and any access would therefore be inappropriate.

It starts synchronization (EC/REC) from the replication source volume to the replication target volume, and when full or differential copying is complete, they are maintained in a state of equivalency.


In the state of maintained equivalency, any upgrades on the replication source volume are successively reflected on the replication target volume, so that any access to the replication target volume would possibly be inappropriate.

The Destination Access Permission function can be used to inhibit access from users to the replication target in such unsuitable moments.



To set the Destination Access Permission function, execute the `swsrpstartsync` command with the `-Xda` option. Until replication processing is stopped by the `swsrpcancel` command neither WRITE nor READ access to the replication target volume can be performed. For changing this setting, once the replication operation has been stopped by the `swsrpcancel` command, execute the `swsrpstartsync` command again.

Whether access is inhibited can be checked by performing the `swsrpstat` command.



6.5.3 Backup and Restoration Using Replication Function

Data can be restored from the replica volume to the original volume if a volume pair has been defined with bi-directional copying by using the `swsrpsetvol` command.


Restoration can be executed according to the following procedures:

1. Execute the `swsrpcancel` command when the target volume for restoration has an EC session.
2. Execute the `swsrpmake` command. The specification of an original volume and the replica volume is reversed at the time that the replication is executed. For the backup operation that uses QuickOPC/SnapOPC/SnapOPC+, execute the `swsrpmake` command without specifying the `-T`, `-C` and `-P` option.

Note

It is not possible to perform restoration using replication if the ETERNUS Disk storage system does not support SnapOPC+ restoration. If restoration is required, copy the data from the copy destination volume manually.

Figure 6.44 Manual Restoration with SnapOPC+

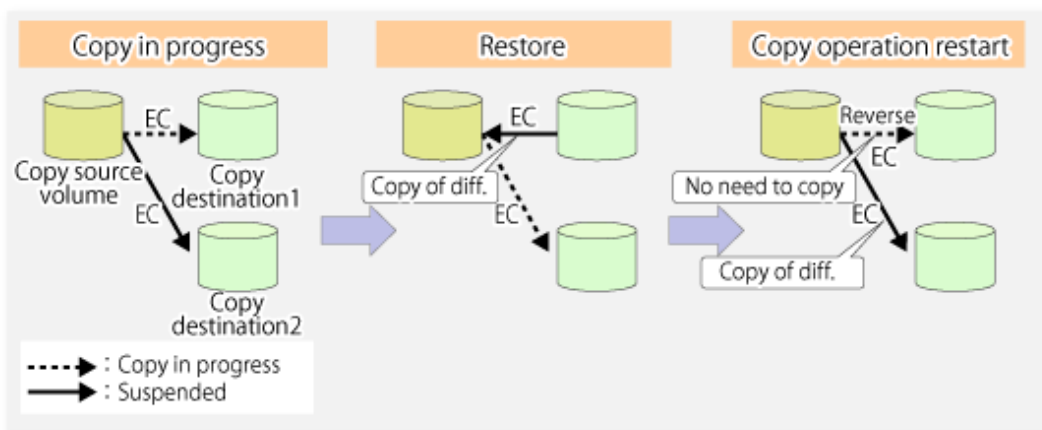


6.5.4 Execution of Synchronization Reversal

Using the `swsrprevsync` command, the physical copy amount can be greatly reduced and the AdvancedCopy time shortened for restore when there is a small amount of differential data.

Additionally, it is also possible to execute restore without cancelling other copy jobs in operations that match multiple copy jobs. Therefore, there is no need for overall copy after restore, so operations can be restarted with the load on ETERNUS Disk storage system kept to a minimum.

Figure 6.45 Example of EC Reversal Operation



6.6 Changing Configuration Information

When a Managed Server or device required for the replication operation has been changed, the information set in AdvancedCopy Manager must also be changed. This section explains how to change this information.

6.6.1 Changing 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 `swsrpdelvol` command and then execute the `swsrpsetvol` command again.

6.6.2 Deleting Source Volume or Replica Volume

To delete the source volume or the replica volume that have been configured, use the `swsrpdelvol` command.

6.6.3 Changing Device Information of Source or Replica Volume

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

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

1. Perform deletion process for the device change on the source/replica volume. For details of this, refer to "[6.6.2 Deleting Source Volume or Replica Volume](#)".
2. Change the device configuration.
3. Perform any additional steps that may be required in the processing of a device. For more information on this, refer to "[6.4.4 Fetching Device Information on Managed Server](#)".
4. Configure the source/replica volume with the new information.

6.7 Stopping Operation

To stop replication processing that is in progress, or to change synchronized replication into snapshot replication, use the `swsrancel` command.



To stop replication processing by SnapOPC+, start with the oldest snap generation.

To forcibly stop the second or subsequent snap generations, use the `-Xforce` option with the `swsrancel` command. This stops the specified snap generation and all earlier snap generations.

If the `-Xforce` and `-Xgroup` options are specified at the same time, the SnapOPC+ sessions in the group are stopped, and all the snap generations and snap generations before it in the group are also stopped.

To stop replication operation, stop the daemon on the Managed Server. In general, the daemon stops when the system stops.

The daemon can be stopped independently if it needs to be stopped for some reason. For more information, refer to "[Chapter 2 Starting and Stopping Daemons](#)".



When the daemon/service on the Managed Server has stopped, all functions of AdvancedCopy Manager running on the Managed Server stop.

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

6.8 SDX Object Replication

This section describes the replication operation of SDX objects.

A good understanding of the basic operation of normal volumes is required for this operation.

For details on PRIMECLUSTER GD, refer to the *PRIMECLUSTER Global Disk Services User's Guide (Solaris(TM) Operating Environment)*.

6.8.1 Units of SDX Object Operations

There are two modes of operation according to the difference in copy units:

1. Replication operation in units of logical volumes (method with PRIMECLUSTER GD Snapshot)
2. Replication operation in units of slices (method without PRIMECLUSTER GD Snapshot)

For details on how to confirm and select the type of operation, refer to the section titled "[3.8.1.1 Configuring Unit of SDX Object Operations](#)". Replication in units of logical volumes (method with PRIMECLUSTER GD Snapshot) is applicable only to replication on single server.

For server-to-server replication, you must use units of slices.

Note

- The two types of operation cannot be performed on a single Managed Server.
- In a replication operation (method that uses PRIMECLUSTER GD Snapshot) of a logical volume unit, the Storage Cluster Continuous Copy function is not available. When using the Storage Cluster Continuous Copy function, use the replication operation (method that does not use PRIMECLUSTER GD Snapshot) of a slice unit.

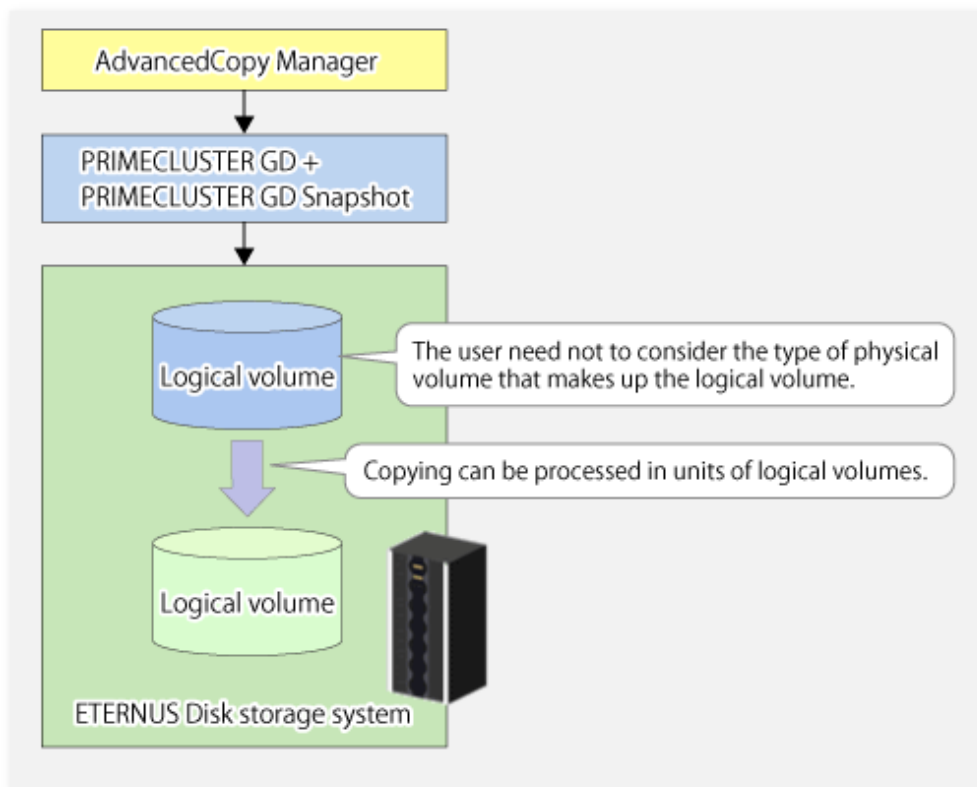
6.8.2 Replication Operation in Units of Logical Volumes (Method with PRIMECLUSTER GD Snapshot)

SDX objects of PRIMECLUSTER GD can be replicated in units of logical volumes.

Replication is performed in linkage with PRIMECLUSTER GD Snapshot.

You can specify a PRIMECLUSTER GD logical volume name as the target volume. Therefore, you need not consider the physical disk structure of the SDX object. For operation in units of slices, all physical slices that make up the mirror volume must be specified, and the physical structure must be considered in the configuration and operation of replication.

Figure 6.46 Backup Operation in Units of Logical Volumes



To specify a PRIMECLUSTER GD logical volume in an AdvancedCopy Manager command, specify the name of the PRIMECLUSTER GD logical volume.

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



Configuration and Conditions of Unsupported SDX Objects

The supported volume types are single volumes and mirror volumes.

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

- Volume of a shadow class
- Stripe volume
- Switch volume
- Volume in a concatenation group

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

If the total number of the slices that make up the original volume and slices that make up all replica volumes associated with the original volume at the start of synchronous replication or snapshot-type replication is 33 or greater, copying cannot be processed. Also, if the original volume is in the invalid state, copying cannot be processed.



To confirm whether the type of volume being used is either single, mirror, stripe, concatenation, or switch, execute the command "sdxinfo -e long" and check the volume type attribute displayed in the TYPE column for the matching volume label in the OBJ column.

[Volume object type attributes]

- single : Single volume (supported)
- mirror : Mirror volume (supported)
- stripe : Stripe volume (not supported)
- concat : Volume in a concatenation group (not supported)
- switch : Switch volume (not supported)

6.8.2.1 Configuring Replication Operation

Note the following about performing the replication operation.

For replication with the function in linkage with PRIMECLUSTER GD Snapshot, AdvancedCopy Manager processes binding and separation of volumes by using the PRIMECLUSTER GD Snapshot function.

Therefore, the original and replica volumes must be SDX objects that can be used as a master object and proxy object respectively for a PRIMECLUSTER GD Snapshot.

The copy source volume is the master object, and the copy destination volume is the proxy object.

Specify the PRIMECLUSTER GD volume structure before registering a volume with AdvancedCopy Manager.

Note the compatibility requirements described below for replication in units of physical slices.

For detailed conditions on what constitutes a master/proxy volume, refer to "Preconditions for Proxy Constitution" in "Cautions" in the *PRIMECLUSTER Global Disk Services User's Guide (Solaris(TM) Operating Environment)*.

- The original volume and replica volume must be the same type of SDX objects. Replication volume information based on logical volumes in a combination that includes slices or volumes of another OS cannot be set. When logical volumes are used, the original volume and replica volume must be the same size.

If the volumes are separated or there is no copy session, AdvancedCopy Manager assumes a non-replicated state. The following conditions result in this state:

- OPC is started by the `swsrpmake` command, and copy processing is completed.

- The `swsrpstartsync` command is executed with the "software-copy" operand specified, and synchronous processing is suspended by the `swsrpmake` command.

If a pair of separated replication volumes is specified, the `swsrpstartsync`, `swsrpmake` or `swsrpdelvol` command automatically clears the separated state.

To clear the separated state without a copy session for operational reasons, delete the replication volumes by using the `swsrpdelvol` command or by using the "sdxproxy" command, referring to the *PRIMECLUSTER Global Disk Services User's Guide (Solaris(TM) Operating System)*.

6.8.2.2 Replication Volumes

To use the Advanced Copy function, the structures of the original volume and replica volume must satisfy the following conditions.

Structures of Copy Source Volumes

To use the Advanced Copy function, the structure of the copy source volume must be one of the following:

- Single volume
- Not the root class
- Mirror volume belonging to a mirror group that has no subordinate groups

For a mirror volume in a mirror group connected to a subordinate group, only software-copy can be used.

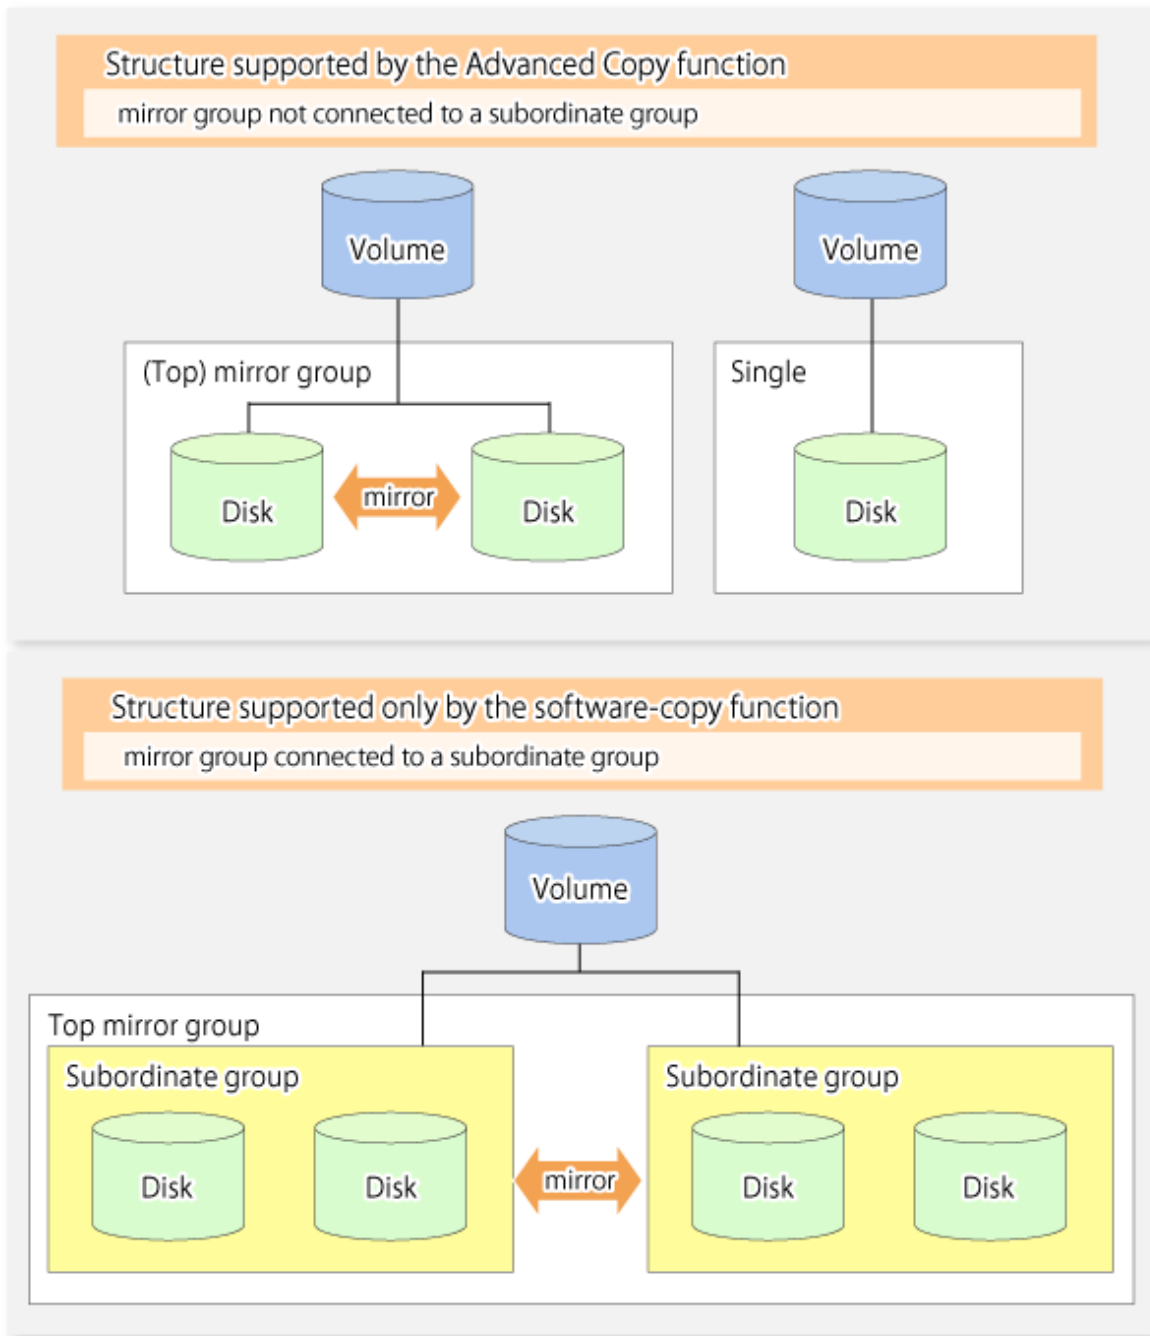
AdvancedCopy Manager does not support the following SDX objects (volumes):

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

If the total number of the slices that make up the original volume plus the slices that make up all replication volumes associated with the original volume at the start of synchronous processing or snapshot-type replication processing is 33 or greater, copying cannot be processed.

Also, if the copy source volume is in the invalid state, copying cannot be processed.

Figure 6.47 Copy Source Volume Structures Supported by Advanced Copy Function




Structures of Copy Destination Volumes

To use the Advanced Copy function, the structure of the copy destination volume must be one of the following:

- Single structure
 This means that the type attribute of the volume is "single" or the type attribute is "mirror" and one disk makes up the disk group.
- Not the root class
- Not a Stripe or Concatenation type

If the copy destination volume is a mirror volume, only the software-copy function can be used. To use the Advanced Copy function in both directions, both volumes must be in a single structure. Neither a stripe nor concatenation volume can be used as the copy destination volume.

Figure 6.48 Copy Destination Volume Structures Supported by Advanced Copy Function



6.8.2.3 Copy Functions Available for Each PRIMECLUSTER GD Structure

The two types of PRIMECLUSTER GD Snapshot copy functions are described below. The availability of the functions depends on the SDX object structure.

- Advanced Copy : Copying with the hardware function of ETERNUS Disk storage system
- Software-copy : Copying with the disk driver of PRIMECLUSTER GD

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

Table 6.6 SDX Object Structures and Available Copy Functions

Replication Structure	Unit	Copy Source/ Destination	SDX Object Structure (*1)	Available Copy Function
Within a server	Logical volume	Copy source	Single	Advanced Copy/ Software-copy

Replication Structure	Unit	Copy Source/ Destination	SDX Object Structure (*1)		Available Copy Function
			Mirror	No subordinate group exists.	Advanced Copy/ Software-copy
				Subordinate groups exist.	Software-copy
		Copy destination	Single		Advanced Copy/ Software-copy
			Mirror	This consists of one disk.	Advanced Copy/ Software-copy
				This consists of multiple disks.	Software-copy

*1: Stripe/concatenation/switch-types of objects cannot be copied.

6.8.2.4 Fetching Device Information from Managed Server

To perform the replication operation with a Managed Server and obtain information relating to a subordinate device of the Managed Server, you must register the server.

For details on how to fetch the device information, refer to "[6.4.4 Fetching Device Information on Managed Server](#)".

6.8.2.5 Configuring Replication Volumes

Note the following about configuring the original and replica volumes.

- Confirm that the volume is of a structure as described in "[6.8.3.1.1 Supported Logical Volume Structures](#)".
- Confirm that the original/replica volume is not a master volume or proxy volume.

Register PRIMECLUSTER GD logical volumes as the original and replica volumes.

```
# /opt/FJSVswsrp/bin/swsrpsetvol -h SRC-SV /dev/sfdsk/CLS1/dsk/VOL1 /dev/sfdsk/CLS1/dsk/VOL2
swsrpsetvol completed
#
```

6.8.2.6 Executing Replication

Note the following about performing replication of SDX objects.

To start or resume synchronous processing ([swsrpstartsync](#) command) or to start snapshot processing ([swsrpmake](#) command), the logical volume of the copy destination must be inactive.

If the logical volume of the copy destination is active when the command is executed, processing to stop the volume is performed within the command. If the copy destination volume is a file system volume, the file system is unmounted during replication pre-processing and the volume is then stopped.

If the volume cannot be stopped (because the volume is accessed by an application), the command is terminated abnormally.

Example of Synchronous Replication

```
# /opt/FJSVswsrp/bin/swsrpstartsync -h SRC-SV /dev/sfdsk/CLS1/dsk/VOL1 /dev/sfdsk/CLS1/dsk/VOL2
FROM=/dev/sfdsk/CLS1/dsk/VOL1, TO=/dev/sfdsk/CLS1/dsk/VOL2 swsrpstartsync completed
Equivalency holding state
# swsrpmake -h SRC-SV /dev/sfdsk/CLS1/dsk/VOL1 /dev/sfdsk/CLS1/dsk/VOL2
FROM=/dev/sfdsk/CLS1/dsk/VOL1, TO=/dev/sfdsk/CLS1/dsk/VOL2 swsrpmake completed
```

Example of Snapshot Replication

```
# /opt/FJSVswsrp/bin/swsrpmake -h SRC-SV /dev/sfdsk/CLS1/dsk/VOL1 /dev/sfdsk/CLS1/dsk/VOL2
FROM=/dev/sfdsk/CLS1/dsk/VOL1, TO=/dev/sfdsk/CLS1/dsk/VOL2 swsrpmake completed
#
```

6.8.2.7 Operational Notes

Do not use the "sdxproxy" command directly to change the state of a volume registered for the replication function. Otherwise, the "sdxproxy" command may create an inconsistency between the management information of the replication function and the actual volume state. In this case, use the [swsrcancel](#) command to cancel the replication operation.


6.8.3 Replication Operation in Units of Slices (Method without PRIMECLUSTER GD Snapshot)

SDX objects of PRIMECLUSTER GD can be replicated in units of slices where replication is performed without PRIMECLUSTER GD Snapshot.

This operation however requires that the physical disks that make up the PRIMECLUSTER GD logical volume be specified.

For details on PRIMECLUSTER GD, refer to the *PRIMECLUSTER Global Disk Services User's Guide (Solaris(TM) Operating System)*.

Figure 6.49 Replication Operation in Units of Slices



Note

- If mirroring between boxes is enabled and copying to the original volume by OPC is required in the event of a box failure, both mirrored systems must be replicated. In such cases, the required capacity of the replica volume is not the size of the logical volume but the size of the physical volume.
- The SDX objects that can be used for operations in slice units are as follows:
 - Single volumes
 - Non-layered mirror volumes
- You can determine if the volume being used is the single, mirror, stripe, concatenation, or switch volume type by executing "sdxinfo -e long". (Check the value in the TYPE column in the line displayed for the volume in the OBJ column.)
[Volume object type attribute]

- single: Single volume (operation possible)
- mirror: Mirror volume (operation possible)
- stripe: Stripe volume (operation not possible)
- concat: Volume in a concatenation group (operation not possible)
- switch: Switch volume (operation not possible)

When specifying a name in the AdvancedCopy Manager command, use the following name format that combines a logical volume name and an AdvancedCopy Manager device name. Refer to "1.4.2 SDX Objects of PRIMECLUSTER GD/SynfinityDISK" for details.

```
/dev/sfdsk/<class-name>/dsk/<volume-name>:<device name where volume is constructed>
```

6.8.3.1 Configuring Replication Operation

For notes on configuring the replication operation, refer to "10.1.5.2 Notes on SDX Object Operations (Replication Management)".

6.8.3.1.1 Supported Logical Volume Structures

Replication with the copy destination volume of a mirror slice is not supported.

Consequently, if either the original volume or replica volume is a mirror slice, you can perform replication only from the mirror slice to a physical slice.

To copy data to a mirror slice, use the "dd" command. Copying must be executed for all mirror slices that make up the logical volume.

6.8.3.2 Fetching Device Information from Managed Server

To perform the replication operation with a Managed Server and obtain information about a subordinate device of the Managed Server, you must register the server.

For details on how to fetch the device information, refer to "6.4.4 Fetching Device Information on Managed Server".

6.8.3.3 Configuring Replication Volumes

Register the mirror slices that make up a mirror volume of PRIMECLUSTER GD as replication volumes.

```
# /opt/FJSVswsrp/bin/swsrpsetvol /dev/sfdsk/CLS1/dsk/VOL1:c1t0d1 /dev/sfdsk/CLS1/dsk/
VOL2:c1t0d2
swsrpsetvol completed
#
```

6.8.3.4 Executing Replication

Example of Snapshot Replication

```
# /opt/FJSVswsrp/bin/swsrpmake /dev/sfdsk/CLS1/dsk/VOL1:c1t0d1 /dev/sfdsk/CLS1/dsk/VOL2:c1t0d2
FROM=/dev/sfdsk/CLS1/dsk/VOL1:c1t0d1, TO=/dev/sfdsk/CLS1/dsk/VOL2:c1t0d2 swsrpmake completed
#
```

States of Objects in Replication Creation

Replication can be performed when the SDX object that makes up the relevant logical volume is in one of the states listed below. If the object is in any other state, replication cannot be performed. (AdvancedCopy Manager checks the state of the SDX object by using the sdxinfo command of PRIMECLUSTER GD/SafeDISK.)

1. The state of the mirror volume is ACTIVE or STOP.
2. The state of the physical disk is ENABLE.
3. The state of the mirror slice is ACTIVE or TEMP.

Replication Pre-processing and Post-processing

Pre-processing and post-processing scripts are not executed when a slice of the SDX object is in the TEMP state.

6.9 VxVM Volume Operation

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

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

- Replication operation in units of logical volume
- Replication operation in units of physical disks

6.9.1 Replication Operation in Units of Logical Volume

Replication can only be performed when the VxVM volume configuration satisfies the conditions required for operation involving units of logical volumes.



Note

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

6.9.1.1 Configuration of Operations Involving VERITAS Volume Structures

6.9.1.1.1 Supported VxVM Volume Structures

For the supported VxVM volume structures, refer to "[1.4 Managing Device on AdvancedCopy Manager](#)".

6.9.1.1.2 Preparing Replica Volume

The following two types of volumes can be used as a replica volume of a VxVM logical volume:


- Logical volume of VxVM
- Slice

For this replica volume, however, use a VxVM logical volume in the same structure as that of the original volume where the replica volume is mounted or in use by another application.

The same volume structure means the following conditions are satisfied:

- The size of the physical disk (LU) is the same.
- The size of the private region is the same.
- The size of the public region is the same.
- Each sub-disk in the public region has a corresponding start offset and size.
- The logical volume and sub-disks have a relationship.

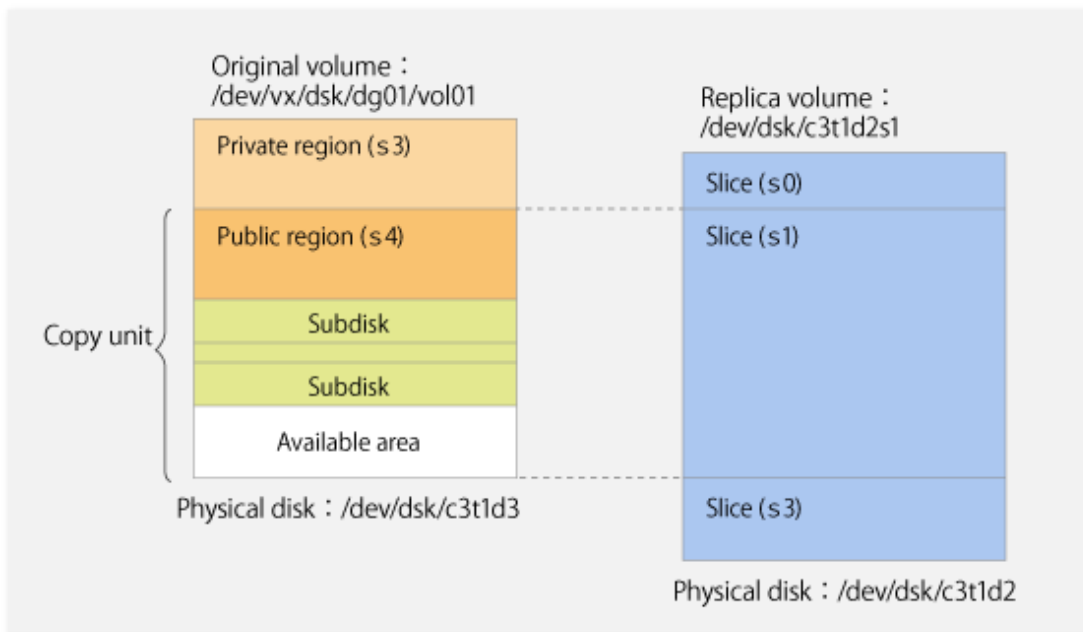
Figure 6.50 Example of Correspondence Between Transaction Volume and Backup Volume (1)



 Note

.....
 If the replica volume is the same size as the public region, the replica volume of a VxVM volume can be made into a slice. However, the replica volume generally cannot be mounted or used by other applications.

Figure 6.51 Example of Correspondence Between Transaction Volume and Backup Volume (2)



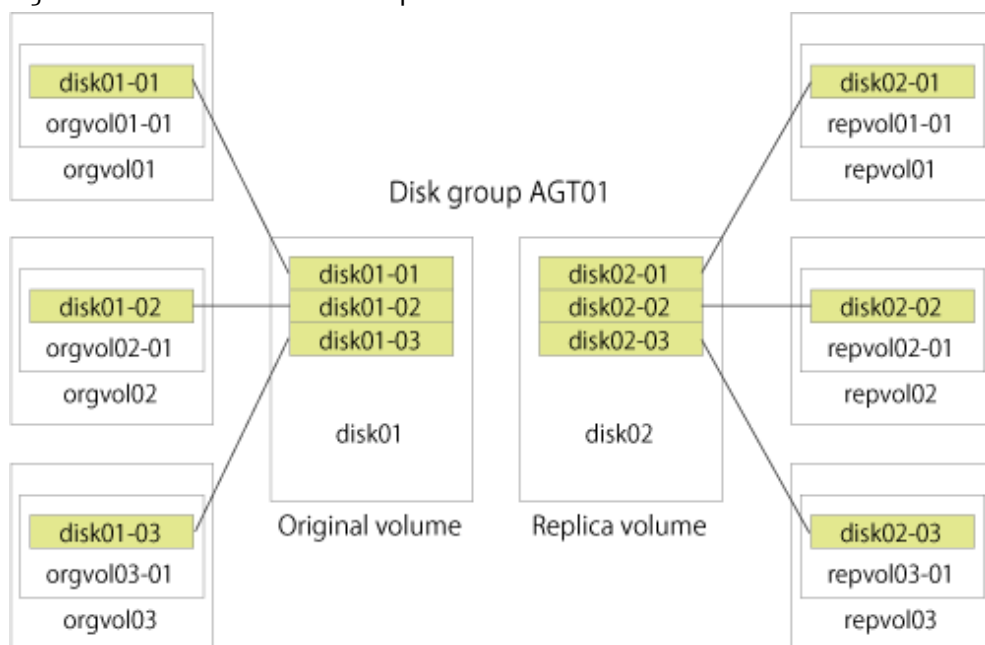
Similarly, a replica volume of a slice type can be made into a VxVM volume. However, the replica volume cannot be mounted or used by other applications.

6.9.1.2 Preparations

6.9.1.2.1 Setting Original Volume and Replica Volume

If the target VxVM logical volume satisfies the condition of logical volumes: VM disk = N:1 (N>1), use one of the logical volume names included with each VM disk to register an original or replica volume, and set the replication volume information. As an example, assume the original volume and replica volume in the following volume structures:

Figure 6.52 Volume Structure Example



- VM disk disk01 (original volume) has logical volumes named orgvol01, orgvol02, and orgvol03, and VM disk disk02 (replica volume) has logical volumes named repvol01, repvol02, and repvol03.
- VM disks disk01 and disk02 have the same slice size. The sub-disk structures of VM disks disk01 and disk02 are the same. The sub-disk structure can be checked by executing "vxprint -AGts".

In the above example, use the logical volume names of orgvol01 and repvol01 when setting the replication volume information.

```
# /opt/FJSVswsrp/bin/swsrpsetvol /dev/vx/dsk/AGT01/orgvol01 /dev/vx/dsk/AGT01/repvol01
swsrpsetvol completed
#
```

The logical volume name specified in the replication volume information setting command is used to specify the VM disk (slice) for replication processing.

This means that orgvol02 or orgvol03 instead of orgvol01 can be set in the above example. This is also true for repvol01.

However, if replication volume information for the combination of orgvol01 and repvol01 is already set, replication volume information using another combination of volumes (e.g., orgvol02 and repvol02) cannot be set.

6.9.1.3 Executing Replication

If the target VxVM logical volume satisfies the condition of logical volumes: VM disk = N:1 (N>2), unmount/mount processing cannot be performed using the pre-processing or post-processing scripts.

For this type of structure, unmount the file system before executing replication, and mount the file system after replication creation.

If the target VxVM logical volume satisfies the condition of logical volumes: VM disk = 1:1, the logical volume is unmounted and mounted in the same way as in operation with normal slices and the following procedure is not required.

Procedure for Synchronous Replication

```
# umount /global/AGT01/repvol01 <-- unmount repvol01
# umount /global/AGT01/repvol02 <-- unmount repvol02
```

```

# umount /global/AGT01/repvol03 <-- unmount repvol03
# /opt/FJVSwsrp/bin/swsrpstartsync /dev/vx/dsk/AGT01/orgvol01 /dev/vx/dsk/AGT01/repvol01
FROM=/dev/vx/dsk/AGT01/orgvol01@svr1,TO=/dev/vx/dsk/AGT01/repvol01@svr1 swsrpstartsync completed

:
(After the equivalence maintenance status)
# umount /global/AGT01/orgvol01 <-- unmount orgvol01
# umount /global/AGT01/orgvol02 <-- unmount orgvol02
# umount /global/AGT01/orgvol03 <-- unmount orgvol03
# /opt/FJVSwsrp/bin/swsrpmake /dev/vx/dsk/AGT01/orgvol01 /dev/vx/dsk/AGT01/repvol01
FROM=/dev/vx/dsk/AGT01/orgvol01@svr1,TO=/dev/vx/dsk/AGT01/repvol01@svr1 swsrpmake completed
# mount -F ufs /dev/vx/dsk/AGT01/orgvol01 /global/AGT01/orgvol01 <-- mount orgvol01
# mount -F ufs /dev/vx/dsk/AGT01/orgvol02 /global/AGT01/orgvol02 <-- mount orgvol02
# mount -F ufs /dev/vx/dsk/AGT01/orgvol03 /global/AGT01/orgvol03 <-- mount orgvol03
# mount -F ufs /dev/vx/dsk/AGT01/repvol01 /global/AGT01/repvol01 <-- mount repvol01
# mount -F ufs /dev/vx/dsk/AGT01/repvol02 /global/AGT01/repvol02 <-- mount repvol02
# mount -F ufs /dev/vx/dsk/AGT01/repvol03 /global/AGT01/repvol03 <-- mount repvol03

```

Procedure for Snapshot Replication

```

# umount /global/AGT01/repvol01 <-- unmount repvol01
# umount /global/AGT01/repvol02 <-- unmount repvol02
# umount /global/AGT01/repvol03 <-- unmount repvol03
# umount /global/AGT01/orgvol01 <-- unmount orgvol01
# umount /global/AGT01/orgvol02 <-- unmount orgvol02
# umount /global/AGT01/orgvol03 <-- unmount orgvol03
# /opt/FJVSwsrp/bin/swsrpmake /dev/vx/dsk/AGT01/orgvol01 /dev/vx/dsk/AGT01/repvol01
FROM=/dev/vx/dsk/AGT01/orgvol01@svr1,TO=/dev/vx/dsk/AGT01/repvol01@svr1 swsrpmake completed
# mount -F ufs /dev/vx/dsk/AGT01/orgvol01 /global/AGT01/orgvol01 <-- mount orgvol01
# mount -F ufs /dev/vx/dsk/AGT01/orgvol02 /global/AGT01/orgvol02 <-- mount orgvol02
# mount -F ufs /dev/vx/dsk/AGT01/orgvol03 /global/AGT01/orgvol03 <-- mount orgvol03
# mount -F ufs /dev/vx/dsk/AGT01/repvol01 /global/AGT01/repvol01 <-- mount repvol01
# mount -F ufs /dev/vx/dsk/AGT01/repvol02 /global/AGT01/repvol02 <-- mount repvol02
# mount -F ufs /dev/vx/dsk/AGT01/repvol03 /global/AGT01/repvol03 <-- mount repvol03

```

To prevent unmount/mount processing of the original volume, back up the volume by using the procedures described below.

Procedure for Synchronous Replication (While Preventing Unmount/Mount Processing of Original Volume)

```

# umount /global/AGT01/repvol01 <-- unmount repvol01
# umount /global/AGT01/repvol02 <-- unmount repvol02
# umount /global/AGT01/repvol03 <-- unmount repvol03
# /opt/FJVSwsrp/bin/swsrpstartsync /dev/vx/dsk/AGT01/orgvol01 /dev/vx/dsk/AGT01/repvol01
FROM=/dev/vx/dsk/AGT01/orgvol01@svr1,TO=/dev/vx/dsk/AGT01/repvol01@svr1 swsrpstartsync
completed

:
(After the equivalence maintenance status)
# lockfs -w /global/AGT01/orgvol01 <-- lock file system orgvol01
# lockfs -w /global/AGT01/orgvol02 <-- lock file system orgvol02
# lockfs -w /global/AGT01/orgvol03 <-- lock file system orgvol03
# /opt/FJVSwsrp/bin/swsrpmake /dev/vx/dsk/AGT01/orgvol01 /dev/vx/dsk/AGT01/repvol01
FROM=/dev/vx/dsk/AGT01/orgvol01@svr1,TO=/dev/vx/dsk/AGT01/repvol01@svr1 swsrpmake completed
# lockfs -u /global/AGT01/orgvol01 <-- unlock file system orgvol01
# lockfs -u /global/AGT01/orgvol02 <-- unlock file system orgvol02
# lockfs -u /global/AGT01/orgvol03 <-- unlock file system orgvol03
# fsck -n -F ufs /dev/vx/dsk/AGT01/repvol01 <-- check file system backvol01
# fsck -n -F ufs /dev/vx/dsk/AGT01/repvol02 <-- check file system backvol02
# fsck -n -F ufs /dev/vx/dsk/AGT01/repvol03 <-- check file system backvol03
# mount -F ufs /dev/vx/dsk/AGT01/repvol01 /global/AGT01/repvol01 <-- mount repvol01
# mount -F ufs /dev/vx/dsk/AGT01/repvol02 /global/AGT01/repvol02 <-- mount repvol02
# mount -F ufs /dev/vx/dsk/AGT01/repvol03 /global/AGT01/repvol03 <-- mount repvol03

```

Procedure for Snapshot Replication (While Preventing Unmount/Mount Processing of Original Volume)

```
# umount /global/AGT01/repvol01 <-- umount repvol01
# umount /global/AGT01/repvol02 <-- umount repvol02
# umount /global/AGT01/repvol03 <-- umount repvol03
# lockfs -w /global/AGT01/orgvol01 <-- lock file system orgvol01
# lockfs -w /global/AGT01/orgvol02 <-- lock file system orgvol02
# lockfs -w /global/AGT01/orgvol03 <-- lock file system orgvol03
# /opt/FJSVswsrp/bin/swsrpmake /dev/vx/dsk/AGT01/orgvol01 /dev/vx/dsk/AGT01/repvol01
FROM=/dev/vx/dsk/AGT01/orgvol01@svr1,TO=/dev/vx/dsk/AGT01/repvol01@svr1 swsrpmake completed
# lockfs -u /global/AGT01/orgvol01 <-- unlock file system orgvol01
# lockfs -u /global/AGT01/orgvol02 <-- unlock file system orgvol02
# lockfs -u /global/AGT01/orgvol03 <-- unlock file system orgvol03
# fsck -n -F ufs /dev/vx/dsk/AGT01/repvol01 <-- check file system backvol01
# fsck -n -F ufs /dev/vx/dsk/AGT01/repvol02 <-- check file system backvol02
# fsck -n -F ufs /dev/vx/dsk/AGT01/repvol03 <-- check file system backvol03
# mount -F ufs /dev/vx/dsk/AGT01/repvol01 /global/AGT01/repvol01 <-- mount repvol01
# mount -F ufs /dev/vx/dsk/AGT01/repvol02 /global/AGT01/repvol02 <-- mount repvol02
# mount -F ufs /dev/vx/dsk/AGT01/repvol03 /global/AGT01/repvol03 <-- mount repvol03
```

Note

Even if "lockfs -w" is used to flush the file system cache and prevent subsequent update, file contents are not guaranteed for the file in Write Open (excluding the database space of Oracle used to support on-line backup).

If metadata renewal is in progress, consistency of data of the file system itself is cannot be guaranteed because "lockfs" is specified. To handle this problem, perform "fsck" processing after replication to check the consistency of the data.

Even if the following message is output in "fsck" processing, there is no problem with the consistency of the file system. Execute "fsck -y" to restore.

```
FILE SYSTEM STATE IN SUPERBLOCK IS WRONG; FIX?
```

6.9.1.4 Changing Configuration Information

To change the volume structure during the replication operation, follow the procedure below.

1. Delete the replication volume. For details on how to delete a replication volume, refer to "[6.6.2 Deleting Source Volume or Replica Volume](#)".
2. Change the structure of the logical volume.
3. With Web Console, execute the discovering all devices and reflecting information (Discover/Reload).
4. Re-configure the replication volume. For details on how to register a replication volume, refer to "[6.4.5 Configuring Source Volume and Replica Volume](#)".
5. Resume the replication operation.

6.9.2 Replication Operation in Units of Physical Disks

When the VxVM volume configuration does not satisfy the conditions for replication operation using units in logical volumes, replication operation can be performed using units of physical disks.

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

Point

For the physical volume which should be synchronized, confirm it by using either Web Console or the `stgxfwcmdispdev` command.

Note

- It is only VM disk of the "nopriv" type that a physical slice becomes the unit of management.
- Before starting this operation, be sure to understand the basic operation of ordinary volumes.
- Snapshot backup is recommended for backing up in units of physical disks. In the case of synchronized backup, commands that access destination volumes, such as VxVM commands, cannot be executed during full copy or differential copy.
- For operation in a cluster configuration, the device name (/dev/(r)dsk/c#t#d#) of the physical disks that comprise the disk group must be the same at all servers that comprise the cluster, and the ETERNUS Disk storage system's disk indicated by the device name must be the same.

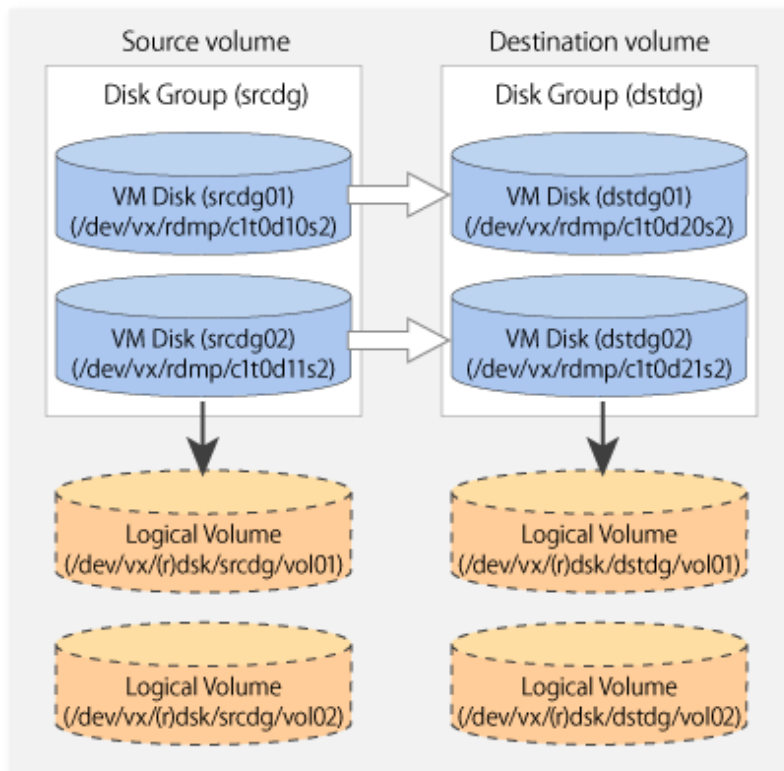
For operation in a SunCluster environment, the VxVM enclosure name must be the same at all servers that comprise the cluster, and the ETERNUS Disk storage system's disk indicated by the enclosure name must be the same.

6.9.2.1 Configuration

When configuring the disk groups of the original and replica volumes, you should note the following:

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

Figure 6.53 Copy Source Volume and Copy Destination Volume



6.9.2.2 Preparations

6.9.2.2.1 Confirming Disk Group Configuration Information File

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

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

6.9.2.2.2 Configuring Original and Replica Volumes

When setting the original and replica volumes, specify all disks in the volume group.

Example of Setting Original and Replica Volumes

```
# /opt/FJSVswsrp/bin/swsrpsetvol /dev/vx/dmp/clt0d10s2 /dev/vx/dmp/clt0d20s2
swsrpsetvol completed
# /opt/FJSVswsrp/bin/swsrpsetvol /dev/vx/dmp/clt0d11s2 /dev/vx/dmp/clt0d21s2
swsrpsetvol completed
#
```

6.9.2.3 Replication

Perform Replication by synchronizing all physical disks in the disk group.

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

Example of Snapshot Backup

```
(Perform pre-processing for the source and target volumes.)

# /opt/FJSVswsrp/bin/swsrpmake -f -t /dev/vx/dmp/clt0d10s2 /dev/vx/dmp/clt0d20s2
FROM=/dev/vx/dmp/clt0d10s2@SV1, TO=/dev/vx/dmp/clt0d20s2@SV1 swsrpmake completed
# /opt/FJSVswsrp/bin/swsrpmake -f -t /dev/vx/dmp/clt0d11s2 /dev/vx/dmp/clt0d21s2
FROM=/dev/vx/dmp/clt0d11s2@SV1, TO=/dev/vx/dmp/clt0d21s2@SV1 swsrpmake completed
#

(Perform post-processing for the source and target volumes.)
```

Example of Synchronous Replication

```
(Perform pre-processing for the target volume.)

# /opt/FJSVswsrp/bin/swsrpstartsync -t /dev/vx/dmp/clt0d10s2 /dev/vx/dmp/clt0d20s2
FROM=/dev/vx/dmp/clt0d10s2@SV1, TO=/dev/vx/dmp/clt0d20s2@SV1 swsrpstartsync completed
# /opt/FJSVswsrp/bin/swsrpstartsync -t /dev/vx/dmp/clt0d11s2 /dev/vx/dmp/clt0d21s2
FROM=/dev/vx/dmp/clt0d11s2@SV1, TO=/dev/vx/dmp/clt0d21s2@SV1 swsrpstartsync completed
#

(After state of equivalency upkeep)
(Perform pre-processing for the source volume.)

# /opt/FJSVswsrp/bin/swsrpmake -f -t /dev/vx/dmp/clt0d10s2 /dev/vx/dmp/clt0d20s2
FROM=/dev/vx/dmp/clt0d10s2@SV1, TO=/dev/vx/dmp/clt0d20s2@SV1 swsrpmake completed
# /opt/FJSVswsrp/bin/swsrpmake -f -t /dev/vx/dmp/clt0d11s2 /dev/vx/dmp/clt0d21s2
FROM=/dev/vx/dmp/clt0d11s2@SV1, TO=/dev/vx/dmp/clt0d21s2@SV1 swsrpmake completed
#

(Perform post-processing for the source and target volumes.)
```

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

Pre-processing and Post-processing for Restoration

Table 6.7 In Case of Normal Operation (Non-clustered Environment)

Volume Type	Pre-processing	Post-processing
Source volume	<ol style="list-style-type: none"> Secure data integrity by stopping access to all logical volumes in the volume group. If file systems are included, unmount all file systems in the volume group. 	If file systems are included, remount the volumes that were unmounted by pre-processing.

Volume Type	Pre-processing	Post-processing
	3. Import the disk group, when the disk group is not imported for replication on single server.	
Target volume	<ol style="list-style-type: none"> 1. Stop access to all logical volumes in the volume group. 2. If file systems are included, unmount all file systems in the volume group. 3. Deport the disk group. 4. Set the disk groups subordinate physical disks to offline. 	<ol style="list-style-type: none"> 1. Online the physical disks that were set to offline during pre-processing. 2. Reconfigure the volume group. 3. If file systems are included, remount the volumes that were unmounted by pre-processing.

Table 6.8 In Case of Operation in Clustered Environment

Volume Type	Pre-processing	Post-processing
Source volume	<ol style="list-style-type: none"> 1. Secure data integrity by stopping access to all logical volumes in the volume group. 2. If file systems are included, unmount all file systems in the volume group. 3. Implement the following tasks depending on the disk group registration status in the clustered system disk group resource. <ul style="list-style-type: none"> - When the Disk Group is not Registered: If the disk group is not imported for replication on single server, import the disk group. - When the Disk Group is Registered: If the disk group resource is not online for replication on single server, bring the disk group resource online. 	If file systems are included, remount the volumes that were unmounted by pre-processing.
Target volume	<ul style="list-style-type: none"> - When not using the VERITAS Cluster Server, or when using a version earlier than the VERITAS Cluster Server 6.0 <ol style="list-style-type: none"> 1. Stop access to all logical volumes in the volume group. 2. If file systems are included, unmount all file systems in the volume group. 3. Implement the following tasks depending on the disk group registration status in the clustered system disk group resource. <ul style="list-style-type: none"> - When the Disk Group is not Registered: Deport the disk group. - When the Disk Group is Registered: Offline the disk group resource. 4. Set the physical disks in the disk groups to offline. 	<ul style="list-style-type: none"> - When not using the VERITAS Cluster Server, or when using a version earlier than the VERITAS Cluster Server 6.0 <ol style="list-style-type: none"> 1. Online the physical disks that were set to offline during pre-processing. 2. Reconfigure the volume group. 3. If file systems are included, remount the volumes that were unmounted by pre-processing.

Volume Type	Pre-processing	Post-processing
	<p>- When using a version later than the VERITAS Cluster Server 6.0</p> <ol style="list-style-type: none"> 1. Stop access to all logical volumes in the volume group. 2. Implement the following tasks to all the file systems in the disk group depending on the registration status of the mounting point as a resource of the clustered system. <ul style="list-style-type: none"> - When the Mount Point is not Registered: Unmount the file system. - When the Mount Point is Registered: Instead of unmounting the file system, offline the mounting resource, and then configure the mounting resource Enabled attribute to off. 3. Implement the following tasks depending on the volume registration status in the clustered system volume resource. <ul style="list-style-type: none"> - When the Disk Group is not Registered: Suspend the volume. - When the Disk Group is Registered: Offline the volume, and then configure the volume resource Enabled attribute to off. 4. Implement the following tasks depending on the disk group registration status in the clustered system disk group resource. <ul style="list-style-type: none"> - When the Disk Group is not Registered: Deport the disk group. - When the Disk Group is Registered: Instead of deporting the disk group, offline the disk group resource, and then configure the mounting resource Enabled attribute to off. 5. Set the physical disks in the disk groups to offline. 	<p>- When using a version later than the VERITAS Cluster Server 6.0</p> <ol style="list-style-type: none"> 1. Online the physical disks that were set to offline during pre-processing. 2. Reconfigure the volume group. 3. Deport the disk group. 4. Implement the following tasks depending on the disk group registration status in the clustered system disk group resource. <ul style="list-style-type: none"> - When the Disk Group is not Registered: Import the disk group. - When the Disk Group is Registered: Instead of importing the disk group, after configuring the Enabled attribute of the disk group resource to on, and then online the disk group resource. 5. Implement the following tasks depending on the volume registration status in the clustered system volume resource. <ul style="list-style-type: none"> - When the Disk Group is not Registered: Start the volume. - When the Disk Group is Registered: Configure the volume resource Enabled attribute to on, and then online the volume resource. 6. Implement the following tasks depending on the state of registration of the mount point as a resource of the clustered system. <ul style="list-style-type: none"> - When the Disk Group is not Registered: Mount file systems that were unmounted during pre-processing. - When the Disk Group is Registered: Instead of mounting the file system, after configure the mounting resource Enabled attribute to on, and then online the mounting resource.

Reconfiguring Disk Group (for Replication on single Server)

Reconfigure the disk group as follows:

1. Pre-commit analysis for restoration

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

Installing volume manager disk header for c1t0d20s2 ...
Installing volume manager disk header for c1t0d21s2 ...
-
```

dstdg's diskgroup configuration is restored (in precommit state).
Diskgroup can be accessed in read only and can be examined using
vxprint in this state.

Run:

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

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

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

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

Reconfiguring Disk Group (for Server-to-server Replication)

1. Import the disk group "srcdg" on the master server as disk group "dstdg" on the target server.

```
# /usr/sbin/vxdg -C -n dstdg import srcdg  
#
```

Note

When the disk group name is the same in the server of the source and destination, the -n option is not specified.

2. Execute recovery processing for the volume in the disk group "dstdg" on the target server.

```
# vxrecover -g dstdg -sb  
#
```

3. Because udid of the copy destination disk is not an adjustment, this is restored.

```
# vxdisk updateudid c1t0d20s2 c1t0d21s2  
#
```

Note

- When performing a system disk exchange, there are cases when conflicting backup configuration information may exist. In such cases, the disk group ID needs to be reset after executing the above command.
- After this operation, where the volumes within a disk group are required to be run in synchronous mode in background, synchronous processing may take some time depending on the volume configuration. It is also possible to use the volumes during this time.

Chapter 7 Storage Cluster Continuous Copy Function


This chapter provides details of the replication operation of AdvancedCopy Manager using the Storage Cluster Continuous Copy function.

7.1 Overview

The Storage Cluster Continuous Copy function is an ETERNUS Disk storage function that maintains consistency on both storage devices by copying in the Primary Storage and Secondary Storage simultaneously. In a Storage Cluster environment, if a failover occurs in the Primary Storage while executing Advanced Copy, this is a function that can continue the copy operation with the Secondary Storage by automatically taking over the Advanced Copy process. For this reason, even if the power is interrupted due to a power cut, or a box failure occurs, the Advanced Copy operation can be continued without stopping the copy session.

Refer to "Storage Cluster Function" in the *ETERNUS SF Storage Cruiser Operation Guide* for Storage Cluster function.

Figure 7.1 Storage Cluster Continuous Copy Function



The replication operation procedure that uses the Storage Cluster Continuous Copy function is the same as the replication operation for volumes (hereinafter referred to as "non-TFOV") that are not the target for Storage Cluster. When executing Advanced Copy from a TFOV to a TFOV, operation of the Storage Cluster Continuous Copy function starts the copy sessions on both the Primary Storage and Secondary Storage. For this reason, Advanced Copy can be executed without considering the storage device status.

On the AdvancedCopy Manager, sessions on the storage device where the TFO status is "Active" are referred to as "Active-Side Continuous Copy Session" and sessions on the storage device where the TFO status is "Standby" are referred to as "Standby-Side Continuous Copy Session".

Point


- Advanced Copy can be executed even if either the Primary Storage or the Secondary Storage is stopped.
- The Storage Cluster Continuous Copy function cannot be used when copying between non-TFOVs. Use the conventional Advanced Copy.
- The Storage Cluster Continuous Copy function cannot be used when copying between a TFOV and non-TFOV. When performing an Advanced Copy between a TFOV and non-TFOV, use the AdvancedCopy Manager Copy Control Module. Refer to "Operations of Advanced Copy Using TFOV" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module* for details.

- Remote Advanced Copy is not supported between TFOVs.
When performing a Remote Advanced Copy between a TFOV and non-TFOV, use the AdvancedCopy Manager Copy Control Module. Refer to "Operations of Advanced Copy Using TFOV" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module* for details.

7.1.1 System Configuration

A system configuration using the Storage Cluster Continuous Copy function is shown below.

Figure 7.2 System Configuration Using Storage Cluster Continuous Copy Function



7.1.2 Operating Requirements

Refer to "Storage Cluster Function" in the *ETERNUS SF Storage Cruiser Operation Guide*.

7.1.3 Support Function

The support range of a Storage Cluster Continuous Copy function is shown below.

Advanced Copy Function	Support Range	
Snapshot type high-speed copies	OPC	A
	QuickOPC	A
	SnapOPC	N/A
	SnapOPC+	N/A
Synchronous high-speed copies	EC	A
	REC	N/A

A=Available, N/A=Not Available

The range of functions supported by the Storage Cluster Continuous Copy function of AdvancedCopy Manager is shown below.

Function	Support Range
Backup operation	N/A
Backup operation using Suspend/Resume function	N/A
Backup operation using QuickOPC function	N/A
Replication operation	A

A=Available, N/A=Not Available

7.2 Operation Design

This section provides details of the replication operation design using the Storage Cluster Continuous Copy function.

Replication operation design is the same as a non-TFOV. Refer to "6.2 Configuration" for details.


Note

- SnapOPC and SnapOPC+ are not supported.
- Make sure that the version of AdvancedCopy Manager's agent on the target server is Version 16.2 or later.

7.3 Flow of Operation

The flow of the replication operation when using the Storage Cluster Continuous Copy function is as described below.

Figure 7.3 Flow of Replication Operation When Using Storage Cluster Continuous Copy Function



7.4 Preparations

This section describes the preparations for using the Storage Cluster Continuous Copy function with replication operations.

Set up using the following procedure:

1. Starting Daemon
Normally, the daemon is configured to be automatically invoked at system startup.
However, if a daemon fails to start for some reason or has stopped, both servers must restart their respective daemons. For information about starting daemons, refer to "[Chapter 2 Starting and Stopping Daemons](#)".
2. Starting Web Console
Start Web Console. For details, refer to "How to Start Web Console" in the *ETERNUS SF Web Console Guide*.
3. Registering Managed Server
The Management Server registers the Managed Server to be managed.
Refer to "[6.4.3 Registering Managed Server](#)" for details.
4. Configuration of Storage Cluster
Set up the Storage Cluster environment.
Refer to "Storage Cluster Function" in the *ETERNUS SF Storage Cruiser Operation Guide* for Storage Cluster function setup.
5. Fetching Device Information on Managed Server
Fetch the device information on a Managed Server.
Refer to "[6.4.4 Fetching Device Information on Managed Server](#)" for details.
6. Configuring Source Volume and Replica Volume
Set the source volume and replica volume that are used in the replication.
Refer to "[6.4.5 Configuring Source Volume and Replica Volume](#)" for details.
When creating groups, refer to "[6.4.7 Creating Groups](#)".

Note

- Register the Advanced Copy license for the respective ETERNUS Disk storage systems on the Primary Storage and Secondary Storage. Refer to "Registration of License" in the *ETERNUS SF Installation and Setup Guide* for how to register the Advanced Copy license.
- The source volume and replica volume must belong to the same TFO group.
- All of the replication volumes registered in the group must belong to the same TFO group.
- Non-TFOVs cannot be registered as source and replica volumes.
- Replication volumes that are not synchronized between the Primary Storage and Secondary Storage cannot be registered.
- When creating groups, do not perform a failover or a failback.

7. Customizing Pre-processing and Post-processing

Pre- and post-processing scripts are customized according to the operation.

Refer to "[6.4.6 Customizing Pre-processing and Post-processing](#)" for details.

7.5 Operations

The replication operation of AdvancedCopy Manager using the Storage Cluster Continuous Copy function is the same as a non-TFOV. Refer to "[6.5 Operations](#)" for details.

Refer to "[Chapter 9 Command References](#)" for the specifications and case examples for each command.

Point

- When the status of the TFO group is "Normal" and Advanced Copy cannot operate on either the Primary Storage or Secondary Storage, the command may terminate abnormally. Check the cause of the error on both the Primary Storage and Secondary Storage to address the error message in the storage device where the error is occurring.
- When the status of the TFO group is "Normal" and its phase is "Maintenance", or when the status of the TFO group is "Halt", the command executes Advanced Copy on the storage device that has an "Active" TFO status and terminates normally. However, Advanced Copy could not be executed on the storage device that has a "Standby" TFO status. In this case, swsrp2882 message may be output.
If a swsrp2882 message is output, after the status of TFO group has become "Normal" and the phase is not "Maintenance", refer to "[7.10 When Recovering Storage Cluster Continuous Copy Function](#)" to take action.
- Advanced Copy cannot be operated during a failover or a failback.
- If the volume that is the target of Storage Cluster or the TFO group of the volume is changed, refer to "[7.11.2 When Changing Storage Cluster Target Volume](#)" and register the replication volume again.
- For TFOVs that are not synchronized between the Primary Storage and Secondary Storage, the following commands cannot be executed:
 - `swsrpstartsync`
 - `swsrpmake`
 - `swsrprevsync`

When the TFO status on the Secondary Storage is "Active", the following commands cannot be executed as well:

- `swsrpstat`
- `swsrpcancel`

To check the status of Advanced Copy and stop the session, use Web Console or AdvancedCopy Manager Copy Control Module commands.

- For using Web Console

Refer to "Display/Forcible Suspend/Forcible Cancel Copy Session" in the *ETERNUS SF Web Console Guide*.

- For using AdvancedCopy Manager Copy Control Module commands
 - accopy query
 - accopy fcancel

Refer to "Command References" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module* for information on the commands.

7.6 Backup Operation in Oracle Solaris Zones

This section describes the backup operation procedure using the Storage Cluster Continuous Copy function for Oracle Solaris zones.

The backup operation procedure that uses the Storage Cluster Continuous Copy function is the same as for non-TFOVs. Refer to "[Appendix F Operation in Oracle Solaris Zones](#)" for details.



The Storage Cluster Continuous Copy function cannot be used with the backup management function. By using the Storage Cluster Continuous Copy function, when performing a backup operation in Oracle Solaris zones, use the replication management function.


7.6.1 Operation Design

"[Figure 7.4 Configuration Example When Using Storage Cluster Continuous Copy Function](#)" shows a configuration example of when the Storage Cluster Continuous Copy function is used.

Perform an operational design by referring to the following section.

- [F.2 Configuration](#)
- [F.3 Software Requirements](#)

Figure 7.4 Configuration Example When Using Storage Cluster Continuous Copy Function



Point

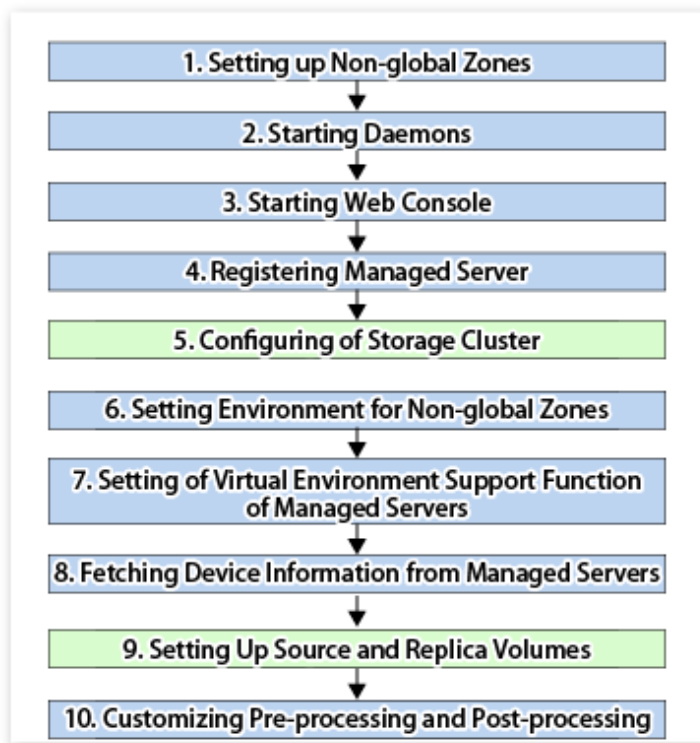
- REC, SnapOPC, and SnapOPC+ are not supported.
- Locate all of the devices that are used in the backup operation in Oracle Solaris Zones in the same TFO group.

7.6.2 Preparations

Below figure shows the flow of preparation tasks.

This section describes steps 5 and 9.

The steps other than 5 and 9 are the same as for non-TFOVs. Refer to the relevant sections of "F.7 Preparations" for details.



7.6.2.1 Configuring of Storage Cluster

Refer to "Storage Cluster Function" in the *ETERNUS SF Storage Cruiser Operation Guide* for Storage Cluster function setup.

7.6.2.2 Setting Up Source and Replica Volumes

Refer to "6.4.5 Configuring Source Volume and Replica Volume" for information on how to set the source and replica volumes for replication.

When using the Storage Cluster Continuous Copy function, note the following points:



Note

- Register the Advanced Copy license for the respective ETERNUS Disk storage systems on the Primary Storage and Secondary Storage. Refer to "Registration of License" in the *ETERNUS SF Installation and Setup Guide* for how to register the Advanced Copy license.
- The source volume and replica volume must belong to the same TFO group.
- All of the replication volumes registered in the group must belong to the same TFO group.
- Non-TFOVs cannot be registered as source and replica volumes.
- Replication volumes that are not synchronized between the Primary Storage and Secondary Storage cannot be registered.
- When setting up, do not perform a failover or a failback.

7.6.3 Operations

Refer to "6.5 Operations". However, "6.4 Preparations" are not required.

Point

- When the status of the TFO group is "Normal" and Advanced Copy cannot operate on either the Primary Storage or Secondary Storage, the command may terminate abnormally. Check the cause of the error on both the Primary Storage and Secondary Storage to address the error message in the storage device where the error is occurring.
- When the status of the TFO group is "Normal" and its phase is "Maintenance", or when the status of the TFO group is "Halt", the command executes Advanced Copy on the storage device that has an "Active" TFO status and terminates normally. However, Advanced Copy could not be executed on the storage device that has a "Standby" TFO status. In this case, swsrp2882 message may be output.
If a swsrp2882 message is output, after the status of TFO group has become "Normal" and the phase is not "Maintenance", refer to ["7.10 When Recovering Storage Cluster Continuous Copy Function"](#) to take action.
- Advanced Copy cannot be operated during a failover or a failback.
- For TFOVs that are not synchronized between the Primary Storage and Secondary Storage, the following commands cannot be executed:
 - [swsrpstartsync](#)
 - [swsrpmake](#)
 - [swsrprevsync](#)

When the TFO status on the Secondary Storage is "Active", the following commands cannot be executed, either:

- [swsrpstat](#)
- [swsrpcancel](#)

To check the status of Advanced Copy and stop the session, use Web Console or AdvancedCopy Manager Copy Control Module commands.

- For using Web Console
Refer to "Display/Forcible Suspend/Forcible Cancel Copy Session" in the *ETERNUS SF Web Console Guide*.
- For using AdvancedCopy Manager Copy Control Module commands
 - [accopy query](#)
 - [accopy fcancel](#)

Refer to "Command References" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module* for information on the commands.

7.6.4 Reconfiguration

Refer to ["F.9 Changing Operating Environment"](#) for information on how to change the operating environment.

Point

If the volume that is the target of Storage Cluster or the volume TFO group is changed, refer to ["7.11.2 When Changing Storage Cluster Target Volume"](#) and register the replication volume again.

7.6.5 Available Commands

When using the Storage Cluster Continuous Copy function, the commands that can be used with the KVM guest are as follows:

- [9.2 Configuration Management Commands](#)
- [9.3 Replication Management Commands](#)
- [9.6 Troubleshooting Data Collection Command](#)

7.6.6 Specifications of Settings File

Refer to "[F.11 Specifications of Settings File](#)" for the specifications of the setup file.

7.7 Backup and Restoration of ZFS File Systems

This section describes the procedure for backing up/restoring the ZFS file system by using the Storage Cluster Continuous Copy function.

The backup/restore procedure that uses the Storage Cluster Continuous Copy function is the same as for non-TFOVs. Refer to "[Appendix G Backup and Restoration of ZFS File Systems](#)" for details.



Storage Cluster must be configured in advance.

- For a replication operation:
A configuration is performed for using AdvancedCopy Manager in advance according to "[6.4 Preparations](#)", however, configure Storage Cluster before "[6.4.4 Fetching Device Information on Managed Server](#)".
- For AdvancedCopy Manager CCM operations:
A configuration is performed for using AdvancedCopy Manager in advance according to "Setup" in "Storage Cluster Continuous Copy Function" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module*, however, configure Storage Cluster before "Creating copy groups".

Refer to "Storage Cluster Function" in the *ETERNUS SF Storage Cruiser Operation Guide* for Storage Cluster function setup.

The following methods are supported for the ZFS file system backup/restore using the Storage Cluster Continuous Copy function.


- [7.7.1 Online Backup Using ZFS Snapshot with AdvancedCopy Manager Function](#)
- [7.7.2 Online Backup Using ZFS Snapshot with ZFS Uses zfs send and zfs receive Commands](#)
- [7.7.3 Offline Backup Using Export and Import Functions of ZFS \(System Stop or Disconnection of Storage Pool\)](#)

Refer to "[G.1.2 Backup Methods](#)" for details of each method.

7.7.1 Online Backup Using ZFS Snapshot with AdvancedCopy Manager Function

A configuration example of when the replication operation uses the Storage Cluster Continuous Copy function is shown below.

Figure 7.5 Configuration Example of Replication Operation Using Storage Cluster Continuous Copy Function



Details on Sample Configuration

Using 1 physical disk that constitutes a storage pool for transactions, the entire disk is registered.


Backups are saved in the backup storage pool.

- Business server (Managed Server): SV
- Backup server(Management Server): BKSV (name for registration in AdvancedCopy Manager)
- Transaction storage pool: STP1
- Backup storage pool: STP1
- Disk

Physical Disk	Connection Server	Device Type	Device Name	Purpose (Storage Pool)
Transaction VOL	SV	LU	c1t1d1	Transaction storage pool (STP1)
Backup VOL	BKSV	LU	c2t1d11	Backup storage pool (STP1)

A configuration example of when the AdvancedCopy Manager CCM operation uses the Storage Cluster Continuous Copy function is shown below.

Figure 7.6 Configuration Example of AdvancedCopy Manager CCM Operation Using Storage Cluster Continuous Copy Function



Details on Sample Configuration

Using 1 physical disk that constitutes a storage pool for transactions, the entire disk is registered.

Backup is saved in the backup storage pool.


- Business server: SV
- CCM Server: BKS
- Transaction storage pool: STP1
- Backup storage pool: STP1
- Disk

Physical Disk	Connection Server	LUN	Device Type	Device Name	Purpose (Storage Pool)
Transaction VOL	SV	0x01	LU	c1t1d1	Transaction storage pool (STP1)
Backup VOL	BKS	0x11	LU	c2t1d11	Backup storage pool (STP1)

Note

- REC between the transaction storage pool and the backup storage pool is not supported.
- Locate all devices in the same TFO group.
- Consider the following points when operating a backup/restore of the ZFS file system by using the backup volume.
 - TFOVs cannot be specified as a backup volume.
 - Use QuickOPC for Advanced Copy from the device that configures the ZFS storage pool to the backup storage pool.
 - Use AdvancedCopy Manager CCM for Remote Advanced Copy from the backup storage pool to the backup volume.

Figure 7.7 Operating Example of Backup/Restore for ZFS File System



Information

- If the device that configures the ZFS Storage pool is TFOV, the usable copy type in Advanced Copy to the backup volume is REC only.
- Refer to the following sections and operate after understanding the design and operation methods.
 - [7.12 Cascade Copies](#)
 - "Operations of Advanced Copy Using TFOV" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module*
 - [G.2 Online Backup Using ZFS Snapshot with AdvancedCopy Manager Function](#)

7.7.1.1 Preparations

Implement the following procedures for preparation:

1. [Fetching Lists of Devices in Transaction Storage Pools](#)
2. [Copy Settings for Backup and Restoration](#)

Fetching Lists of Devices in Transaction Storage Pools

The method for acquiring the device list within the transaction storage pool is the same as for non-TFOVs. Refer to ["Fetching Lists of Devices in Transaction Storage Pools"](#) in ["G.2.3 Preparations"](#) for details.

Copy Settings for Backup and Restoration

Execute the copy settings for the transaction storage pool and the backup storage pool (for backup and restoration).

Use the following combinations to register them in AdvancedCopy Manager:

Table 7.1 Copy Groups

Used for	Copy Source	Copy Destination	Group Name (ex.)
Backup/Restoration	Transaction storage pool	Backup storage pool	STP1

Table 7.2 Registered Devices

Device Name	Devices to Be Registered	Example (*1)	
		Device Name	LUN
Transaction storage pool	Devices that you confirmed in "Fetching Lists of Devices in Transaction Storage Pools" in "G.2.3 Preparations" to be devices constituting the transaction storage pool	c1t1d1	0x01
Backup storage pool	Devices that constitute the backup storage pool	c2t1d11	0x11

*1: When using replication management function, use device name. When using AdvancedCopy Manager CCM, use LUN number.

For Using Replication Management Function

Register copy sources indicated at ["Table 7.1 Copy Groups"](#) as copy destination volumes, and copy destinations as copy destination volumes. Then register the devices indicated at ["Table 7.2 Registered Devices"](#) respectively.

Since these are controlled synchronously in processing for backup and restoration, you have to register all devices that constitute a storage pool as a single group. For details on controlling groups, refer to ["6.4.7 Creating Groups"](#).

[Example]

Specify the source volumes and copy destination volumes to create the 3 groups shown in the chart.

```
# /opt/FJSVswsrp/bin/swsrpsetvol -o BOTH -Xgroup STP1 /dev/dsk/c1t1d1@SV /dev/dsk/c2t1d11@BKSV
swsrpsetvol completed
# /opt/FJSVswsrp/bin/swsrpvoinfo -L
Server Original-Volume      Size      Replica-Volume      Size      Copy      Op-Server
Group
BKSV /dev/dsk/c1t1d1@SV      4.0 Gbyte /dev/dsk/c2t1d11@BKSV 4.0 Gbyte bi-direction
both STP1
#
```

For Using AdvancedCopy Manager CCM

Register each combination of copy sources and copy destinations.

Since these are controlled synchronously in processing for backup and restoration, you have to register all devices that constitute a storage pool as a single group. For details on controlling groups, refer to ["Creating Copy Groups"](#) and ["Adding Copy Pairs"](#) in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module*.

1. Creation of copy groups

Create a copy group for each copy setting.

[Example]

Create the 3 copy groups shown in the chart. (Copying type: OPC)

```
# acgroup create -g STP1 -type OPC -a DX440
Successful completion.
#
```

2. Registration of copy pairs

In the created copy groups, register the combinations of copy sources and copy destinations. Register all devices that constitute a storage pool.

[Example]

In the created copy groups, register each device shown in the chart.

```
# acpair add -g STP1 -p DX440/0x01:DX440/0x11
Successful completion.
#
```

7.7.1.2 Operation

This section describes procedures for backing up and restoring ZFS storage pools. Refer to "[Appendix G Backup and Restoration of ZFS File Systems](#)" for details.

Point

- When the status of the TFO group is "Normal" and Advanced Copy cannot operate on either the Primary Storage or Secondary Storage, the command may terminate abnormally. Check the cause of the error on both the Primary Storage and Secondary Storage to address the error message in the storage device where the error is occurring.
- When the status of the TFO group is "Normal" and its phase is "Maintenance", or when the status of the TFO group is "Halt", the command executes Advanced Copy on the storage device that has an "Active" TFO status and terminates normally. However, Advanced Copy could not be executed on the storage device that has a "Standby" TFO status. In this case, swsrp2882 message may be output.
If a swsrp2882 message is output, after the status of TFO group has become "Normal" and the phase is not "Maintenance", refer to "[7.10 When Recovering Storage Cluster Continuous Copy Function](#)" to take action.
- Advanced Copy cannot be operated during a failover or a failback.
- If the volume that is the target of Storage Cluster or the TFO group of the volume is changed, refer to "[7.11.2 When Changing Storage Cluster Target Volume](#)" and register the replication volume again.
- For TFOVs that are not synchronized between the Primary Storage and Secondary Storage, the following commands cannot be executed:
 - [swsrpstartsync](#)
 - [swsrpmake](#)
 - [swsrprevsync](#)

When the TFO status on the Secondary Storage is "Active", the following commands cannot be executed as well:

- [swsrpstat](#)
- [swsrpcancel](#)

To check the status of Advanced Copy and stop the session, use Web Console or AdvancedCopy Manager Copy Control Module commands.

- For using Web Console
Refer to "Display/Forcible Suspend/Forcible Cancel Copy Session" in the *ETERNUS SF Web Console Guide*.
- For using AdvancedCopy Manager Copy Control Module commands
 - [accopy query](#)
 - [accopy fcancel](#)

Refer to "Command References" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module* for information on the commands.

Backup

This section shows the procedure for a online backup.

No.	Description	Controlling Server
1	Export the backup storage pool.	Backup server
2	Create a snapshot of the transaction storage pool.	Business server
3	Using AdvancedCopy Manager, back up data to a device in the backup storage pool from a device in the transaction storage pool.	Business server
4	Check the status of the backup storage pool.	Backup server

1. Export the backup storage pool.

For more information on export of ZFS storage pools, refer to the operating system manual.

[Example]

```
# zpool list
NAME    SIZE  USED   AVAIL   CAP  HEALTH  ALTROOT
STP1    345G  121M   345G    0%   ONLINE  -
# zpool export STP1
#
```

2. Create snapshots of the transaction storage pool.

a. Delete the snapshot that is not restored to the transaction storage pool.

To delete ZFS snapshots, refer to the operating system manual.

b. Execute the sync command to synchronize the transaction storage pool with the file system.

c. Create a file system snapshot on the transaction storage pool.

To create ZFS snapshots, refer to the operating system manual.

d. Execute the sync command to synchronize the transaction storage pool with the file system.

[Example]

```
# zfs destroy STP1@SNAP0
# sync
# zfs snapshot -r STP1@SNAP0
# sync
#
```

3. Using AdvancedCopy Manager, back up data to a device in the backup storage pool from a device in the transaction storage pool.

For Using Replication Management Function

Specify backup/restore group "STP1" for the -Xgroup option and execute the [swsrpmake](#) command.

[Example]

```
# /opt/FJVSwsrp/bin/swsrpmake -Xgroup STP1
GROUP=STP1 swsrpmake completed
#
```

For Using AdvancedCopy Manager CCM

Specify backup/restore group "STP1" for the -g option and execute the `acopc start` command for AdvancedCopy Manager CCM.

[Example]

```
# /opt/FJSSVccm/bin/acopc start -g STP1
DX440/0x01:DX440/0x11
# DATE : 2012/01/01 00:00:00 - << OPC Started >>
#
From:BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303034232
3/Olu=01/Adr_high=0/Adr_low=0/size_high=0/size_low=0
#
To:BoxID=303045343030304d3323232323234534353053323041232323234b44343033303633393030342323/
Olu=11/ Adr_high=0/Adr_low=0
# Standby-Session : DX440SEC/0x02:DX440SEC/0x12
Succeeded : 1
Failed : 0
#
```

 Note

Since processing is carried out for all devices that constitute the storage pool, do not specify any copy pairs.

4. Check the status of the backup storage pool.
 - a. Import the backup storage pool on the backup server.
 - b. Export the backup storage pool on the backup server.

For more information on import/export of the ZFS storage pool, refer to the operating system manual.

[Example]

```
# zpool import
pool: STP1
id: 15271135801818262407
state: ONLINE
action: The pool can be imported using its name or numeric identifier.
config:

        STP1      ONLINE
        c2t1d11  ONLINE
# zpool import -f STP1
# zpool list
NAME  SIZE  USED   AVAIL  CAP  HEALTH  ALTROOT
STP1  345G  345M   345G   0%   ONLINE  -
# zpool export STP1
#
```

 Note

If import/export has failed, re-execute it from step 2.

Restoration

This section shows the procedure for restoration.

No.	Description	Controlling Server
1	Export the transaction storage pool.	Business server
2	Export the backup storage pool.	Backup server
3	Using AdvancedCopy Manager, restore data to a device in the transaction storage pool from a device in the backup storage pool.	Business server or Backup server

No.	Description	Controlling Server
4	Import the transaction storage pool.	Business server
5	Roll back the transaction storage pool to a restored snapshot.	Business server
6	Mount the target ZFS file systems.	Business server

1. If any transaction storage pool exists, export it.

For more information on export of ZFS storage pools, refer to the operating system manual.

[Example]

```
# zpool list
NAME    SIZE    USED    AVAIL    CAP    HEALTH    ALTROOT
STP1    345G    345M    345G    0%    ONLINE    -
# zpool export STP1
#
```

2. Export the backup storage pool.

For more information on export of ZFS storage pools, refer to the operating system manual.

[Example]

```
# zpool list
NAME    SIZE    USED    AVAIL    CAP    HEALTH    ALTROOT
STP1    345G    121M    345G    0%    ONLINE    -
# zpool export STP1
#
```

3. Using AdvancedCopy Manager, restore data to a device in the transaction storage pool from a device in the backup storage pool.

When Restoring Data Using Replication Management Function

Specify backup/restore group "STP1" for the -Xreverse option and the -Xgroup option and execute the [swsrpmake](#) command.

[Example]

```
# /opt/FJSVswsrp/bin/swsrpmake -Xgroup STP1 -Xreverse
GROUP=STP1 swsrpmake completed
#
```

When Restoring Data Using AdvancedCopy Manager CCM

Specify backup/restore group "STP1" for the -r option and the -g option and execute the `acopc start` command for AdvancedCopy Manager CCM.

[Example]

```
# /opt/FJSVccm/bin/acopc start -g STP1 -r
DX440/0x11:DX440/0x01
# DATE : 2012/01/01 00:00:00 - << OPC Started >>
# From:
BoxID=303045343030304d332323232323234534353053323041232323234b44343033303633393030342323/
Olu=11/Adr_high=0/Adr_low=0
# To:
BoxID=303045343030304d332323232323234534353053323041232323234b44343033303633393030342323/
Olu=01/Adr_high=0/Adr_low=0/size_high=0/size_low=0
# Standby-Session : DX440SEC/0x12:DX440SEC/0x02
Succeeded : 1
Failed : 0
#
```



Since processing is carried out for all devices that constitute the storage pool, do not specify any copy pairs.

4. In the transaction storage pool, import the device restored at step 3 in the restore procedure.

For information on how to import a ZFS storage pool, refer to the operating system manuals.

[Example]

```
# zpool import STP1
#
```

5. Roll back the transaction storage pool to a restored snapshot.

To see how to roll back the ZFS storage pool to a snapshot, refer to the operating system manual.

[Example]

For rolling back to the "SNAP1"

```
# zfs list -t snapshot
NAME          USED  AVAIL  REFER  MOUNTPOINT
STP1@SNAP1    0      -    19K    -
STP1@SNAP2   16M     -    18K    -
STP1@SNAP3   32M     -    18K    -
# zfs rollback -r STP1@SNAP1
#
```

6. Mount the target ZFS file systems.

If the ZFS file systems have not been mounted automatically by step 5, mount them manually.

For information on how to mount ZFS file systems, refer to the operating system manuals.


[Example]

```
# zfs mount STP1/data1
# zfs mount STP1/data2
#
```

7.7.2 Online Backup Using ZFS Snapshot with ZFS Uses `zfs send` and `zfs receive` Commands

A configuration example of when the replication operation uses the Storage Cluster Continuous Copy function is shown below.

Figure 7.8 Configuration Example of Replication Operation Using Storage Cluster Continuous Copy Function



Details on Sample Configuration

Using 1 physical disk that constitutes a storage pool for transactions, the entire disk is registered.

Including the backup storage pool, backup can preserve a maximum of 2 generations.

- Business server (Managed Server): SV-1
- Backup server (Management Server): BKS (name for registration in AdvancedCopy Manager)
- Transaction storage pool: STP1
- Backup storage pool: STP1B
- Disk

Physical Disk	Connection Server	Device Type	Device Name	Purpose (Storage Pool)
Transaction VOL	SV-1	LU	c1t1d1	Transaction storage pool (STP1)
Backup VOL-1	BKSV	LU	c2t1d11	Backup storage pool (STP1B)
Backup VOL-2	BKSV	Slice	c2t1d12s1	Backup volume for older generations

Note

With the method for performing an online backup using the `zfs send/zfs receive` commands and the ZFS snapshot function, the AdvancedCopy Manager Copy Control Module cannot be used. When performing an online backup by using the Storage Cluster Continuous Copy function, use the replication management function of the AdvancedCopy Manager.

Point

- REC between the backup storage pool and the backup volume is not supported.
- Locate all devices in the same TFO group.

7.7.2.1 Preparations

Refer to "[G.3.3 Preparations](#)" for details.

7.7.2.2 Operation

- Backup

Refer to "[Full Backup](#)" and "[Differential Backup](#)" in "[G.3.4 Operation](#)" for details.

- Restore

Refer to "[Restoration from Latest Generations](#)" and "[Restoration from Older Generations](#)" in "[G.3.4 Operation](#)" for details.

However, when backing up from previous generations, the following procedure is required.

1. When executing Advanced Copy from the backup storage pool to the backup volume, stop that session.
2. Restore from the backup volume to the transaction storage pool.

Refer to "[Appendix G Backup and Restoration of ZFS File Systems](#)" for details.

Point

- When the status of the TFO group is "Normal" and Advanced Copy cannot operate on either the Primary Storage or Secondary Storage, the command may terminate abnormally. Check the cause of the error on both the Primary Storage and Secondary Storage to address the error message in the storage device where the error is occurring.
- When the status of the TFO group is "Normal" and its phase is "Maintenance", or when the status of the TFO group is "Halt", the command executes Advanced Copy on the storage device that has an "Active" TFO status and terminates normally. However, Advanced Copy could not be executed on the storage device that has a "Standby" TFO status. In this case, `swsrp2882` message may be output.
If a `swsrp2882` message is output, after the status of TFO group has become "Normal" and the phase is not "Maintenance", refer to "[7.10 When Recovering Storage Cluster Continuous Copy Function](#)" to take action.
- Advanced Copy cannot be operated during a failover or a failback.
- If the volume that is the target of Storage Cluster or the TFO group of the volume is changed, refer to "[7.11.2 When Changing Storage Cluster Target Volume](#)" and register the replication volume again.

- For TFOVs that are not synchronized between the Primary Storage and Secondary Storage, the following commands cannot be executed:

- `swsrpstartsync`
- `swsrpmake`
- `swsrprevsync`

When the TFO status on the Secondary Storage is "Active", the following commands cannot be executed as well:

- `swsrpstat`
- `swsrpcancel`

To check the status of Advanced Copy and stop the session, use Web Console or AdvancedCopy Manager Copy Control Module commands.

- For using Web Console

Refer to "Display/Forcible Suspend/Forcible Cancel Copy Session" in the *ETERNUS SF Web Console Guide*.

- For using AdvancedCopy Manager Copy Control Module commands


- `accopy query`
- `accopy fcancel`

Refer to "Command References" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module* for information on the commands.

7.7.3 Offline Backup Using Export and Import Functions of ZFS (System Stop or Disconnection of Storage Pool)

A configuration example of when the replication operation uses the Storage Cluster Continuous Copy function is shown below.


Figure 7.9 Configuration Example of Replication Operation Using Storage Cluster Continuous Copy Function.



Details on Sample Configuration

- Business server (Managed Server): SV-1 (name for registration in AdvancedCopy Manager)
- Backup server (Management Server): BKSV (name for registration in AdvancedCopy Manager)
- Transaction storage pool: STP
- Disk

Physical Disk	Connection Server	Device Name	Purpose (Storage Pool)
Transaction VOL	SV-1	c1t1d1	STP
Backup VOL-1	BKSV	c1t1d11	
Backup VOL-2	BKSV	c1t1d12	



Details on Sample Configuration

- Business server: SV-1
- CCM Server: BKSX
- Transaction storage pool: STP
- Disk

Physical Disk	LUN Number	Connection Server	Device Name	Purpose (Storage Pool)
Transaction VOL	0x02	SV-1	c1t1d2	STP
Backup VOL	0x12	BKSX	c1t1d12	

Note

- Backup/Restore of the system area is not supported.
- When using AdvancedCopy Manager CCM
 - Connect the copy destination volume to a different server from the copy source volume.
 - Multiple generations cannot be copied.

Point

- REC between the transaction storage pool and the backup volume is not supported.
- Locate all devices in the same TFO group.

7.7.3.1 Preparations

Refer to "[G.4.3 Preparations](#)" for details.

7.7.3.2 Operation

Refer to "[Backup and Restoration of Data Areas](#)" in "[G.4.4 Operation](#)" for details.

"[Backup and Restoration of System Areas](#)" cannot use the Storage Cluster Continuous Copy function.

Point

- When the status of the TFO group is "Normal" and Advanced Copy cannot operate on either the Primary Storage or Secondary Storage, the command may terminate abnormally. Check the cause of the error on both the Primary Storage and Secondary Storage to address the error message in the storage device where the error is occurring.
- When the status of the TFO group is "Normal" and its phase is "Maintenance", or when the status of the TFO group is "Halt", the command executes Advanced Copy on the storage device that has an "Active" TFO status and terminates normally. However, Advanced Copy could not be executed on the storage device that has a "Standby" TFO status. In this case, swsrp2882 message may be output.
If a swsrp2882 message is output, after the status of TFO group has become "Normal" and the phase is not "Maintenance", refer to "[7.10 When Recovering Storage Cluster Continuous Copy Function](#)" to take action.
- Advanced Copy cannot be operated during a failover or a failback.
- If the volume that is the target of Storage Cluster or the TFO group of the volume is changed, refer to "[7.11.2 When Changing Storage Cluster Target Volume](#)" and register the replication volume again.
- For TFOVs that are not synchronized between the Primary Storage and Secondary Storage, the following commands cannot be executed:
 - `swsrpstartsync`
 - `swsrpmake`
 - `swsrprevsync`

When the TFO status on the Secondary Storage is "Active", the following commands cannot be executed as well:

- `swsrpstat`
- `swsrpcancel`

To check the status of Advanced Copy and stop the session, use Web Console or AdvancedCopy Manager Copy Control Module commands.

- For using Web Console
Refer to "Display/Forcible Suspend/Forcible Cancel Copy Session" in the *ETERNUS SF Web Console Guide*.
- For using AdvancedCopy Manager Copy Control Module commands
 - `accopy query`
 - `accopy fcancel`

Refer to "Command References" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module* for information on the commands.

7.7.4 Maintenance

Refer to "[G.5 Maintenance](#)" for details.

7.8 Oracle VM Server for SPARC

This section describes the procedure for backing up/restoring the virtual environment of Oracle VM Server for SPARC by using the Storage Cluster Continuous Copy function.

The backup/restore procedure that uses the Storage Cluster Continuous Copy function is the same as for non-TFOVs. Refer to "[Appendix H Operation in Oracle VM Server for SPARC](#)" for details.

Note

The backup management function cannot use the Storage Cluster Continuous Copy function. When performing an operation in Oracle VM Server for SPARC by using the Storage Cluster Continuous Copy function and AdvancedCopy Manager's agent, use the replication management function.

7.8.1 Operation Design

Perform an operational design by referring to the following procedure.

- [H.1 Overview](#)
- [H.2 Software Requirements](#)

Point


- REC, SnapOPC, and SnapOPC are not supported.
- Locate all of the devices that are used in the operation in Oracle VM Server for SPARC in the same TFO group.

7.8.2 Preparations

Below figure shows the flow of preparation tasks.

This section describes steps 4 and 6.

The steps other than 4 and 6 are the same as for non-TFOVs. Refer to the relevant sections of "[F.7 Preparations](#)" for details,



7.8.2.1 Configuring of Storage Cluster

Refer to "Storage Cluster Function" in the *ETERNUS SF Storage Cruiser Operation Guide* for Storage Cluster function setup.

7.8.2.2 Setting Up Source and Replica Volumes

Refer to "[6.4.5 Configuring Source Volume and Replica Volume](#)" for information on how to set the source and replica volumes for replication.

When using the Storage Cluster Continuous Copy function, note the following points:



Note

- Register the Advanced Copy license for the respective ETERNUS Disk storage systems on the Primary Storage and Secondary Storage. Refer to "Registration of License" in the *ETERNUS SF Installation and Setup Guide* for how to register the Advanced Copy license.
- The source volume and replica volume must belong to the same TFO group.
- All of the replication volumes registered in the group must belong to the same TFO group.
- Non-TFOVs cannot be registered as source and replica volumes.
- Replication volumes that are not synchronized between the Primary Storage and Secondary Storage cannot be registered.
- When setting up, do not perform a failover or a failback.

7.8.3 Operation

Refer to "[6.5 Operations](#)". However, "[6.4 Preparations](#)" are not required.

When using the Storage Cluster Continuous Copy function, note the following points:



Point

- When the status of the TFO group is "Normal" and Advanced Copy cannot operate on either the Primary Storage or Secondary Storage, the command may terminate abnormally. Check the cause of the error on both the Primary Storage and Secondary Storage to address the error message in the storage device where the error is occurring.
- When the status of the TFO group is "Normal" and its phase is "Maintenance", or when the status of the TFO group is "Halt", the command executes Advanced Copy on the storage device that has an "Active" TFO status and terminates normally. However, Advanced Copy could not be executed on the storage device that has a "Standby" TFO status. In this case, swsrp2882 message may be output.
If a swsrp2882 message is output, after the status of TFO group has become "Normal" and the phase is not "Maintenance", refer to "[7.10 When Recovering Storage Cluster Continuous Copy Function](#)" to take action.
- Advanced Copy cannot be operated during a failover or a failback.
- If the volume that is the target of Storage Cluster or the TFO group of the volume is changed, refer to "[7.11.2 When Changing Storage Cluster Target Volume](#)" and register the replication volume again.
- For TFOVs that are not synchronized between the Primary Storage and Secondary Storage, the following commands cannot be executed:
 - `swsrpstartsync`
 - `swsrpmake`
 - `swsrprevsync`

When the TFO status on the Secondary Storage is "Active", the following commands cannot be executed as well:

- `swsrpstat`
- `swsrpcancel`

To check the status of Advanced Copy and stop the session, use Web Console or AdvancedCopy Manager Copy Control Module commands.

- For using Web Console

Refer to "Display/Forcible Suspend/Forcible Cancel Copy Session" in the *ETERNUS SF Web Console Guide*.

- For using AdvancedCopy Manager Copy Control Module commands

- accopy query
- accopy fcancel

Refer to "Command References" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module* for information on the commands.



7.9 Operating During Maintenance

When using the Storage Cluster Continuous Copy function, even if either the Primary Storage or Secondary Storage is stopped, Advanced Copy operation on the other storage can be continued. Further, when starting an ETERNUS Disk storage system that has been stopped, Primary Storage and Secondary Storage Advanced Copy automatically synchronizes.

A description of how the Storage Cluster Continuous Copy function operates in a maintenance state is as follows.

Refer to "Maintenance" in "Storage Cluster Function" in the *ETERNUS SF Storage Cruiser Operation Guide* for the Storage Cluster environment of maintenance procedure.

Figure 7.10 When Primary Storage Is Stopped


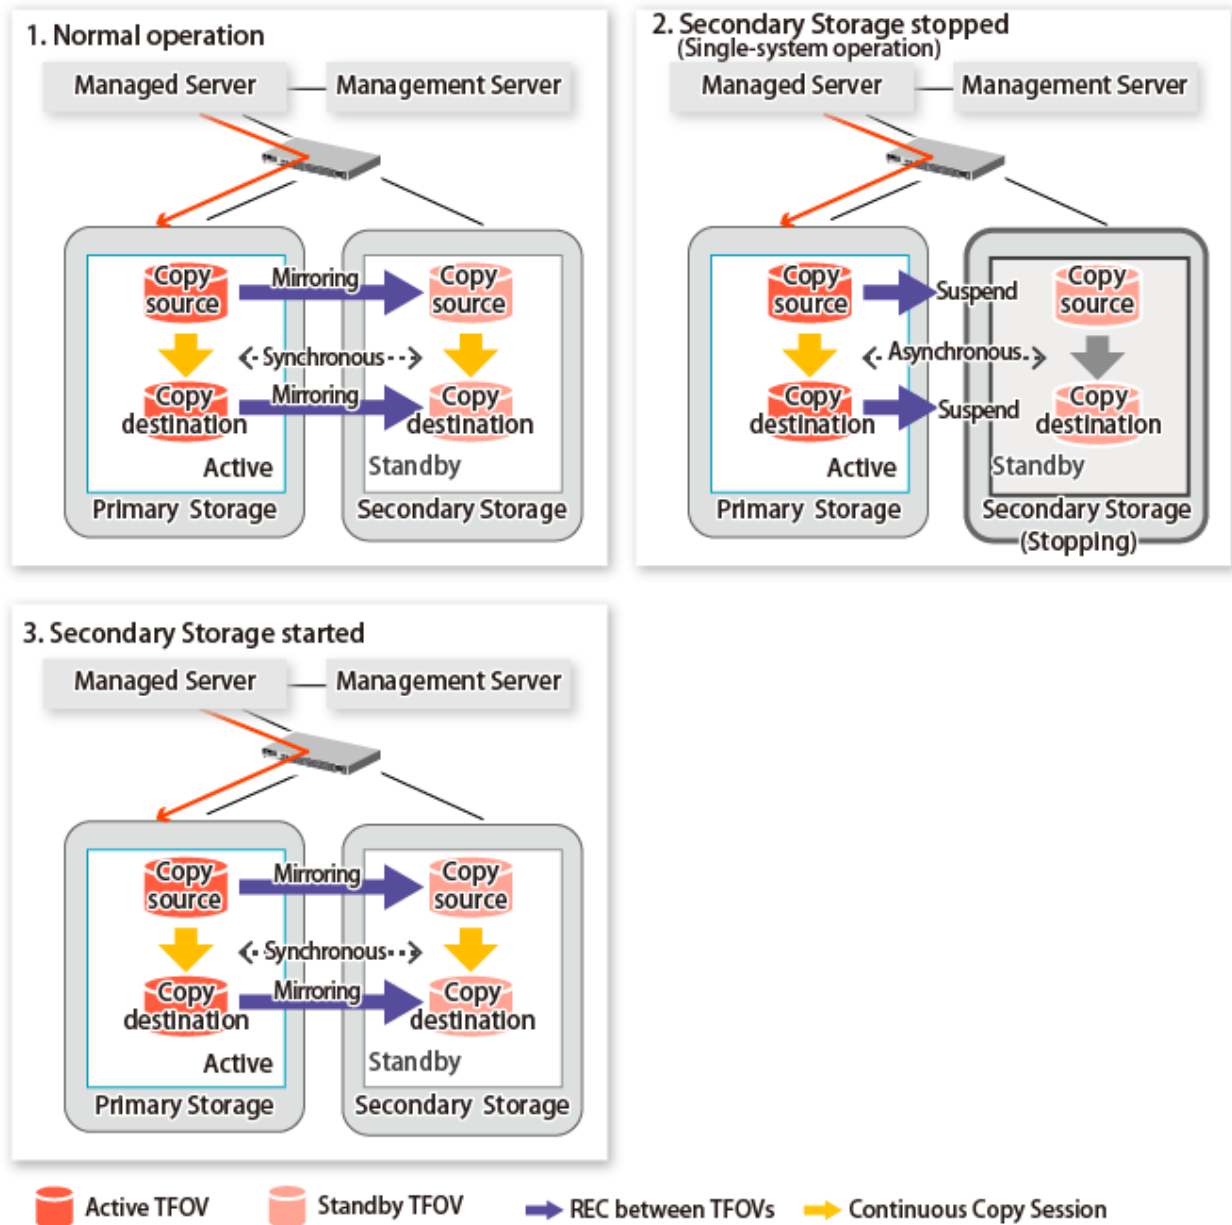


Figure 7.11 When Secondary Storage Is Stopped




Point

- This command starts or stops copying while one side is operating and terminates normally, but there may be times when the swsrp2882 message is output. No measures need to be taken in regard to the swsrp2882 message.
- Even after the stopped ETERNUS Disk storage system is started, it may not be possible to synchronize the Primary Storage and Secondary Storage sessions. In addition, if a failover or a failback is performed in a state where a session is only on either the Primary Storage or the Secondary Storage, the session transitions to an Error Suspend state. For this reason, refer to "7.10 When

Recovering Storage Cluster Continuous Copy Function" and check that the status of the Primary Storage and the Secondary Storage are the same.

Figure 7.12 When Failback Is Carried Out in State Where There Are Only Sessions on Storage Where TFO Status Is "Active"



7.9.1 Recovery from Hardware Failure

This shows the state transition, with the Storage Cluster Continuous Copy function, where either of the Primary Storage or Secondary Storage volumes have a hardware failure.

Refer to "Maintenance" in "Storage Cluster Function" in the *ETERNUS SF Storage Cruiser Operation Guide* for the details of recovery procedure.

When using the Storage Cluster Continuous Copy function, Advanced Copy operation can continue with only a single storage device in the following cases. Stop only the Error Suspend state session that exists within the storage device where failure occurred.

- Hardware failure occurred on the Primary Storage and Automatic Failover was performed
- Hardware failure occurred on the storage device where the TFO status is "Standby"

For cases other than those above, since Advanced Copy cannot continue on a storage device where the TFO status is "Standby" only, the Primary Storage and Secondary Storage session states transitions to Error Suspend. When sessions for both the Primary Storage and Secondary Storage have transitioned to Error Suspend, stop both sessions.

Note

- If a TFO pairing is canceled while recovering from a hardware failure such as a RAID failure and RAID close, since the TFOV between the Primary Storage and Secondary Storage is not synchronized, the following commands cannot be executed until the TFO pairing is restored:
 - `swsrpstartsync`
 - `swsrpmake`
 - `swsrprevsync`

When the TFO status on the Secondary Storage is "Active", the following commands cannot be executed, either:

- `swsrpstat`
- `swsrpcancel`

Back up the copy source volume before cancelling the TFO pairing as required.

- In a storage device where the TFO status is "Active", if there are bad sectors in the copy source logical volume, and if **Status** of the **Volume** tab on the **TFO Group Detail** screen in Web Console is other than "Active", recovering the data according to the usage state of the copy source logical volume is required.
Perform recovery by following the steps outlined below.

1. Use ETERNUS Web GUI to check the Advanced Copy status and the copy process error codes.
 - a. On the [Display status] menu, click [Advanced Copy status display] in the status display.
 - b. At "Session status", click the "Number of active sessions" link for the relevant copy type.
 - c. Refer to the value in the "Error code" column of the relevant copy process.
 2. Refer to "Maintenance" in "Storage Cluster Function" in the *ETERNUS SF Storage Cruiser Operation Guide* to recover the Storage Cluster environment, and recover the bad sector.
 3. Recover data according to the error code checked in step 1.
 - For "0xBA":
Recover data from the copy destination logical volume.
 - Other than "0xBA":
Contact Fujitsu Technical Support.
-

Figure 7.13 When Automatic Failover Is Performed Due to Hardware Failure Occurring on Primary Storage Where TFO Status Is "Active"




Figure 7.14 When Manual Failover Is Performed Due to Hardware Failure Occuring on Storage Device Where TFO Status Is "Active"





Figure 7.15 Hardware Failure Occurred on Storage Device Where TFO Status Is "Standby"



7.10 When Recovering Storage Cluster Continuous Copy Function

In a state where either the Primary Storage or Secondary Storage is stopped, if the copy is started or stopped, there may be times when the Primary Storage or Secondary Storage sessions are not synchronized.

Take measures after confirming the Advanced Copy status and the Storage Cluster Continuous Copy Session status. The methods for confirming and dealing with the Advanced Copy status and the Storage Cluster Continuous Copy Session status are shown below.

Point

- In the case of the replication management function commands, this operation cannot be executed in relation to storage where the TFO status is "Standby". For this reason, when performing operations on the storage where the TFO status is "Standby", use the commands of the AdvancedCopy Manager Copy Control Module function.

- If the command is executed in a state where the session of the Primary Storage and the Secondary Storage are not synchronized, the swsrp0740 message (error code: 103) may be output. If this message is output, after forcibly stopping the copy with one of the following methods, then take action with "Table 7.3 Corrective Action":
 - For using command
Specify the storage device where the TFO status is "Standby" with the -a option and forcible cancel the copy by the accopy fcancel command. Refer to "Command References" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module* for information on the command.
 - For using Web Console
Refer to "Display/Forcible Suspend/Forcible Cancel Copy Session" in the *ETERNUS SF Web Console Guide*, and forcible cancel the copy.

Advanced Copy status and Standby-Side Continuous Copy Session status

1. Confirm the status of the Standby-Side Continuous Copy Session.
 - When checking by command
Check using the `swsrpstat` command.
The Standby-Side Continuous Copy Session status can be confirmed from the **Standby-Session** column.
 - When confirming from Web Console
Refer to "Create/Display/Change/Delete Replication Group" in the *ETERNUS SF Web Console Guide*, and display the replication group. The Standby-Side Continuous Copy Session status can be confirmed from the **Storage Cluster Standby Copy Session** column.
2. Confirm the status of the Active-Side Continuous Copy Session from the storage device where the TFO status is "Standby".
 - When checking by command
Check using the `accopy query` command. Refer to "Command References" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module* for information on the command.
Specify the storage device where the TFO status is "Standby" with the -a option and execute the command. The Active-Side Continuous Copy Session status can be confirmed from the **TFO Info** column displayed when specifying the -tfoinfo option. Refer to "When Recovering Storage Cluster Continuous Copy Function" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module* for details.
 - When confirming from Web Console
Refer to "Display/Forcible Suspend/Forcible Cancel Copy Session" in the *ETERNUS SF Web Console Guide*, and display the Copy Session. Active-Side Continuous Copy Session status can be confirmed from the **Storage Cluster/Session Status** column.

Countermeasures

Countermeasures are shown below.

Table 7.3 Corrective Action

Advanced Copy Status of Storage Where TFO Status Is "Active"	Continuous Copy Session Status	Action
<ul style="list-style-type: none"> - readying - snap 	When the Standby-Session column for the <code>swsrpstat</code> command, or the Storage Cluster Standby copy session column for Web Console is "not-exist".	<p>Execute the <code>swsrpmake</code> command.</p> <p>When taking measures on Web Console, refer to "Start OPC/QuickOPC/SnapOPC/SnapOPC+" in the <i>ETERNUS SF Web Console Guide</i>, and start the copy.</p>
<ul style="list-style-type: none"> - sync - equivalent 	When the Standby-Session column for the <code>swsrpstat</code> command, or the Storage Cluster Standby copy session column for Web Console is "not-exist".	<ol style="list-style-type: none"> 1. Execute the <code>swsrpmake</code> command to temporarily suspend EC (Suspend). 2. Execute the <code>swsrpstartsync</code> command to resume EC. <p>When taking measures on Web Console, refer to "Suspend EC" and "Start/Resume EC" in the <i>ETERNUS</i></p>

Advanced Copy Status of Storage Where TFO Status Is "Active"	Continuous Copy Session Status	Action
		<i>SF Web Console Guide</i> , and to temporarily suspend/ resume copying.
replicated	When the Standby-Session column for the swsrpstat command, or Storage Cluster Standby copy session column for Web Console is "not-exist".	<ol style="list-style-type: none"> Execute the swsrpstartsync command to resume EC. Execute the swsrpmake command to temporarily suspend EC (Suspend). <p>When taking measures on Web Console, refer to "Start/Resume EC" and "Suspend EC" in the <i>ETERNUS SF Web Console Guide</i>, and to temporarily resume / suspend copying.</p>
failed	When the Standby-Session column for the swsrpstat command, or Storage Cluster Standby copy session column for Web Console is "not-exist".	<ol style="list-style-type: none"> Execute the swsrpcancel command, and stop the copy. Execute the swsrpmake command or the swsrpstartsync command, and start the copy. <p>When taking measures on Web Console, refer to "Stop OPC/QuickOPC/SnapOPC/SnapOPC+" and "Start OPC/QuickOPC/SnapOPC/SnapOPC+", or "Stop EC" and "Start/Resume EC" in the <i>ETERNUS SF Web Console Guide</i>, and to temporarily stop/resume copying.</p>
----	When the TFO Info column for the accopy query command, or the Storage Cluster/Session Status column for Web Console is "unpaired".	<ol style="list-style-type: none"> Specify the storage device where the TFO status is "Standby" with the -a option and forcibly cancel the copy sessions with the accopy fcancel command. Refer to "Command References" in the <i>ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module</i> for information on the command. Execute the swsrpmake command or the swsrpstartsync command, start the copy. <p>When taking measures on Web Console, refer to "Display/Forcible Suspend/Forcible Cancel Copy Session" and "Start OPC/QuickOPC/SnapOPC/SnapOPC +", or "Start/Resume EC" in the <i>ETERNUS SF Web Console Guide</i>, and to temporarily forcible cancel/ resume copying.</p>

7.11 When Changing Configurations/Settings

This section describes the procedure for changing the configuration of the replication operation and the setting of the replication volume information using the Storage Cluster Continuous Copy function.

The procedure for configuration and setting is the same as a non-TFOV.
Refer to ["6.6 Changing Configuration Information"](#).

7.11.1 Deleting Storage Cluster Environment

When deleting the Storage Cluster environment, delete the replication volume settings for the replication operation.

Refer to ["6.6.2 Deleting Source Volume or Replica Volume"](#) for the deletion procedure of the replication volume information.

7.11.2 When Changing Storage Cluster Target Volume

When changing the logical volume number for the replication source/destination volume of the Primary Storage replication source/destination volume, configure the replication source/destination volume after deleting the reproduction information.

The procedure for changing the source/destination volume settings is as follows:

1. Deleting the source volume or replica volume
Delete the replication volume information by the `swsrpdelvol` command.
2. Fetching Device Information on Managed Server
Refer to "Fetching Device Information on Managed Server" in "7.4 Preparations", and update the device information.
3. Setting up the source volumes and replica volumes
Refer to "Configuring Source Volume and Replica Volume" in "7.4 Preparations", and reset the source volume and replica volumes.

7.12 Cascade Copies

This section describes cascade copy when using the Storage Cluster Continuous Copy function.

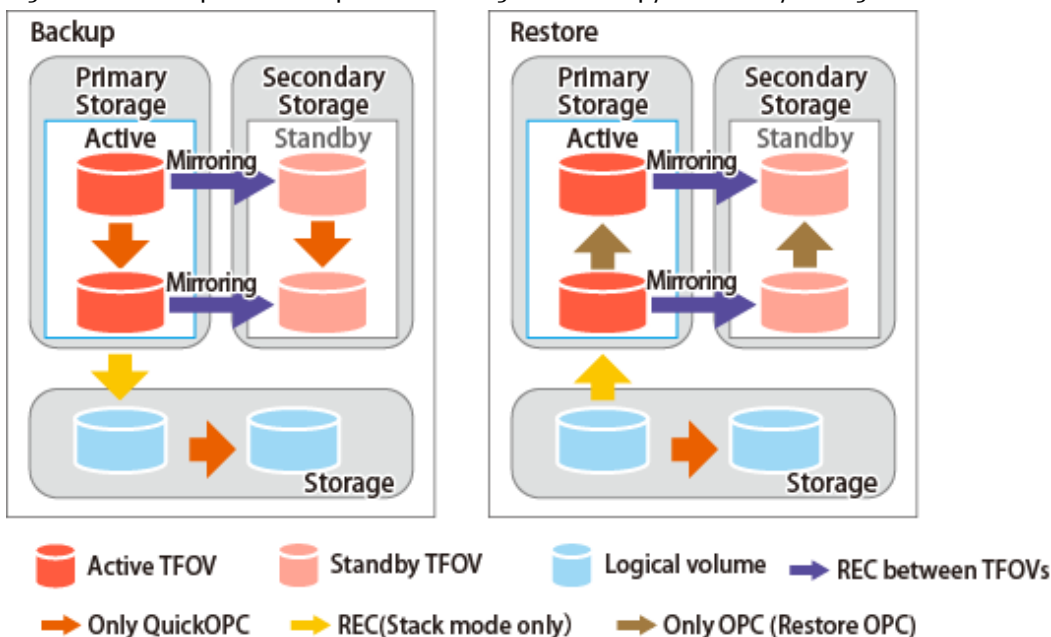
A cascade copy links the copy destination to the processing of another copy as a copy source. Refer to the ETERNUS Disk storage system manuals for details of the cascade copies.

Using cascade copy, it is possible to broaden application scenarios, such as backups that combine duplication to remote areas and sub-center side backups. When using the Storage Cluster Continuous Copy function, configuring cascade copies to different boxes as the backup target are supported. In this way, it can be operated as a secondary backup.

Point

- Use QuickOPC for intra-box copy.
- Use AdvancedCopy Manager Copy Control Module for REC. Refer to the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module* for details. Use Stack mode for the data transfer mode in REC.
- Cascade copy (QuickOPC) can be performed with the REC target as the copy source.


Figure 7.16 Example of Backup/Restore Using Cascade Copy on Primary Storage



Point

- Cascade copy can be performed from the Secondary Storage via REC.
When restoring, put the REC for the storage device where the TFO status is "Standby" into suspend state, and then perform REC.

Figure 7.17 Example of Backup/Restore Using Cascade Copy on Secondary Storage



Chapter 8 Operation Maintenance

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

8.1 Maintenance of Database

AdvancedCopy Manager holds various repositories of management information on the usage of its functions. They are named generically and are referred to as a database in this chapter.

A database is important information. You are recommended to periodically back up a database that is frequently updated, when the space available for repositories becomes insufficient after long-time usage of the database that is different from originally planned, as a safeguard against running out of space. (Note that it differs from other types of databases such as ORACLE etc.)

The following table lists instances when backup of the database is needed.

Table 8.1 Databases on Management Server

Type	When Backup to Be Created	Backup Method (References)	Size
Database	The configuration information has changed as a result of adding a Managed Server or loading of new device information, or changing the Management Server information (server name, IP address, or port number). In addition, information on the backup management operation or the replication management operation has changed.	"8.1.1.3.1 Backing up Database"	Refer to "8.1.1.3.1 Backing up Database".
Registry	The installation of the manager program has been performed on a Management Server (a Management Server transaction has been created in a clustered system), or Management Server information (server name, IP address, or port number) has changed.	"8.1.1.4.1 Backing Up Registry" or "8.1.2 Making Package Backup of Databases"	Refer to "8.1.1.4.1 Backing Up Registry".

Table 8.2 Databases on Managed Server

Type	When Backup to Be Created	Backup Method (References)	Size
Registry	A Managed Server has been added or Management Server information (server name, IP address, or port number) has changed.	"8.1.1.4.1 Backing Up Registry" or "8.1.2 Making Package Backup of Databases"	Refer to "8.1.1.4.1 Backing Up Registry".
Backup management list	A backup transaction has been completed or information on the backup management operation has changed.	"8.1.1.1.1 Backing up Backup Management List" or "8.1.2 Making Package Backup of Databases"	Refer to "8.1.1.1.1 Backing up Backup Management List".
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	Check the size of the pre-processing and post-processing script file for backup management.

Type	When Backup to Be Created	Backup Method (References)	Size
		"8.1.2 Making Package Backup of Databases"	
Replication management list	Information on the replication management operation has changed.	"8.1.1.2.1 Backing up Replication Management List" or "8.1.2 Making Package Backup of Databases"	Refer to "8.1.1.2.1 Backing up Replication Management List".
Pre-processing and post-processing scripts for replication management	A pre-processing or post-processing script has changed.	Saving a database by using the copy command or "8.1.2 Making Package Backup of Databases"	Check the size of the pre-processing and post-processing script file for replication management.
Policy file for management information package backup	The policy file has changed.	Saving a database by using the copy command or "8.1.2 Making Package Backup of Databases"	Check the size of the policy file for management information package backup.



Point

You can create backups of databases either individually or as a whole.

If a Managed Server also operates as a Management Server, the Managed Server contains a database for the Management Server as well as for the Managed Server.

8.1.1 Making Individual Backups of Databases

8.1.1.1 Maintaining Backup Management List

This section provides information on the maintenance of a backup management list. A backup management list is maintained on the server that performs the backup.

8.1.1.1.1 Backing up Backup Management List

Execute the `swstresback` command to back up a backup management list.

Backup Data Size of Backup Management List

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/logicalNodeName/data
```

Point

When you execute the `swstresback` command by using the same directory used before to save data, the file system in which these directories exist should be at least three times greater than the size returned in query above.
This is because the `swstresback` command executes backup in the state holding the data saved before.

8.1.1.1.2 Restoring Backup Management List

Execute the `swstresrst` command to restore a backup management list.

8.1.1.2 Maintaining Replication Management List

This section describes the maintenance of a replication management list.

8.1.1.2.1 Backing up Replication Management List

To back up a replication management list, back up the repository of the Management Server. For information about backing up a repository, refer to "8.1.1.3.1 Backing up Database" of the Management Server.

Point

- Operational 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 `swsrpsetvol` or `swsrpdvol` command.

8.1.1.2.2 Restoring Replication Management List

Perform the following steps to restore a replication management list:

1. Restore the repository to the Management Server. For information about restoring a repository, refer to "8.1.1.3.2 Restoring Database" of the Management Server.
2. On the Managed Server, execute the `swsrprecoverres` command with the `-r` option.

Point

- When above 2 steps are performed, operational information for replication management of a repository is copied to a Managed Server, where it is then available for use.
- When this information for replication management in the repository does not need to be changed, and only the operation information for replication management on the Managed Server is incorrect (for some reason), perform only step 2.

8.1.1.3 Maintaining Database

This section describes the maintenance of a database used by a Management Server.

8.1.1.3.1 Backing up Database

"8.7.2.1 Backup of Common Control and AdvancedCopy Manager Repository" to backup the ETERNUS SF Manager database.

Point

- Use this save method only when the configuration information is changed due to the addition of a Managed Server, or fetching of new device information. It is recommended to save a database regularly (for example, once a day).

- To safeguard against an input-output error in an external file while saved data is being acquired, it is recommended to keep at least two generations of saved data.



Backup Data Size of Database

Calculate the size using the following formula:

$$60 + (\text{number of devices to be managed} \times 0.001) \text{ megabytes}$$

8.1.1.3.2 Restoring Database

"[8.7.2.2 Restore of Common Control and AdvancedCopy Manager Repository](#)" to restore the ETERNUS SF Manager database. Execute the command on the Management Serve.

8.1.1.4 Maintaining Registry

This section describes the maintenance of a registry. Perform this procedure on a Management Server or a Managed Server.

8.1.1.4.1 Backing Up Registry

In a command prompt window, back up the following files using the copy command, etc.

File Names

For Non-clustered Systems

```
/etc/opt/swstorage/swnode.ini
```

For Clustered Systems

```
/etc/opt/swstorage/logicalNodeName/swnode.ini
```

Example

Save it in the /backup directory.

```
# cp /etc/opt/swstorage/swnode.ini /backup
```

Backup Data Size of Registry

Check the size of the file to be backed up.

8.1.1.4.2 Restoring Registry

In a command prompt window, use the copy command or another method to restore the following files.

File Names

For Non-clustered Systems

```
/etc/opt/swstorage/swnode.ini
```

For Clustered Systems

```
/etc/opt/swstorage/logicalNodeName/swnode.ini
```

Example

Restore the following file saved in the /backup directory:

```
# cp /backup/swnode.ini /etc/opt/swstorage
```


8.1.2 Making Package Backup of Databases

This section describes the function of making a backup of management information (e.g., databases) from just one Managed Server or Management Server.

With a package backup of management information, you can make a backup of management information for all Managed Servers and Management Server from just one Managed Server or Management Server. Executing a package backup of management information runs backup commands (making backups of backup management resources and databases) for the databases of other Managed Servers and Management Server, and collects the backup data and management information created by these backup commands.

The server on which you can execute a package backup of management information is called the "package-backup-of-the-management-information execution server". The servers whose data is to be backed up are called the "package-backup-of-the-management-information target servers".

Figure 8.1 Management Server and Managed Server



Example

Assume that the package backup of management information function has been used to back up the management information on a Management Server, Managed Server A, and Managed Server B onto Managed Server B in the diagram above. In this case, Managed Server B is the backup execution server. The Management Server, Managed Server A, and Managed Server B are the backup target servers. That is, Managed Server B operates as both the backup execution server and backup target server.

The following table shows the types of management information that is handled in a package backup of management information.

Table 8.3 Types of Management Information Backed Up by Package Backup of Management Information

Function	Type of Data
Database	- Database
Registry	- Registry
Backup management	- Backup management list, or backup management list and recovery control file - Pre-processing and post-processing scripts - Transaction volume locking specification file for backups - Backup volume locking specification file for backups - Backup-restored volume locking specification file
Replication management	- Pre-processing and post-processing scripts - Copy-source volume locking specification file - Copy-destination volume locking specification file
Package backup of management information	- Policy file



Note

In an environment where the Managed Server with IPv4 address and the Managed Server with IPv6 address coexist, register the Management Server as a "package-backup-of-the-management-information execution server".

8.1.2.1 Preliminaries

Before creating a package backup of management information, you must perform the following:

1. Determine the management information package backup target servers and the types of management information to be backed up.
2. Determine the management information package backup execution server.
3. Create a directory where the management information package backup execution server stores the backup data. The size of the directory must not be less than the total of the following sizes:
 - The required size for a database is the size of the database backup data multiplied by 1.2.
For the backup data size, refer to "Maintaining Database" in the *ETERNUS SF AdvancedCopy Manager Operation Guide* for the operating system running on the Management Server.
 - The required size for the registries is the total of the data sizes of the registry backup data on the management information package backup target servers.
For the backup data size, refer to "Maintaining Database" in the *ETERNUS SF AdvancedCopy Manager Operation Guide* for the operating system running on the management information package backup target server.
 - The required size for backup management is the total of the sizes below for each management information package backup target server:
For the size of each component of backup data, refer to "Maintaining Database" in the *ETERNUS SF AdvancedCopy Manager Operation Guide* for the operating system running on the management information package backup target server.
 - Size of the backup data for the backup management lists (size of the backup data for the backup management lists plus the size of the backup data for the recovery control files) multiplied by 1.2
 - Size of the pre-processing and post-processing scripts for backup management
 - Size of the transaction volume locking specification file for backups for backup management
 - Size of the backup volume locking specification file for backups for backup management
 - Size of the backup-restored volume locking specification file for backup management
 - The required size for replication management is the total of the following sizes for each management information package backup target server:
 - Size of the pre-processing and post-processing scripts for replication management
 - Size of the copy-source volume locking specification file for replication management
 - Size of the copy-destination volume locking specification file for replication management
4. Check the size of the working directory for the management information package backup target servers (or the shared disk for shared data for AdvancedCopy Manager in a clustered system). The sizes described below are needed only when a package backup of management information is being created. Use the largest value of the sizes of the databases to be backed up from the management information package backup target servers. The required sizes are as follows:
 - The required size for the databases is the size of the database backup data multiplied by 2.2.
For the backup data size, refer to "Maintaining Database" in the *ETERNUS SF AdvancedCopy Manager Operation Guide* for the operating system running on the Management Server.
 - The required size for backup management is the size of the backup data for the backup management lists (size of the backup data for the backup management lists plus the size of the backup data for the recovery control files) multiplied by 2.2.
For the backup data size, refer to "Maintaining Database" in the *ETERNUS SF AdvancedCopy Manager Operation Guide* for the operating system running on the management information package backup target server.
 - The working directory does not require any space for registry, replication management, and policy file backups.

5. Execute the `stgmgrinfo set` command on the management information package backup execution server (determined in step 2).
6. Create a policy file on the management information package backup execution server.

Table 8.4 Management Information Package Backup Execution Server and Policy File Paths

Management Information Package Backup Execution Server		Policy File Path
OS	Operation Type	
Windows	Non-cluster	<Environment directory>\etc\stgmgrinfo_policy.ini
	Cluster	<AdvancedCopy-Manager-shared-data>\etc\opt\swstorage\etc\stgmgrinfo_policy.ini
Solaris, Linux, HP-UX, AIX	Non-cluster	/etc/opt/FJSVswstf/stgmgrinfo_policy.ini
	Cluster	/etc/opt/FJSVswstf/logical/NodeName/stgmgrinfo_policy.ini

Note

- Do not specify more than one Managed Server as a management information package backup execution server.
- When this server is used for executing transactions on the Managed Server in a clustered system, create a directory for storing backup data on a shared disk.
- For databases and backup management, backup commands are executed on the corresponding management information package backup target servers. In addition, for database backups and backup management (ie, backup management lists) backups, a temporary capacity that is approximately 2.2 times greater than the total size of the backups is required to save the backup data in an archive log file.

Table 8.5 Policy File Format

Section Name	Key Name	Description(*)
Server name	BkupData	Specifies whether to make a backup of the backup management. Yes: Make a backup of the backup management. No: Do not make a backup of the backup management.
	ReplData	Specifies whether to make a backup of the replication management. Yes: Make a backup of the replication management. No: Do not make a backup of the replication management.

Specifiable only when the applicable function is operating on the server (Management Server or Managed Server).

*: If this key is not specified or if "Yes" is not in the form "Y" (uppercase) + "es" (lowercase), then "No" is set.

Example

Set up a policy file as follows:

- The Management Server named ACMMGR creates backups for the backup management.
- The Managed Server named ACMAGT1 creates backups for backup management.
- The Managed Server named ACMAGT2 creates backups for replication management.

```
[ ACMMGR ]
BkupData=Yes
[ ACMAGT1 ]
```

```
BkupData=Yes
[ACMAGT2]
ReplData=Yes
```

Note

- Create a policy file only if it is 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 `stgmgrinfobkup` command on the management information package backup execution server stores the backup data in the directory under the following structure.

After the `stgmgrinfobkup` command is executed, refer to "8.7.2.1 Backup of Common Control and AdvancedCopy Manager Repository" to back up the ETERNUS SF Manager database.

Table 8.6 Directory Structure for Execution of Management Information Package Backup and Data Types

Directory		Data Type		
*1	Server name	RegiData		- Registry
		BkupData	FJSVswsts	- Backup management list, or backup management list and recovery control file
			Scripts	- Pre-processing and post-processing scripts for backup management
			data	- Transaction volume locking specification file for backups, for backup management - Backup volume locking specification file for backups, for backup management - Backup-restored volume locking specification file for backups, for backup management
	ReplData	Scripts	- Pre-processing and post-processing scripts for replication management	
		data	- Copy-source volume locking specification file for replication management - Copy-destination volume locking specification file for replication management	
	MgrInfoData		- Policy file for management information package backup	

*1 is the backup data storage directory specified by the `stgmgrinfoset` command.

Example

When *1 is `/ACMBkupData` and the server name is `ACMAGT1`, the pre-processing and post-processing scripts for backup management are put in the following directory:

```
/ACMBkupData/ACMAGT1/BkupData/Scripts
```

Initial Operations

The policy file backup operations are described below. Backup of registries which are not included in the backup targets when a policy file is used to perform backup are also detailed. The initial operations are as follows:

1. Execute the [stgmgrinfobkup](#) command on the management information package backup execution server to make backups of the registries on all of the Managed Server.
2. Execute the [stgmgrinfobkup](#) command on the management information package backup execution server to make a backup of the policy file.
3. Refer to "[8.7.2.1 Backup of Common Control and AdvancedCopy Manager Repository](#)" to back up the ETERNUS SF Manager database.

Normal-use Operations

Execute the [stgmgrinfobkup](#) command on the management information package backup execution server when you need to make backups of the target data.

Afterwards, refer to "[8.7.2.1 Backup of Common Control and AdvancedCopy Manager Repository](#)" to back up the ETERNUS SF Manager database.

For the timing of the backups, refer to "Maintaining Database" in the *ETERNUS SF AdvancedCopy Manager Operation Guide* for the applicable operating system.

Operations When Server Is Added or Deleted

Operations When Server Added

1. On the management information package backup execution server, set up a policy for the added Managed Server in the policy file.
2. Execute the [stgmgrinfobkup](#) command on the management information package backup execution server to make a backup of the registry of the added Managed Server.
3. Execute the [stgmgrinfobkup](#) command on the management information package backup execution server to make a backup of the policy file.
4. Refer to "[8.7.2.1 Backup of Common Control and AdvancedCopy Manager Repository](#)" to back up the ETERNUS SF Manager database.

Operations When Server Deleted

1. On the management information package backup execution server, delete the information about the Managed Server that was deleted from the policy file.
2. In the backup data directory on the management information package backup execution server, delete the directory containing the name of the deleted Managed Server.
3. Execute the [stgmgrinfobkup](#) command on the management information package backup execution server to make a backup of the policy file.
4. Refer to "[8.7.2.1 Backup of Common Control and AdvancedCopy Manager Repository](#)" to back up the ETERNUS SF Manager database.

Operations When Server Name Is Changed

1. On the management information package backup execution server, delete the information for the Managed Server from the policy file.
2. In the backup data directory on the management information package backup execution server, delete the directory containing the name of the Managed Server.
3. On the management information package backup execution server, set up a policy for the renamed Managed Server in the policy file.
4. Execute the [stgmgrinfobkup](#) command on the management information package backup execution server to make a backup of the registry of the renamed Managed Server.
5. Execute the [stgmgrinfobkup](#) command on the management information package backup execution server to make a backup of the policy file.

6. Refer to "[8.7.2.1 Backup of Common Control and AdvancedCopy Manager Repository](#)" to back up the ETERNUS SF Manager database.

8.1.2.3 Restoration Operations

This section describes how to use the backup data created by the `stgmgrinfobkup` command and "[8.7.2.1 Backup of Common Control and AdvancedCopy Manager Repository](#)" to execute restoration.

Perform the operations below as a root user or with an account that belongs to the Administrators group.

Restoring Registry

1. Execute the `stgmgrinfo` command on the management information package backup execution server to determine which directory contains the backup data.
2. On the management information package backup execution server, copy the registry backup to the following directory:

For Windows:

```
<Directory-determined-in-step1>\<server-whose-registry-to-be-restored>\RegiData
```

For Solaris, Linux, HP-UX or AIX:

```
<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 Registry" in the *ETERNUS SF AdvancedCopy Manager Operation Guide* for the operating system running on the target server.

Restoring Database

Refer to "[8.7.2.2 Restore of Common Control and AdvancedCopy Manager Repository](#)" to restore the ETERNUS SF Manager database on the Management Server.

Restoring Replication Management List

1. Perform the procedure of "[Restoring Database](#)" to restore a database.

If you have performed it, this step is not needed.

2. On the target server, execute the `swsrecoverres` command with the `-r` option.

For details on the command, refer to "`swsrecoverres (Resource Adjustment Command)`" in the *ETERNUS SF AdvancedCopy Manager Operation Guide* for the operating system running on the target server.

Restoring Backup Management List, or Backup Management List and Recovery Control File

1. Execute the `stgmgrinfo` command on the management information package backup execution server to determine which directory contains the backup data.
2. On this server, copy the backup of the backup management list or the backups of the backup management list and the recovery control file to the following directory:

For Windows:

```
<Directory-determined-in-step1>\<server-whose-list/file-to-be-restored>\BkupData\FJSVswsts
```

For Solaris, Linux, HP-UX or AIX:

```
<Directory-determined-in-step1>/<server-whose-list/file-to-be-restored>/BkupData/FJSVswsts
```

3. Transfer the backup of the backup management list or the backups of the backup management list and the recovery control file copied in step 2 to the target server in binary mode.

4. Expand the backup of the backup management list or the backups of the backup management list and recovery control file transferred in step 3 on the target server.
 - On a Windows version server:

Expand the backup of the backup management list using the stgarcext (Backup Data Expansion Command) in the *ETERNUS SF AdvancedCopy Manager Operation Guide (for Windows)*.
 - On a Solaris/Linux/HP-UX/AIX version server:

Expand the backup of the backup management list or the backups of the backup management list and recovery control file by using the tar command supported by the operating system.
5. Restore the backup management list, or the backup management list and recovery control file on the target server. For details on restoring, refer to "Restoring Backup Management List and Recovery Control File" in the *ETERNUS SF AdvancedCopy Manager Operation Guide* for the operating system running on the target server. The Backup-Directory specified in the resource restoration command is the directory expanded in step 4.

Restoring Backup Management Pre-processing/Post-processing Scripts

1. Execute the [stgmgrinfo](#) command on the management information package backup execution server to determine which directory contains the backup data.
2. On this server, copy the backups of the backup management pre-processing/post-processing scripts for backup management to the following directory:

For Windows:

```
<Directory-determined-in-step1>\<server-whose-scripts-to-be-restored>\BkupData\Scripts
```

For Solaris, Linux, HP-UX or AIX:

```
<Directory-determined-in-step1>/<server-whose-scripts-to-be-restored>/BkupData/Scripts
```

3. Transfer the backups of the backup management pre-processing/post-processing scripts copied in step 2 to the target server in binary mode.
4. Transfer the backups of the backup management pre-processing/post-processing scripts copied in step 3 to the directory in which the scripts are 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 of Backup and Restoration" in the *ETERNUS SF AdvancedCopy Manager Operation 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 [stgmgrinfo](#) command on the management information package backup execution server to determine which directory contains the backup data.
2. On this server, copy the backups of the locking specification files for backup management to the following directory:

For Windows:

```
<Directory-determined-in-step1>\<server-whose-files-to-be-restored>\BkupData\data
```

For Solaris, Linux, HP-UX or AIX:

```
<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 Operation Guide (for Windows)*.

Restoring Pre-processing and Post-processing Scripts for Replication Management

1. Execute the `stgmgrinfo` command on the management information package backup execution server to determine which directory contains the backup data.
2. On this server, copy the backups of the pre-processing and post-processing scripts for replication management to the following directory:

For Windows:

```
<Directory-determined-in-step1>\<server-whose-scripts-to-be-restored>\ReplData\Scripts
```

For Solaris, Linux, HP-UX or AIX:

```
<Directory-determined-in-step1>/<server-whose-scripts-to-be-restored>/ReplData/Scripts
```

3. Transfer the backups of the pre-processing and post-processing scripts for replication management copied in step 2 to the target server in binary mode.
4. Copy the backups of the pre-processing and post-processing scripts for replication management transferred in step 3 to the directory in which the scripts are 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 "[Appendix C Pre-processing and Post-processing of Replication](#)" in the *ETERNUS SF AdvancedCopy Manager Operation Guide* for the operating system running on the target server.
5. On a Solaris/Linux/HP-UX/AIX 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 `stgmgrinfo` command on the management information package backup execution server to determine which directory contains the backup data.
2. On this server, copy the backups of the locking specification files for replication management to the following directory:

For Windows:

```
<Directory-determined-in-step1>\<server-whose-files-to-be-restored>\ReplData\data
```

For Solaris, Linux, HP-UX or AIX:

```
<Directory-determined-in-step1>/<server-whose-files-to-be-restored>/ReplData/data
```

3. Transfer the backups of the locking specification files for replication management copied in step 2 to the target server in binary mode.
4. Copy the backups of the locking specification files for replication management transferred in step 3 to the target server. For the directory in which the locking specification files for replication management are used (that is, the directory to which the files are to be copied), refer to "Pre-processing and Post-processing of Replication" in the *ETERNUS SF AdvancedCopy Manager Operation Guide (for Windows)*.

Restoring Policy File for Management Information Package Backup

1. Execute the `stgmgrinfo` command on the management information package backup execution server to determine which directory contains the backup data.
2. On this server, copy the policy file located in the following directory to the policy file for the management information package backup command:

For Windows:

```
<Directory-determined-in-step1>\<server-whose-file-to-be-restored>\MgrInfoData
```

For Solaris, Linux, HP-UX or AIX:

```
<Directory-determined-in-step1>/<server-whose-file-to-be-restored>/MgrInfoData
```

For information on the directory for the policy file for the management information package backup command, refer to step 6 in "8.1.2.1 Preliminaries".

8.2 Troubleshooting

8.2.1 Collecting Troubleshooting Information

8.2.1.1 Collecting Error Survey Information

Collecting Error Survey Information on Management Server

The collecting error survey information on the Management Server can be performed using either the command or Web Console.

When using the command, login to the Management Server and execute the `esfsnap` command.

When using Web Console, refer to "Collecting Troubleshooting Information" in the *ETERNUS SF Web Console Guide*.

Send the collected error survey information to Fujitsu Technical Support.



- If an error occurs during installing ETERNUS SF, the `esfsnap` command may not exist. In this case, collect the following data.
 - System log
 - The files under the following directories.
 - `/var/opt/FJSVswtrc/trc`
 - `/var/opt/FJSVswstf`
 - `/etc/opt/swstorage`
 - In a clustered system, execute the command to collect troubleshooting information on all the nodes making up the clustered system.
-

Collecting Error Survey Information on Managed Server

Login to the Managed Server and execute the `acmcapture` command on each server.

Send the collected error survey information to Fujitsu Technical Support.



- If a fault occurs during node degeneracy in a load sharing degenerate function on the Symfoware Server, collect files under the directory specified by the `-altdir` option.
 - In a clustered system, execute the command to collect troubleshooting information on all the nodes making up the clustered system.
-

8.3 Troubleshooting: Hardware Error During Backup

This section describes the actions to perform if a hardware error, etc. occurs during a backup using AdvancedCopy Manager.

8.3.1 Overview

If a hardware error, etc. occurs during backup, the following symptoms may be displayed:

- The `swstbackstat` command displays "failed" or "?????" in the **Status** column.
- The `swstsyncstat` command displays "failed" or "?????" in the **Status** column.
- The `swstreststat` command displays "failed" or "?????" in the **Status** column.
- The `swsthistdisp` command displays "failed" or "?????" in the **Status** column.
- The `swstexecstat` command displays "sync(failed)", "snapshot(failed)", "restore(failed)", or "?????" in the **Status** column.
- The `swstrkstat` command displays "failed" or "?????" in the **Status** column.
- The message `swst0634`, `swst0719`, or `swst0740` is displayed during command execution.

If the execution status display command and history information display command display "failed" in the **Status** column, a hardware error has occurred during physical copying, and the copy processing has failed.

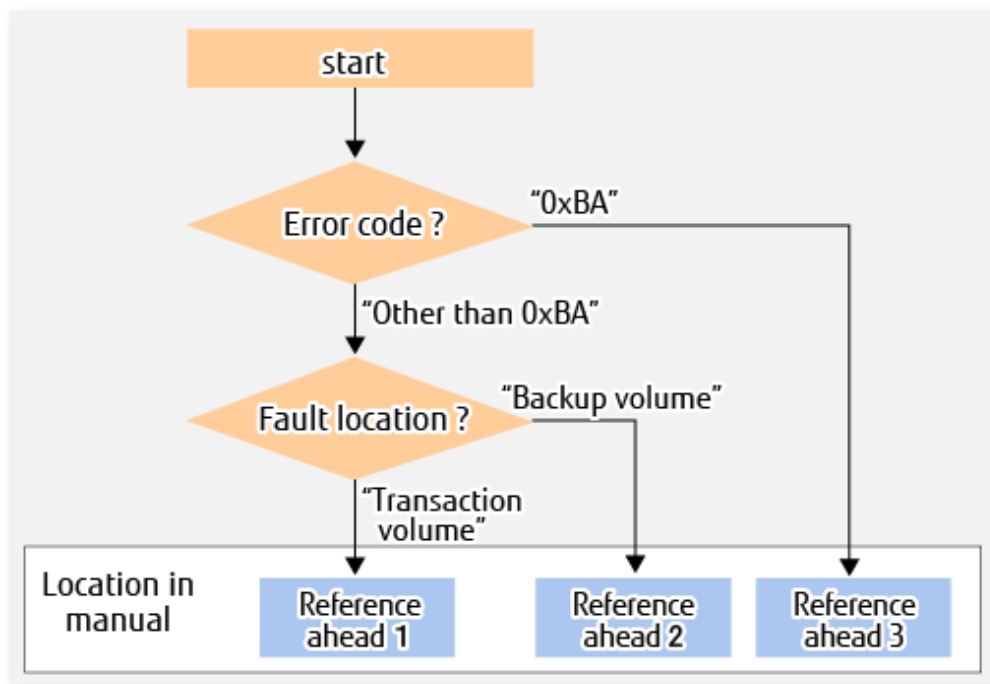
If the message `swst0634`, `swst0719`, or `swst0740` is displayed during command execution and the same error occurs even if the command is re-executed, a hardware error, etc. has occurred, causing Advanced Copy to fail.

If any of these messages, statuses or symptoms is generated, a hardware error may have occurred in the transaction volume or backup volume. Check the hardware status, and if a hardware error is found, fix the cause of the error, then re-execute the processing.

8.3.2 Troubleshooting

The following figure shows the flow of troubleshooting tasks if a hardware or similar fault occurs.

Figure 8.2 Troubleshooting Flow (for Faults During Backup)



Reference ahead 1: [8.3.2.1 Hardware Error on Transaction Volume](#)

Reference ahead 2: [8.3.2.2 Hardware Error on Backup Volume](#)

Reference ahead 3: [8.3.2.3 Troubleshooting If Bad Sector Occurred in Transaction Volume](#)

Note

- Refer to "8.3.1 Overview" for details of the **Status** column and "Fault location".
If the **Status** column is "?????", check if the copy processing is in the error suspend state ("failed") or the hardware suspend state ("halt") using ETERNUS Web GUI.
If the copy processing is in either of these states, take the action indicated in the above troubleshooting flow.
In other cases, take the action checked in the following points.
 - If a device is not accessible:
Check if the device exists.
 - If there is anything unusual with Managed Server, switches, etc.:
Contact Fujitsu Technical Support.
- Use ETERNUS Web GUI to check the copy process error codes. Use the following method to check the error codes:
 - On the [Display status] menu, click [Advanced Copy status display] in the status display.
 - At "Session status", click the "Number of active sessions" link for the relevant copy type.
 - Refer to the value in the "Error code" column of the relevant copy process.

The following table shows the meanings of the error codes.

Table 8.7 Meanings of Error Codes

Error Code	Meaning
0xBA	If a) or b) below applies, a bad sector was created in the transaction volume. <ul style="list-style-type: none">a. QuickOPC has not yet performed physical copying and tracking is in progressb. EC is in the suspend status (replication established status) <p>Note: If a bad sector is created in a transaction volume when a) or b) applies, the ETERNUS Disk storage system automatically changes the copy processing to the error suspend state. This prevents a restart of QuickOPC or EC resume and prevents the backup volume from being overwritten with invalid transaction volume data.</p>
Other than 0xBA	An error other than the above occurred.

8.3.2.1 Hardware Error on Transaction Volume

When a hardware fault occurs in the transaction volume, perform the repair work according to the following procedures.

Note

- The backup history information for which a copy failed cannot be used to perform a restoration.
 - If there is no normal (that is, **Status** is "succeeded") backup history information, restoration is not possible.
1. If an error occurred during physical copying, cancel the processing in which the error occurred, with one of the following commands:
 - If an error occurred during backup (OPC) physical copy: `swsthistdel`
 - If an error occurred during synchronous processing (EC): `swstcancelsync`
 - If an error occurred during restore (OPC) physical copy: `swstcancelrest`
 - If an error occurred during tracking processing (QuickOPC): `swstcanceltrk`

2. Execute the execution status display commands ([swstbackstat](#), [swstsyncstat](#), [swstreststat](#), [swsttrkstat](#) and [swsthistdisp](#)) to verify that no other errors have occurred.
3. Fix the hardware error in the transaction volume.
4. Execute the [swstrestore](#) command to restore the transaction volume.

8.3.2.2 Hardware Error on Backup Volume

When a hardware fault occurs in backup volume, perform repair work according to the following procedures.

1. Cancel the processing in which the error occurred, with one of the following commands:
 - If an error occurred during backup (OPC) physical copy: [swsthistdel](#)
 - If an error occurred during synchronous processing (EC): [swstcancelsync](#)
 - If an error occurred during restoration (OPC) physical copy: [swstcancelrest](#)
 - If an error occurred during tracking processing (QuickOPC): [swstcanceltrk](#)If the processing cannot be cancelled using one of the above commands, use ETERNUS Web GUI to cancel it.
2. Execute the [swstsrsemitch](#) command.
3. Cancel the processing in which the error occurred. At this point, the execution status display commands ([swstbackstat](#), [swstreststat](#) and [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 cancelled, so no response is required.
 - If an error occurred during backup (OPC) physical copy: [swsthistdel](#)
 - If an error occurred during restoration (OPC) physical copy: [swstcancelrest](#)
4. Execute the execution status display commands ([swstbackstat](#), [swstsyncstat](#), [swstreststat](#), [swsttrkstat](#) and [swsthistdisp](#)) to verify that no other errors have occurred.
5. Execute the [swstdevinfofet](#) command to delete the backup volume in which the error occurred.
6. Execute the [swstdevinfofet](#) command to register a new backup volume. If the backup volume on which the error occurred is to be repaired and be reused, execute the [swstdevinfofet](#) command after executing the [stgxfwcmsetdev](#) command on the Management Server.
7. Re-execute the processing in which the error occurred.

8.3.2.3 Troubleshooting If Bad Sector Occurred in Transaction Volume

If a bad sector occurred in the transaction volume, use the following procedure to restore the transaction volume:

1. Cancel the copy processing for which the bad sector occurred.
 - If the bad sector occurred during the suspend state of synchronous processing (EC): [swstcancelsync](#)
 - If the bad sector occurred during QuickOPC tracking: [swstcanceltrk](#)
2. Execute the [swstbackstat](#), [swstsyncstat](#), [swstreststat](#), [swsttrkstat](#) and [swsthistdisp](#) commands to check for other errors.
3. Restoration is performed by overwriting the area containing the bad sector. Select the appropriate method, in accordance with the usage or use **status** of the transaction volume, from the methods below.
 - Restoration method 1:
If the area can be reconstructed from high-level software (file system, DBMS, or similar), reconstruct the area.
 - Restoration method 2:
If the area containing the bad sector is an area that is not being used, such as an unused area or a temporary area, use a system command (for example, the UNIX dd command or the Windows format command) to write to the area.

- Restoration method 3:

Execute the [swstore](#) command to restore the transaction volume from a backup volume. (If the backup volume for the copy processing for which the bad sector occurred contains backup history information, restoration is also possible from that backup volume.)



The transaction volume cannot be restored by this method, if there is no backup history information.

8.4 Troubleshooting: Hardware or Other Error During Replication

This section describes the corrective action to take when a hardware or other type of error occurs during AdvancedCopy Manager replication.

8.4.1 Overview

If a hardware error, etc. occurs during the replication, the following symptoms are displayed:

- The [swsrpstat](#) command displays "failed" or "?????" in the **Status** column.
- The message swsrp2606 is output during command execution.
- The message swsrp0634, swsrp0719, or swsrp0740 is displayed during command execution.

If the operation status display command displays "failed" in the **Status** column and swsrp2606 is displayed during command execution, a hardware error has occurred during physical copying, and the copy processing has failed.

If the message swsrp0634, swsrp0719, or swsrp0740 is displayed during command execution, and the same error occurs even if the command is re-executed, a hardware error, etc. has occurred, causing Advanced Copy to fail.

If a hardware error occurs during remote copying, the following symptoms may be displayed:

- The [swsrpstat](#) command displays "halt" in the **Status** column.
- The message swsrp2648 is displayed during command execution.


If the operation status display command displays "halt" in the **Status** column, and the message swsrp2648 is displayed during command execution, a hardware error has occurred during physical copying, and the remote copy processing may have failed.

If any of these messages, statuses or symptoms is generated, a hardware error or path close may have occurred in the source or replica volume. Check the hardware status, and if a hardware error is found, fix the cause of the error and re-execute the command.

8.4.2 Troubleshooting

The following figure shows the troubleshooting flow when a hardware or similar fault occurs.

Figure 8.3 Troubleshooting Flow (When Fault Occurs During Replication)



Reference ahead 1: [8.4.2.1 Hardware Error on Replication Volume](#)

Reference ahead 2: [7.10 When Recovering Storage Cluster Continuous Copy Function](#)

Reference ahead 3: [8.4.2.2 Troubleshooting If Bad Sector Occurred in Copy Source Volume](#)

Reference ahead 4: [8.4.2.3 Troubleshooting When Lack of Free Physical Space Has Occurred in Copy Destination Volume](#)

Reference ahead 5: [8.4.2.4 Error \(halt\) on Remote Copy Processing](#)

Note

- Refer to ["8.4.1 Overview"](#) for details of the **Status** column and "Fault location".

If the **Status** column is "?????", check if the copy processing is in the error suspend state ("failed") or the hardware suspend state ("halt") using ETERNUS Web GUI.

If the copy processing is in either of these states, take the action indicated in the above troubleshooting flow.

In other cases, take the action checked in the following points.

- If a device is not accessible:
 - Check if the device exists.
- If there is anything unusual with Managed Server, switches, etc.:
 - Contact Fujitsu Technical Support.
- Use ETERNUS Web GUI to check the error codes. Use the following two methods to check.
 - Checking with the `swsrpstat` command
 - Execute the command with the `-O` option.
 - Checking with ETERNUS Web GUI
 1. On the [Display status] menu, click [Advanced Copy status display] in the status display.
 2. At "Session status", click the "Number of active sessions" link for the relevant copy type.
 3. Refer to the value in the "Error code" column of the relevant copy process.

The following table shows the meanings of the error codes.

Table 8.8 Meanings of Error Codes

Error Code	Meaning
0xB2	The Storage Cluster Continuous Copy Sessions of the Primary Storage and the Secondary Storage were not able to synchronize, and Advanced Copy was not able to continue.
0xBA	If a) or b) below applies, a bad sector was created in the transaction volume. <ul style="list-style-type: none"> a. QuickOPC has not yet performed physical copying and tracking is in progress b. EC/REC is in the suspend status (replication established status) <p>Note: If a bad sector is created in a transaction volume when a) or b) applies, the ETERNUS Disk storage system automatically changes the copy processing to the error suspend state. This prevents a restart of QuickOPC or EC/REC resume and prevents the copy destination volume from being overwritten with invalid copy source volume data.</p>
0x1E, 0x2E, 0xBB	A lack of free space has occurred in the copy destination volume.
Other than above	An error other than the above occurred.

8.4.2.1 Hardware Error on Replication Volume

When a hardware error occurs in a duplicate volume, perform the repair work on the error according to the following procedures. If the Storage Cluster Continuous Copy function is used, refer to "7.9.1 Recovery from Hardware Failure".

1. Execute the [swsrpcancel](#) command to cancel the processing in which the error occurred. If the processing cannot be cancelled from the operation server when inter-server replication is performed, cancel it from a non-operational server. If the processing cannot be cancelled by using the command, use ETERNUS Web GUI to cancel it.
2. Execute the [swsrprecoverres](#) command.
3. Execute the [swsrpstat](#) command to verify that no other errors have occurred.
4. Execute the [swsrpdelvol](#) command to delete the replication volume in which the error occurred.
5. Execute the [swsrpsetvol](#) command to register a new replication volume. If the replication volume on which the error occurred is to be repaired, execute the [swsrpsetvol](#) command after executing the [stgxfwcmsetdev](#) command on the Management Server.
6. Re-execute the processing in which the error occurred.

8.4.2.2 Troubleshooting If Bad Sector Occurred in Copy Source Volume

If a bad sector occurred in the copy source volume, use the following procedure to restore the copy source volume: If the Storage Cluster Continuous Copy function is used, refer to "7.9.1 Recovery from Hardware Failure".

1. Execute the [swsrpcancel](#) command to cancel processing for which the error occurred. If inter-server replication was being performed and cancellation is not possible from the active server, cancel processing from the inactive server. If processing cannot be cancelled using commands, use ETERNUS Web GUI to cancel it.
2. Execute the [swsrpstat](#) command to check for other errors.
3. Restoration is performed by overwriting the area containing the bad sector. Select the appropriate method, in accordance with the usage or use status of the copy source volume, from the methods below.
 - Restoration method 1:
If the area can be reconstructed from high-level software (file system, DBMS, or similar), reconstruct the area.

- Restoration method 2:
If the area containing the bad sector is an area that is not being used, such as an unused area or a temporary area, use a system command (for example, the UNIX dd command or the Windows format command) to write to the area.
- Restoration method 3:
Execute the [swsrpmake](#) command to restore the data from the copy destination volume. (Restoration is also possible from the copy destination volume of the copy process for which the bad sector occurred.)

8.4.2.3 Troubleshooting When Lack of Free Physical Space Has Occurred in Copy Destination Volume

Use the following procedure to recover the copy destination volume:

1. Execute the [swsrpcancel](#) command to cancel the copy session in which the error occurred.
If the cancellation is not possible from the operation server when an inter-server replication is being performed, cancel it from a non-operation server.
If the copy session cannot be cancelled using command, use ETERNUS Web GUI to cancel it.
2. Check the status of the copy destination volume and initialize it.
 - If the copy destination volume is TPV
Use Storage Cruiser or ETERNUS Web GUI to check the status of the copy destination volume and initialize it.
For the operation procedure when using Storage Cruiser, refer to "Display Volume" and "Delete Reserved Volume or Forcible Delete/Format Volume" in the *ETERNUS SF Web Console Guide*.
 - If the copy destination volume is FTV
Use Storage Cruiser or ETERNUS Web GUI to check the status of the copy destination volume and initialize it.
For the operation procedure when using Storage Cruiser, refer to "Display FTV" and "Format FTV" in the *ETERNUS SF Web Console Guide*.
 - If the copy destination volume is SDV
Use Storage Cruiser, the [swstsdv](#) command, or ETERNUS Web GUI to check the status of the copy destination volume and initialize it.
For the operation procedure when using Storage Cruiser, refer to "Display Volume" and "Delete Reserved Volume or Forcible Delete/Format Volume" in the *ETERNUS SF Web Console Guide*.
The operation procedure when using the [swstsdv](#) command is as follows:
 1. Execute the command with the "stat" subcommand and check the status of the SDV.
 2. Execute the command with the "init" subcommand and initialize the SDV.
3. Recreate partitions (slices) in the copy destination volume.

The following factors may have caused a capacity shortage of the physical space in the copy destination volume:

- a. The estimate of the required physical space for the copy destination volume is not adequate.
- b. Although the estimate of the required physical space for the copy destination volume is adequate, because large amounts of updates have been performed in the copy destination volume in which the copy session does not exist, the physical space of the copy destination volume is being wasted.

When corresponding to the above "a", re-estimate the physical space required for the copy destination volume and consider the disk expansion.

- If the copy destination volume is TPV
Use Storage Cruiser or ETERNUS Web GUI to check the status of the Thin Provisioning pool and expand the capacity of the Thin Provisioning pool.
For the operation procedure when using Storage Cruiser, refer to "Display Thin Provisioning Pool" and "Expand Capacity of/Format/Change Threshold Value of/Delete Thin Provisioning Pool" in the *ETERNUS SF Web Console Guide*.

- If the copy destination volume is FTV

Use Storage Cruiser or ETERNUS Web GUI to check the status of the Tier pool and expand the sub-pool capacity of the Tier pool. For the operation procedure when using Storage Cruiser, refer to "Display Tier Pool" and "Expand Capacity of Sub-Pool in Tier Pool" in the *ETERNUS SF Web Console Guide*.

- If the copy destination volume is SDV

Use the `swstsdv` command or ETERNUS Web GUI to check the status of the SDP and expand the SDP capacity.

Creating an exclusive volume named Snap Data Pool Volume (SDPV) enables the SDP and the created SDPV is automatically incorporated in the SDP. Creating an SDPV of the physical capacity that is assigned to a copy destination volume expands the SDP capacity.

The operation procedure when using the `swstsdv` command is as follows:

1. Execute the command with the "poolstat" subcommand and check the status of the SDP.
2. Create the SDPV with ETERNUS Web GUI.

8.4.2.4 Error (halt) on Remote Copy Processing

The REC restart (Resume) method varies, depending on the halt status.

Execute the `swsrpstat` command with the -H option specified to check the halt status, and then implement the relevant countermeasure.

- For "halt(use-disk-buffer)" or "halt(use-buffer)"

This status means that data is saved to the REC Disk buffer or REC buffer because data cannot be transferred due to a path closure (halt).

In order to restart REC, perform path recovery before a space shortage occurs for the REC Disk buffer or REC buffer.

After recovery, the ETERNUS Disk storage system restarts REC automatically.

If a space shortage has already occurred for the REC Disk buffer or REC buffer, the "halt(sync) or halt (equivalent)" status shown below occurs. Implement the countermeasures for that status.

- For "halt(sync) or halt(equivalent)"

This status means that data transfer processing was discontinued due to a path closure (halt).

The REC restart method differs for different REC Recovery modes.

For Automatic Recovery Mode

1. Remove the cause that made all paths close (halt).
2. ETERNUS Disk storage system automatically restarts (Resume) REC.

For Manual Recovery Mode

1. Remove the cause that made all paths close (halt).
2. Execute the `swsrpmake` command to forcibly suspend the REC that is in the halt status.

```
[For volume units]
swsrpmake -j < replication source volume name > <replication destination
volume name >
```

```
[For group units]
swsrpmake -j -Xgroup <group name>
```

3. Execute the `swsrpstartsync` command to restart (Resume) the REC. The -t option must be specified if REC is being restarted after a forcible suspend.

```
[For volume units]
swsrpstartsync -t <replication source volume name > <replication destination
volume name >
```

[For group units]
swsrpstartsync -t -Xgroup <group name>

8.5 Changing Operating Environment

Server information batch change operations preserve the consistency of server information in systems managed by AdvancedCopy Manager. All Managed Server services and transactions in the systems that include a Management Server and Management Server transactions must be running when these changes are made.



Note

- If the `stgxfwcmmodsrv` command is used to change IP addresses, or similar.

Do not execute backup management commands or replication management commands at the Management Server that executes the `stgxfwcmmodsrv` command, nor at any of the Managed Servers managed by Management Server transactions, either before or during execution of `stgxfwcmmodsrv` command.

- If the `stgxfwcmmodsrv` command terminates abnormally, implement recovery in accordance with the action indicated in the error message. If backup management or replication management commands are accidentally executed either before or during execution of the server information change commands.
 - If the backup management and replication management commands terminate abnormally, wait until the `stgxfwcmmodsrv` command terminates normally, then re-execute the backup management and replication management commands.
 - Change to the value of 800 characters or less before executing the `stgxfwcmmodsrv` command when the value of the environment variable `PATH` is more than 800 characters.
-

8.5.1 Changing IP Address of Management Server

Change (including addition and deletion) the IP address of a Management Server as follows:

If only a physical IP address needs to be changed and the Management Server uses a logical IP address, then perform only steps 3 and 4. Otherwise, perform all steps.

1. On the Management Server, change the system IP address, and restart the Management Server
2. On the Management Server, create a server information change instruction file, and then execute the `stgxfwcmmodsrv` command with the `-f` option specified to change the IP address. When using the replication management function, execute the `swsrprecoverres` command with the `-r` option and the `-h` option on the Management Server. Specify the server name of the Managed Server running the replication operation with the `-h` option.
3. If the Storage Cruiser's agent has been installed on the Management Server implement the following procedures:
 1. Stop the Storage Cruiser's agent daemon.
 2. Change the starting IP address of the Storage Cruiser's agent.
 3. Start the Storage Cruiser's agent daemon.

Refer to "Starting and Stopping Agent" in the *ETERNUS SF Storage Cruiser Operation Guide* for information on how to start and stop the Storage Cruiser's agent. Refer to "setagtip (Agent Information Change Command)" in the *ETERNUS SF Storage Cruiser Operation Guide* for information on how to change the activation IP address of the Storage Cruiser's agent.

4. Refer to "Change Server Information" in the *ETERNUS SF Web Console Guide* and change the IP address information.
5. When the Storage Cruiser fault monitoring function is in use, change the SNMP Trap sending destination IP address of each device being monitored. Refer to "sanadmsh chtraprcvip (SNMP Trap Transmission Place IP Address Change Command)" in the *ETERNUS SF Storage Cruiser Operation Guide* for details.



See

For information about making the above change in a cluster operation, refer to "[8.5.6.1 Changing IP Address of Management Server Transaction](#)".



Note

If you change or delete the IP version of the Management Server, take the following actions for the Managed Servers registered with pre-changed or pre-deleted IP version:

- When changing the IP version (change from IPv4 to IPv6 or from IPv6 to IPv4):
 1. Delete the target Managed Server
 2. Change the IP address of the Management Server according to the procedure.
 3. Re-register the target Managed Server with changed IP version.
- When deleting the IP version (change from both IPv4 and IPv6 to either one of IPv6 and IPv4):
 1. Change the IP address of the target Managed Server to the changed IP version.
 2. Change the IP address of the Management Server according to the procedure.

8.5.2 Changing IP Address of Managed Server

Change the IP address of a Managed Server as follows:

- When the Storage Cruiser's agent has been installed on a Managed Server
 1. On the target Managed Server, stop the AdvancedCopy Manager daemons. For more information on stopping daemons, refer to "[Chapter 2 Starting and Stopping Daemons](#)".
 2. On the target Managed Server, stop the Storage Cruiser's agent daemon. Refer to "Starting and Stopping Agent" in the *ETERNUS SF Storage Cruiser Operation Guide* for information on how to stop it.
 3. On the target Managed Server, change the system IP address.
 4. On the target Managed Server, restart the AdvancedCopy Manager daemons. For more information on starting a daemons, refer to "[Chapter 2 Starting and Stopping Daemons](#)".
 5. Refer to "setagtip (Agent Information Change Command)" in the *ETERNUS SF Storage Cruiser Operation Guide* to change the Storage Cruiser agent's starting IP address.
 6. On the target Managed Server, start the Storage Cruiser's agent daemon. Refer to "Starting and Stopping Agent" in the *ETERNUS SF Storage Cruiser Operation Guide* for information on how to start it.
 7. Refer to "Change Server Information" in the *ETERNUS SF Web Console Guide* and change the IP address information.
- When the Storage Cruiser's agent has not been installed on the Managed Server
 1. On the target Managed Server, stop the AdvancedCopy Manager daemons. For more information on stopping daemons, refer to "[Chapter 2 Starting and Stopping Daemons](#)".
 2. On the target Managed Server, change the system IP address.
 3. On the target Managed Server, start the AdvancedCopy Manager daemons. For more information on stopping daemons, refer to "[Chapter 2 Starting and Stopping Daemons](#)".
 4. When AdvancedCopy Manager is being operated using Web Console, refer to "Change Server Information" in the *ETERNUS SF Web Console Guide* and change the IP address information. When AdvancedCopy Manager is being operated using only Command Line Interface, after creating a server information change instruction file with the Management Server, execute the `stgxfwcmmodsrv` command with the -f option and change the IP address. When using the replication management function, execute the `swsrecoverres` command with the -r option and the -h option on the Management Server. Specify the server name of the Managed Server running the replication operation with the -h option.



See

For information about performing the above change in cluster operation, refer to ["8.5.6.2 Changing IP Address of Managed Server Transaction"](#).

8.5.3 Changing Port Number of Management Server or Managed Server

Change the port number of a Management Server or Managed Server as follows:

When changing port numbers for AdvancedCopy Manager communication services:

1. On the target Management Server or Managed Server, change the port number of the AdvancedCopy Manager communication daemon (stgxfws) specified in /etc/services.
2. On the target Management Server or Managed Server, restart the AdvancedCopy Manager daemons. For information on restarting daemons, refer to ["Chapter 2 Starting and Stopping Daemons"](#).
3. On the Management Server, create a server information change instruction file, and then execute the `stgxfwcmmodsrv` command with the `-f` option specified to change the port number. When using the replication management function, execute the `swsrprecoverres` command with the `-r` option and the `-h` option on the Management Server. Specify the server name of the Managed Server running the replication operation with the `-h` option.
4. Refer to "Reload Server Configuration Information" in the *ETERNUS SF Web Console Guide*, and reload the configuration information of the Management Server or Managed Server.

When changing port numbers for other than AdvancedCopy Manager communication services:

Refer to "Changing Port Numbers (Linux)" in the *ETERNUS SF Express Operation Guide*. Replace "Express" with "AdvancedCopy Manager" to read the manual.



See

For information about performing the above change in cluster operation, refer to ["8.5.6.3 Changing Port Number of Management Server Transaction or Managed Server Transaction"](#).

8.5.4 Changing Server Name of Management Server or Managed Server

Change the server name of a Management Server or Managed Server as follows:

1. Change the system server name at the Management Server or Managed Server that is to be changed.
If the server name to be changed is only that of a Management Server or Managed Server managed by AdvancedCopy Manager, this step is not required.
2. On the Management Server, create a server information change instruction file, and then execute the `stgxfwcmmodsrv` command with the `-f` option specified to change the server name. When using the replication management function, execute the `swsrprecoverres` command with the `-r` option and the `-h` option on the Management Server. Specify the server name of the Managed Server running the replication operation with the `-h` option.
3. Refer to "Reload Server Configuration Information" in the *ETERNUS SF Web Console Guide* and reload the configuration information of the Management Server or Managed Server.



See

For information about performing the above change in a cluster operation, refer to ["8.5.6.4 Changing Server Name of Management Server Transaction or Managed Server Transaction"](#).

8.5.5 Changing Device Configuration

Changing the device configuration refers to any of the following:

- Changing the mount name
- Adding and deleting partitions
- Change from File System operation to Raw Device operation
- Capacity change of partitions
- Start position for change of partitions
- Changing the box identifier (Box ID)
- Recreating LUNs

8.5.5.1 Backup Management

Change the device configuration by using the following procedure. This processing is executed for all volumes on the disk containing the volume whose device configuration is to be changed.

1. Check whether there is a transaction volume on the target disk that needs to be backed up. If there is, delete the transaction volume. For information on this, refer to "[3.6.2.1 Deleting Transaction Volume](#)".
2. Check whether a backup volume exists in the target disk. If it does, delete the backup volume. For information on this, refer to "[3.6.2.2 Deleting Backup Volume](#)".
3. Change the device configuration.
4. Because the post-change device configuration needs to be reflected in the AdvancedCopy Manager data set, execute device information change processing. Refer to "[3.4.4 Fetching Device Information on Managed Server](#)", for details on changing the device information.
5. Reset the deleted transaction or backup volume in the backup. Refer to "[3.4.6 Setting Operation Type for 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 and backup volumes are deleted (in steps 1 and 2), history information cannot be deleted, and synchronized backup processing cannot be cancelled. This could possibly prevent the deletion of the transaction and backup volumes.

In this case, set the emergency operation mode, and delete the history information or cancel synchronized backup processing. For details of the emergency operation mode, refer to "[swsthistdel](#)" and "[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. If there are, delete them. For information about this, refer to "[6.6.2 Deleting Source Volume or Replica Volume](#)".
2. Change the device configuration.
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 "[6.4.4 Fetching Device Information on Managed Server](#)", for details about changing the device information.
4. Reset a deleted replica volume in the replication operation. Refer to "[6.4.5 Configuring Source Volume and Replica Volume](#)", for details about adding the replica volume.

Point

If the device configuration is changed (in step 2 above) before the replication volumes are deleted (in step 2), replication processing cannot be cancelled. This could possibly prevent the deletion of the replication volumes.

In this case, set the emergency operation mode, and cancel replication processing.

For details on the emergency operation mode, refer to "[swsrpcancel](#)".

8.5.6 Changing Configuration in Cluster Operation

This section describes how to make a change to the operational configuration in cluster operation.

8.5.6.1 Changing IP Address of Management Server Transaction

Refer to "Changing IP Address for Management Server Transaction" in "Changing Operation Environment (for Solaris/Linux)" in the *ETERNUS SF Cluster Environment Setup Guide*.

8.5.6.2 Changing IP Address of Managed Server Transaction

Change the IP address of a Managed Server transaction by performing the following steps:

1. Stop the Managed Server transaction.
For details on how to stop it, refer to the relevant cluster software manual.
2. Change the IP address resource.
For details on how to change an IP address resource, refer to the relevant cluster software manual.
3. Start the Managed Server transaction.
For details on how to start it up, refer to the relevant cluster software manual.
4. When AdvancedCopy Manager is being operated using Web Console, refer to "Change Server Information" in the *ETERNUS SF Web Console Guide* and change the IP address information. When AdvancedCopy Manager is being operated using only Command Line Interface, after creating a server information change instruction file with the Management Server, execute the [stgxfwcmmodsrv](#) command with the -f option and change the IP address. When using the replication management function, execute the [swsrprecoverres](#) command with the -r option and the -h option on the Management Server. Specify the server name of the Managed Server running the replication operation with the -h option.

8.5.6.3 Changing Port Number of Management Server Transaction or Managed Server Transaction

For changing the port number of Management Server transaction, refer to "Changing Port Number of Management Server Transactions" in "Changing Operation Environment (for Solaris/Linux)" in the *ETERNUS SF Cluster Environment Setup Guide*.

Change the port number of a Managed Server transaction by performing the following steps:

1. Stop the Management Server transaction or Managed Server transaction.
For details on how to stop it, refer to the relevant cluster software manual.
2. On the primary node and secondary node of the target Managed Server transaction, change the port number of the AdvancedCopy Manager communication daemon ([stgxfws_logicalNodeName](#)) specified in */etc/services*.
3. Start the targeted Managed Server transaction.
For details on how to start it up, refer to the relevant cluster software manual.
4. Create a server information change instruction file at the Management Server that manages the Managed Server. Then, execute the [stgxfwcmmodsrv](#) command with the -f option specified to change the port number. When using the replication management function, execute the [swsrprecoverres](#) command with the -r option and the -h option on the Management Server. Specify the server name of the Managed Server running the replication operation with the -h option.
5. Refer to "Reload Server Configuration Information" in the *ETERNUS SF Web Console Guide* and reload the configuration information of the Management Server or Managed Server.

8.5.6.4 Changing Server Name of Management Server Transaction or Managed Server Transaction

For changing the server name of Management Server transaction, refer to "Changing Server Name of Management Server Transaction" in "Changing Operation Environment (for Solaris/Linux)" in the *ETERNUS SF Cluster Environment Setup Guide*.

Change the server name of a Managed Server transaction by performing the steps below.

1. Create a server information change instruction file at the Management Server that manages the Managed Server. Then, execute the `stgxfwcmmodsrv` command with the `-f` option specified to change the port number. When using the replication management function, execute the `swsrprecoverres` command with the `-r` option and the `-h` option on the Management Server. Specify the server name of the Managed Server running the replication operation with the `-h` option.
2. Refer to "Reload Server Configuration Information" in the *ETERNUS SF Web Console Guide* and reload the configuration information of the Management Server or Managed Server.

8.5.6.5 Changing Logical Node Name

For changing the logical node name of Management Server transaction, refer to "Changing Logical Node Name" in "Changing Operation Environment (for Solaris/Linux)" in the *ETERNUS SF Cluster Environment Setup Guide*.

Change the logical node name of a Managed Server transaction by performing the following steps:

1. Check whether the target Managed Server transaction is operating.
If the transaction is not operating, start the Managed Server transaction.
For details on how to start it up, refer to the relevant cluster software manual.
2. Delete the configuration of the Managed Server transaction.
When using the backup management function, delete the transaction volume and backup volume by referring to "3.6.2 Deleting Device".
When it is defined as a copy source or copy destination volume by using the replication management function, delete the copy source or copy destination volume settings by referring to "6.6.2 Deleting Source Volume or Replica Volume".
3. Save the following AdvancedCopy Manager data.
 - Pre-processing/post-processing scripts
Save the pre-processing and post-processing scripts for backup management and replication management.For a "package-backup-of-the-management-information execution server" of AdvancedCopy Manager, save the following data:
 - Policy file for package backup of the management information
For details about the policy file path, refer to "8.1.2.1 Preliminaries".
4. Delete `/etc/opt/swstorage/logical/NodeName/swnode.ini` on the node operating the transaction.
5. Delete the cluster environment of the target transaction.
For details about how to delete, refer to "Deletion for Solaris/Linux/HP-UX version Clustered Systems" in the *ETERNUS SF Cluster Environment Setup Guide*.

Note

Do not perform the following step in "Deletion for Solaris/Linux/HP-UX version Clustered Systems" in the *ETERNUS SF Cluster Environment Setup Guide*.

- Step 1 described in "Deletion Procedure"

6. Create the cluster environment of the target transaction.

For details on how to create, refer to "Customization for Solaris/Linux/HP-UX version Clustered Systems" in the *ETERNUS SF Cluster Environment Setup Guide*.



Note

When you perform the cluster set up, specify the name of the new logical node.

7. Start the target transaction.
For details on how to start it up, refer to the relevant cluster software manual.
8. Restore the following AdvancedCopy Manager data.
 - Pre-processing/post-processing scripts
Restore the pre-processing and post-processing scripts for backup management and replication management.For a "package-backup-of-the-management-information execution server" of AdvancedCopy Manager, restore the following data:
 - Policy file for package backup of the management information
For details about the policy file path, refer to ["8.1.2.1 Preliminaries"](#).
9. Change the server name.
Perform the procedure described in ["8.5.6.4 Changing Server Name of Management Server Transaction or Managed Server Transaction"](#).
10. To use the backup management or replication management function, rebuild the backup management or replication management environment. (For backup management, perform the ["3.4.5 Configuring Environment for Backup Operation Server"](#) task and subsequent tasks. For replication management, perform the ["6.4.5 Configuring Source Volume and Replica Volume"](#) task and subsequent tasks.)
11. For a "package-backup-of-the-management-information execution server" of AdvancedCopy Manager, execute the [stgmgrinfo](#) command to register as its server.

8.5.6.6 Changing Transaction Name

Change the name of a Managed Server transaction by performing the following steps:

1. Check whether the target Managed Server transaction is operating.
If the transaction is not operating, start the Managed Server transaction.
For details about how to start it up, refer to the relevant cluster software manual.
2. Save the following AdvancedCopy Manager data.
 - AdvancedCopy Manager management file
On the node operating the transaction, save `/etc/opt/swstorage/logicalNodeName/snode.ini`.
 - Backup management list
Make a backup of the backup management list.
For details about how to do this, refer to ["8.1.1.1.1 Backing up Backup Management List"](#).
 - Pre-processing/post-processing scripts
Save the pre-processing and post-processing scripts for backup management and replication management.For a "package-backup-of-the-management-information execution server" of AdvancedCopy Manager, save the following data:
 - Policy file for package backup of the management information
For details about the policy file path, refer to ["8.1.2.1 Preliminaries"](#).
3. Delete `/etc/opt/swstorage/logicalNodeName/snode.ini` on the node operating the transaction.
4. Delete the cluster environment of the target transaction.
For details on how to do this, refer to "Deletion for Solaris/Linux/HP-UX version Clustered Systems" in the *ETERNUS SF Cluster Environment Setup Guide*.



Note

Do not perform the following step of "Deletion for Solaris/Linux/HP-UX version Clustered Systems" in the *ETERNUS SF Cluster Environment Setup Guide*.

- Step 1 described in "Deletion Procedure"

5. On the clustered system, change the name of cluster transaction.
For details on how to do this, refer to the relevant cluster software manual.

6. Create the cluster environment of the target transaction.

For details on how to do this, refer to "Customization for Solaris/Linux/HP-UX version Clustered Systems" in the *ETERNUS SF Cluster Environment Setup Guide*.



Note

When you perform the cluster set up, do not change any information other than the transaction name.

7. Start the target transaction.
For details on how to start it up, refer to the relevant cluster software manual.

8. Restore the following AdvancedCopy Manager data.

- AdvancedCopy Manager management file

On the node operating the transaction, restore `/etc/opt/swstorage/logical/NodeName/swnode.ini`.

- Backup management list

Restore the backup management list.

For details about how to do this, refer to "[8.1.1.1.2 Restoring Backup Management List](#)".

- Pre-processing/post-processing scripts

Restore the pre-processing and post-processing scripts for backup management and replication management.

- When using the replication management function, restore the environment by using the `swsrprecoverres` command with the `-r` option.

For a "package-backup-of-the-management-information execution server" of AdvancedCopy Manager, restore the following data:

- Policy file for package backup of the management information

For details about the policy file path, refer to "[8.1.2.1 Preliminaries](#)".

9. For a "package-backup-of-the-management-information execution server" of AdvancedCopy Manager, execute the `stgmgrinfo` command to register as its server.

8.5.6.7 Changing Shared Disk for Shared Data

Change a shared disk that is used for Managed Server transaction by performing the following steps:

1. Check whether the target Managed Server transaction is operating.
If the transaction is not operating, start the Managed Server transaction.
For details about how to start it up, refer to the relevant cluster software manual.

2. Save the following AdvancedCopy Manager data.

- AdvancedCopy Manager management file

On the node operating the transaction, save `/etc/opt/swstorage/logical/NodeName/swnode.ini`.

- Backup management list

Make a backup of the backup management list.

For details about how to do this, refer to "[8.1.1.1.1 Backing up Backup Management List](#)".

- Pre-processing/post-processing scripts
Save the pre-processing and post-processing scripts for backup management and replication management.

For a "package-backup-of-the-management-information execution server" of AdvancedCopy Manager, save the following data:

- Policy file for package backup of the management information
For details about the policy file path, refer to "8.1.2.1 Preliminaries".

3. Delete `/etc/opt/swstorage/logical/NodeName/swnode.ini` on the node operating the transaction.
4. Delete the cluster environment of the target transaction.

For details about how to do this, refer to "Deletion for Solaris/Linux/HP-UX version Clustered Systems" in the *ETERNUS SF Cluster Environment Setup Guide*.

Note

Do not perform the following step of "Deletion for Solaris/Linux/HP-UX version Clustered Systems" in the *ETERNUS SF Cluster Environment Setup Guide*.

- Step 1 described in "Deletion Procedure"

5. Create the cluster environment of the target transaction.

For details about how to do this, refer to "Customization for Solaris/Linux/HP-UX version Clustered Systems" in the *ETERNUS SF Cluster Environment Setup Guide*.

Note

- To change the name of a mount point, specify the new mount point name when you perform the cluster set up.
- To change a physical disk, specify the new physical disk when you mount it.
- Change the configuration of a shared disk in the clustered system.

6. Start the target transaction.

For details about how to start it up, refer to the relevant cluster software manual.

7. Restore the following AdvancedCopy Manager data.

- AdvancedCopy Manager management file

On the node operating the transaction, restore `/etc/opt/swstorage/logical/NodeName/swnode.ini`.

- Backup management list

Restore the backup management list.

For details about how to restore do this, refer to "8.1.1.2 Restoring Backup Management List".

- Pre-processing/post-processing scripts

Restore the pre-processing and post-processing scripts for backup management and replication management.

- When using the replication management function, restore the environment by using the `swsrprecoverres` command with the `-r` option.

For a "package-backup-of-the-management-information execution server" of AdvancedCopy Manager, restore the following data:

- Policy file for package backup of the management information
For details about the policy file path, refer to "8.1.2.1 Preliminaries".

8. For a "package-backup-of-the-management-information execution server" of AdvancedCopy Manager, execute the `stgmgrinfo set` command to register as its server.

8.5.7 Operation on Model Upgrade for ETERNUS Disk Storage System

Perform the procedure of "Operation on Model Upgrade for ETERNUS Disk Storage System" in the *ETERNUS SF Web Console Guide*.

8.5.8 Changing the SSL Version Settings for Communicating with Supported Devices

If you change the SSL version setting for communicating with the following devices that are registered for this product, restart the ETERNUS SF Manager if the protocol Maintenance-Secure TLSv1.2 is disabled, or TLSv1.2 and TLSv1.1 are disabled:

- ETERNUS DX S4 series
- ETERNUS DX S3 series
- ETERNUS AF All-Flash Arrays
- ETERNUS DX200F

8.6 Action When Managed Server Information Cannot Be Deleted

If it becomes necessary to delete Managed Server information under the circumstances specified below, this section provides a way to delete the Managed Server information forcibly from the Management Server.

- The operating system on the Managed Server cannot be started because of a hardware failure or another similar reason. From that point, the server is not used for operations.
- The Managed Server has been physically removed without performing the Managed Server deletion procedure ("[3.6.4 Deleting Managed Server](#)").
- The Managed Server is managed by AdvancedCopy Manager and has been physically removed without being deleted in AdvancedCopy Manager.

For cases other than these, refer to "[3.6.4 Deleting Managed Server](#)" to delete the Managed Server.



- When this procedure is performed, the specified Managed Server is forcibly deleted regardless of the backup or replication status. Therefore, this procedure should be used only for Managed Server that are not planned to be used afterwards.
- Do not execute a backup or replication while this procedure is performed.
- The components and management information of AdvancedCopy Manager on the forcibly deleted Managed Server are not deleted. When the forcibly deleted server is to be used again, uninstall the AdvancedCopy Manager from the deleted server before it is reconnected to the network. If its server is reconnected to the networks without uninstalling first AdvancedCopy Manager, the AdvancedCopy Manager functions of other servers on the network may not work correctly.

The procedure to delete forcibly a Managed Server is described below:

Confirm that the Managed Server is not used afterwards or that it is removed before starting the procedure.

1. Check the status of AdvancedCopy Manager on the Managed Server to be deleted forcibly.
If AdvancedCopy Manager is operable, use the normal deletion procedure for deleting the Managed Server instead of the forcible deletion procedure. For details of the normal deletion procedure, refer to "[3.6.4 Deleting Managed Server](#)".

How to Check Status of Managed Server

On the Management Server, execute the `stgxfwcmsetdev` command with `-n` option. Specify the target Managed Server with an operand of `-n` option.

After executing the command, check the status of Managed Server as follows.

Go to step 2, only when the Managed Server cannot be started.

- When the Managed Server cannot be started:

The following message is displayed:

"swstf5405 The acquisition of device information caused a failure.
Reason=swstf2111 Connection to Storage Server failed."

- When the Managed Server can be started:

The following message is displayed:

"swstf5408 There is no change in information currently being managed."

Alternatively the command may be waiting for input of the device information confirmation. If the command is waiting for input, enter 'q' or 'n' to prevent updating the device information and end the stgxfwcmsetdev command execution.

2. Stop the Advanced Copy sessions.

Even if the Managed Server information has been deleted, the Advanced Copy sessions for the connected disks (Logical Units) may not be stopped.

For the unnecessary Advanced Copy sessions, stop its sessions by using ETERNUS Web GUI.

3. Back up the database.

For the procedures, refer to "[8.7.2.1 Backup of Common Control and AdvancedCopy Manager Repository](#)".

However, for the "pg_dump" command, read the commands below.

```
/opt/FJSVesfcm/postgres/bin/pg_dump -U esfpostgres -a --attribute-inserts -p portNumber -t escs.\
\"BASE_\"* -t escs.\"REP_REPVOL\" -t escs.\"BR_\"* -f $BAK_DIR/Common/fileName esfdb
```

4. Delete the Managed Server information forcibly.

Execute the [stgforcedelsrv](#) command.



Note

If the stgforcedelsrv command terminated abnormally, do not proceed from step 5 onwards.

If the steps after step 5 are executed an incorrect database is restored and normal system operation may not be performed adequately afterwards.

5. Restore the database.

Restore the database from the directory that the database is stored. Its directory must be a directory of "after-dir" specified with -a option in [stgforcedelsrv \(Managed Server forcibly deletion command\)](#) at step 4.

- Stop the ETERNUS SF Manager Web service.

Execute the following shell script to stop the service.

```
/opt/FJSVesfcm/bin/stop-webservice.sh
```

- Restore a backup data. The following procedure must be implemented by a user with administrator privileges for the operating system.

- Execute the following command to restore the backup data.

If you did not change the port number, specify 15432 for *portNumber*. If you did change the port number, specify the changed port number for *portNumber*.

The *fileName* is the file name specified when the backup data was created.

```
/opt/FJSVesfcm/postgres/bin/psql -U esfpostgres -p portNumber -f $BAK_DIR/Common/
fileName esfdb
```

- The prompt is displayed, and waiting for input. Input the following strings.

```
master01!
```

- Start the ETERNUS SF Manager Web service.

Execute the following shell script to start the service.

```
/opt/FJSVesfcm/bin/start-webservice.sh
```

6. Perform the following operation according to the status of the forcibly deleted Managed Server.
 - If a replication operation has been done:
 - If the forcibly deleted Managed Server was part of a server-to-server replication:

On the Management Server, execute the `swsrprecoverres` command with the `-h` and `-r` options. For the operand of `-h` option, specify the name of the server used for the replication with the forcibly deleted Managed Server. If several servers where used for replication, the `swsrprecoverres` command has to be executed for each of those servers.
 - If the forcibly deleted Managed Server was not part of a server-to-server replication:

No operation is required.
 - If a backup operation has been done:

No operation is required.
7. Take the following actions to reflect to Web Console depending on how the agent installed in the forcibly deleted Managed Server is used:
 - If both of AdvancedCopy Manager's agent and Storage Cruiser's agent are used:

With Web Console, execute the **Reload Conf.** operation to update to the latest information.
 - If only AdvancedCopy Manager's agent is used:

With Web Console, delete the Managed Server that you have deleted.

8.7 Maintaining Database of ETERNUS SF Manager

This section explains the maintenance for database of ETERNUS SF Manager.

8.7.1 Saving Storage Cruiser's Manager Environment

To save a Storage Cruiser's manager environment, refer to "Backup of Manager" in the *ETERNUS SF Storage Cruiser Operation Guide*.

8.7.2 Maintaining Common Control and AdvancedCopy Manager Repository

This section explains the maintenance of common control and AdvancedCopy Manager repository.

8.7.2.1 Backup of Common Control and AdvancedCopy Manager Repository

Backup the common control and AdvancedCopy Manager repository.



Information

Notation in Procedure

Directory Name	Explanation
\$BAK_DIR	Backup destination directory.

1. Stop the ETERNUS SF Manager Web service.
Execute the following shell script to stop the service.

```
/opt/FJSVesfcm/bin/stop-webservice.sh
```

2. Create a backup data.

1. Create the following backup directories:

```
mkdir -p $BAK_DIR/Common
```

2. Execute the following command to create a backup data with OS Administrator privilege.

If you did not change the port number, specify 15432 for *portNumber*. If you did change the port number, specify the changed port number for *portNumber*.

The file name of the created backup data is *fileName*.

```
/opt/FJSVesfcm/postgres/bin/pg_dump -U esfpostgres -C --attribute-inserts -p portNumber -f
$BAK_DIR/Common/fileName esfdb
```

3. The prompt is displayed, and waiting for input. Input the following strings.

```
master01!
```

3. Start the ETERNUS SF Manager Web service.

Execute the following shell script to start the service.

```
/opt/FJSVesfcm/bin/start-webservice.sh
```

8.7.2.2 Restore of Common Control and AdvancedCopy Manager Repository

Restore the common control and AdvancedCopy Manager repository.



Information

Notation in Procedure

Directory Name	Explanation
\$BAK_DIR	Backup destination directory.

1. Stop the ETERNUS SF Manager Web service.

Execute the following shell script to stop the service.

```
/opt/FJSVesfcm/bin/stop-webservice.sh
```

2. Restore a backup data. The following procedure must be implemented by a user with administrator privileges for the operating system.

1. Execute the following command.

If you did not change the port number, specify 15432 for *portNumber*. If you did change the port number, specify the changed port number for *portNumber*.

```
/opt/FJSVesfcm/postgres/bin/psql -U esfpostgres -p portNumber postgres
```

2. The prompt is displayed, and waiting for input. Input the following strings in order.

```
master01!
drop database esfdb;
\q
```

3. Execute the following command to restore the backup data.

If you did not change the port number, specify 15432 for *portNumber*. If you did change the port number, specify the changed port number for *portNumber*.

The *fileName* is the file name specified when the backup data was created.

```
/opt/FJSVesfcm/postgres/bin/psql -U esfpostgres -p portNumber -f $BAK_DIR/Common/fileName
postgres
```

4. The prompt is displayed, and waiting for input. Input the following strings.

```
master01!
```

3. Start the ETERNUS SF Manager Web service.

Execute the following shell script to start the service.

```
/opt/FJSVesfcm/bin/start-webservice.sh
```



Note

After restoring the shared control repository, it takes approximately 2 minutes before Web Console becomes available.

8.7.3 Maintaining License Management Database

This section explains the maintenance of license management database.

8.7.3.1 Backup of License Management Database

Save the license management database.



Information

Notation in Procedure

Directory Name	Explanation
\$BAK_DIR	Backup destination directory.

1. Create the following backup directories:

```
mkdir -p $BAK_DIR/LM
```

2. Execute the following command to create a backup data with OS Administrator privilege.
The file name of the created backup data is *fileName*.

```
/opt/FJSVesflm/bin/esflm backup $BAK_DIR/LM/fileName
```

8.7.3.2 Restore of License Management Database

Restore the license management database.



Information

Notation in Procedure

Directory Name	Explanation
\$BAK_DIR	Backup destination directory.

Execute the following command with administrator privileges for the operating system.
The *fileName* is the file name specified when the backup data was created.

```
/opt/FJSVesflm/bin/esflm restore $BAK_DIR/LM/fileName
```

8.7.4 Saving AdvancedCopy Manager CCM Environment Settings Files

If AdvancedCopy Manager CCM is used by the server where ETERNUS SF Manager is installed, refer to "How to Back Up Environment Settings Files" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module* to save an AdvancedCopy Manager CCM environment settings files.

If AdvancedCopy Manager CCM environment settings files are already saved, it is unnecessary to perform this task.

Chapter 9 Command References

Note

Set the following paths in the environment variable PATH to execute the commands.

- /usr/sbin
- /usr/bin

9.1 Backup Management Commands

This section describes operations using Backup Management Commands.

Note

- Only a root user can execute all the commands of backup management. When command execution is carried out by other users, a message "swst0601 This Command cannot be run. This is because the user is not a root user." is displayed and then, an error end.
- Backup operations in Logical Units (disk units) are not supported. An error occurs if the [swstdevinfo](#) command is used to specify a Logical Unit (disk).
- The backup operations of AdvancedCopy Manager using the Storage Cluster Continuous Copy function are not supported.

Note

Operating Environment for Command Execution

The table below lists the operating environment requirements for different commands.

Table 9.1 Operating Environment for Executing Backup Management Commands

Command	Operating Environment of AdvancedCopy Manager
swstbackstat	The daemons or services of the Management Server and target Managed Server must be operating in order to execute this command, if the Management Server has either of the following parameter designations: <ul style="list-style-type: none">- -h option specified- No -h option and device name are specified
swstbackup	To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating. When the Managed Server runs on Windows and is in EC operation, the daemons or services of the Managed Server must be operating.
swstbkpoldel	To execute this command on a Managed Server, the daemons or services of the Management Server must be operating. To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating.
swstbkpoldisp	The daemons or services of the Management Server and target Managed Server must be operating in order to execute this command, if the Management Server has either of the following parameter designations: <ul style="list-style-type: none">- -h option specified- No -h option and device name are specified

Command	Operating Environment of AdvancedCopy Manager
swstbkpolset	To execute this command on a Managed Server, the daemons or services of the Management Server must be operating. To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating.
swstcancelrest	To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating.
swstcancelsync	To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating. When the Managed Server runs on Windows and is in EC operation, the daemons or services of the Managed Server must be operating.
swstcanceltrk	To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating.
swstdevdisp	The daemons or services of the Management Server and target Managed Server must be operating in order to execute this command, if the Management Server has either of the following parameter designations: <ul style="list-style-type: none"> - -h option specified - No -h option and device name are specified To execute this command with the -o option specified, the daemons or services of the Management Server must be operating.
swstdevinforeset	To execute this command on a Managed Server, the daemons or services of the Management Server must be operating. To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating.
swstexecstat	The daemons or services of the Management Server and target Managed Server must be operating in order to execute this command, if the Management Server has either of the following parameter designations: <ul style="list-style-type: none"> - -h option specified - No -h option and device name are specified
swsthistdel	To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating.
swsthistdisp	The daemons or services of the Management Server and target Managed Server must be operating in order to execute this command, if the Management Server has either of the following parameter designations: <ul style="list-style-type: none"> - -h option specified - No -h option and device name are specified
swstresback	-
swstresrst	The daemons or services of the Management Server must be operating. When the Managed Server runs on Windows, the daemons or services of the Managed Server must be always operating.
swstrestore	To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating.
swstreststat	The daemons or services of the Management Server and target Managed Server must be operating in order to execute this command, if the Management Server has either of the following parameter designations:

Command	Operating Environment of AdvancedCopy Manager
	<ul style="list-style-type: none"> - -h option specified - No -h option and device name are specified
swtstrsemth	<p>To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating.</p> <p>To execute this command on a Managed Server, the daemons or services of the Management Server must be operating.</p> <p>When the Managed Server runs on Windows, the daemons or services of the Managed Server must be always operating.</p>
swststartsync	<p>To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating.</p> <p>When the Managed Server runs on Windows, the daemons or services of the Managed Server must be operating.</p>
swstsvrdisp	<p>To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating.</p>
swstsvrset	<p>To execute this command on a Managed Server, the daemons or services of the Management Server must be operating.</p> <p>To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating.</p>
swstsyncstat	<p>The daemons or services of the Management Server and target Managed Server must be operating in order to execute this command, if the Management Server has either of the following parameter designations:</p> <ul style="list-style-type: none"> - -h option specified - No -h option and device name are specified
swsttrkstat	<p>The daemons or services of the Management Server and target Managed Server must be operating in order to execute this command, if the Management Server has either of the following parameter designations:</p> <ul style="list-style-type: none"> - -h option specified - No -h option and device name are specified

Note

In Cluster Operation

In cluster operation, execute a command by performing the following steps:

1. If necessary, specify the logical node name to the transaction in the environment variable SWSTGNODE. Specify the environment variable as follows:

```
# SWSTGNODE=logicalNodeName
# export SWSTGNODE
```

2. Execute the command.

The table below lists the commands available in cluster operation, notes whether environment variables must be configured, and specifies the nodes on which commands are executable. For notes and cautions on command execution, refer to the Remarks column.

Table 9.2 Commands When Cluster Operation Is Used

Command	Environment Variable	Command Execution Node	Remarks
swstsvrset	Required	Transaction operating node	-
swstsvrdisp	Required	Transaction operating node	-
swstdevinfo	Required	Transaction operating node	-
swstdevdisp	Required	Transaction operating node	-
swstbkpolset	Required	Transaction operating node	-
swstbkpoldisp	Required	Transaction operating node	-
swstbkpoldel	Required	Transaction operating node	-
swstbackup	Required	Transaction operating node	-
swstbackstat	Required	Transaction operating node	-
swstrestore	Required	Transaction operating node	-
swstcancelrest	Required	Transaction operating node	-
swstreststat	Required	Transaction operating node	-
swsthistdisp	Required	Transaction operating node	-
swsthistdel	Required	Transaction operating node	-
swststartsync	Required	Transaction operating node	-
swstcancelsync	Required	Transaction operating node	-
swstsyncstat	Required	Transaction operating node	-
swstexecstat	Required	Transaction operating node	-
swstcanceltrk	Required	Transaction operating node	-
swstrkstat	Required	Transaction operating node	-
swstsrsemch	Required	Transaction operating node	-
swstresback	Required	Transaction operating node	-
swstresrst	Required	Transaction operation node	-

 Note

In a Management Server, when a display-type command that has 'disp' or 'stat' at the end of the command name is executed, the display contents depend on the existence of the -h option and device name. For details, refer to the following table.

Table 9.3 Display Contents of Display-type Commands in Management Server

-h Option	Device Name Specified	Display Contents
Does not exist	No	Displays all Managed Server information registered using the Backup management functionality. If an error occurs on any of the Managed Servers during display processing, the processing is terminated.
	Yes	Displays only the specified Management Server device information.
Exists	No	Displays all Managed Server information specified using the -h option.
	Yes	Displays only the specified device information for the Managed Server specified using the -h option.

9.1.1 Environment Definition Commands

This section describes the environment definition commands for backup management.

9.1.1.1 swstsvrset (Managed Server Configuration Information Setting Command)

NAME

swstsvrset - sets configuration information of a Managed Server

SYNOPSIS

To Execute on Managed Server

```
/opt/FJJSVswsts/bin/swstsvrset
```

To Execute on Management Server

```
/opt/FJJSVswsts/bin/swstsvrset [-h serverName]
```

DESCRIPTION

This command sets configuration information of a Managed Server. Execution of this command declares that the Managed Server on which the command is executed is a server that performs backup operations.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdisprv](#) command.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Set the configuration information of a Managed Server:

```
# /opt/FJJSVswsts/bin/swstsvrset
swstsvrset completed
#
```

- Set the configuration information of a Managed Server named job1 from a Management Server:

```
# /opt/FJJSVswsts/bin/swstsvrset -h job1
swstsvrset completed
#
```

NOTES

- This command can only be executed if none of the following commands are running:
 - [swstsvrset](#)
 - [swstbkpolset](#)
 - [swstbkpoldel](#)
 - [swstbackup](#)
 - [swstrestore](#)

- [swstcancelrest](#)
- [swststartsync](#)
- [swstcancelsync](#)
- [swstdevinforeset](#)
- [swsthistdel](#)
- [swstcanceltrk](#)
- [swstresback](#)
- [swstresrst](#)
- [swstsrsemitch](#)

9.1.1.2 swstsvrdisp (Managed Server Configuration Information Display Command)

NAME

swstsvrdisp - displays configuration information of a Managed Server

SYNOPSIS

To Execute on Managed Server

```
/opt/FJJSVswsts/bin/swstsvrdisp
```

To Execute on Management Server

```
/opt/FJJSVswsts/bin/swstsvrdisp [-h serverName]
```

DESCRIPTION

This command displays configuration information of a Managed Server that has been set using the [swstsvrset](#) command. The following information is displayed.

Title	Description
Storage-Server	Displays a Managed Server name.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdispsrv](#) command.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Display configuration information of a Managed Server named Work1:

```
# /opt/FJJSVswsts/bin/swstsvrdisp
Storage-Server = Work1
#
```

9.1.1.3 swstdevinfoaset (Device Information Setting Command)

NAME

swstdevinfoaset - sets the device information

SYNOPSIS

To Execute on Managed Server

```
/opt/FJJSVswsts/bin/swstdevinfoaset -t | -b | -o deviceName
```

To Execute on Management Server

```
/opt/FJJSVswsts/bin/swstdevinfoaset [-h serverName] -t | -b | -o deviceName
```

DESCRIPTION

Set a volume that has been allocated to the ETERNUS Disk storage system, as a transaction or backup volume.

A transaction volume refers to a volume that stores transaction data to be backed up. A backup volume refers to a volume used to store a backup of this transaction data.

If you want to perform the backup operation on a regular transaction volume, refer to "[3.2.2 Selecting Backup Source](#)" for information on a transaction volume, and "[3.2.4 Preparing Backup Volume](#)" for information on a backup volume.

PRIMECLUSTER GD logical volumes can also be registered.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdisprv](#) command.

-t

Register the specified device as a transaction volume.

Cannot be specified at the same time as the -b and -o options.

-b

Register the specified device as a backup volume.

Cannot be specified at the same time as the -t and -o options.

-o

Deregister the specified device if it has been registered as a transaction or backup volume.

Cannot be specified at the same time as the -t and -b options.

OPERANDS

deviceName

Specify an AdvancedCopy Manager device name.

Refer to "[1.4 Managing Device on AdvancedCopy Manager](#)" for the devices to be specified.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Register a device as a transaction volume:

```
# /opt/FJSVswsts/bin/swstdevinfoaset -t /dev/dsk/clt0d0s6  
swstdevinfoaset completed  
#
```

- Register a device as a backup volume:

```
# /opt/FJSVswsts/bin/swstdevinfoaset -b /dev/dsk/clt0d2s6  
swstdevinfoaset completed  
#
```

- Deregister a device that has been registered as a transaction volume:

```
# /opt/FJSVswsts/bin/swstdevinfoaset -o /dev/dsk/clt0d0s6  
swstdevinfoaset completed  
#
```

NOTES

- A device cannot be registered as a transaction volume if:
 - Another device with the same box information (box identifier, OLU number, EXTENT starting position, and EXTENT size) has already been registered as a transaction volume.
 - The space of the device to be registered overlaps the space of an already registered working or backup volume.
 - When the specified volume is a logical volume of VxVM in an unsupported structure. For the supported logical volumes of VxVM, refer to "[1.4 Managing Device on AdvancedCopy Manager](#)".
 - The device is a Logical Unit (disk).
 - The license for the ETERNUS Disk storage system that contains the specified transaction volume is one of the following:
 - The license is not registered in the license management function.
 - The license is expired.
 - The device is a TFOV.
- A device cannot be registered as a backup volume if:
 - Another device with the same box information (box identifier, OLU number, EXTENT starting position, and EXTENT size) has already been registered as a backup volume.
 - The space of the device to be registered overlaps the space of an already registered work or backup volume.
 - The database space of Symfaware is built on the device to be configured.
 - The device to be configured is already mounted.
 - The device to be set up is registered as a replication destination volume in replication management.
 - The device to be set up is registered as a replication source volume of the duplicate volume information in which bidirectional copy is possible in replication management.
 - The specified volume is a logical volume of VxVM in an unsupported structure. For the supported logical volumes of VxVM, refer to "[1.4 Managing Device on AdvancedCopy Manager](#)".
 - The device is a Logical Unit (disk).
 - The device is an SDX object.
 - The device is a TFOV.
- You must delete related backup history information before:
 - Changing a backup volume to a transaction volume

- Cancelling the setting of a backup volume
- You must delete a backup policy specified on a transaction volume before:
 - Changing a transaction volume to a backup volume
 - Cancelling the setting of a transaction volume
- The following changes are not possible unless the related synchronous processing (in an EC session) is cancelled:
 - Changing a setting in a suspended transaction volume
 - Changing a setting in a suspended backup volume
- The related tracking processing (in an OPC session) must be cancelled in advance to perform the following changes:
 - Changing the settings of a transaction volume in tracking processing
 - Changing the settings of a backup volume in tracking processing
- For details on how to use SDX objects of PRIMECLUSTER GD, refer to "[Configuration and Conditions of Unsupported SDX Objects](#)" and "[10.1.5.1 Notes on SDX Object Operations \(Backup Management\)](#)".
- This command can be executed only while none of the following commands are running:
 - `swstsvrset`
 - `swstdevinfo`
 - `swstbkpolset`
 - `swstbkpoldel`
 - `swstbackup`
 - `swstrestore`
 - `swstcancelrest`
 - `swststartsync`
 - `swstcancelsync`
 - `swsthistdel`
 - `swstcanceltrk`
 - `swstsrsemtrk`
 - `swstresback`
 - `swstresrst`

9.1.1.4 swstdevdisp (Device Usage Status Display Command)

NAME

swstdevdisp - displays the device information

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsts/bin/swstdevdisp [-t] [-b [-u]] [-o] [deviceName]
```

To Execute on Management Server

```
/opt/FJSVswsts/bin/swstdevdisp [-h serverName] [-t] [-b [-u]] [-o] [deviceName]
```

DESCRIPTION

This command displays information about a transaction or backup volume or any other definable device, which has been set using the [swstdevinfo](#) command.

The following information is displayed.

Title	Description
Server	Displays a Managed Server name.
Device	Displays the AdvancedCopy Manager device name .
Size	Displays the size of a partition allocated to a device. For a logical volume of VxVM, the size of the public area (slice) that makes up the volume is displayed instead of the volume size.
Device-Mode	Displays a character string by which a device can be identified: "Transaction" : Represents a transaction volume. "Backup" : Represents a backup volume. "Other" : Represents a device that can be registered as a working or backup volume. A backup volume is indicated as "Backup (used)" if it is being used or "Backup (free)" if it is not used. During the backup process with -suspend or -T option specified, the backup volume remains in use after history information is deleted. When the -u option is specified, "Backup (used-S)" is displayed for a backup volume in the suspend status for which history information was deleted, or "Backup (used-T)" is displayed for a backup volume in the tracking status for which history information was deleted.
Mount-Point (Method)	Displays the mount point of a device. The file system type of the mount point is displayed inside the parentheses.
Backup-Engine	Displays the operation type: "AdvancedCopy" : Backup operation using AdvancedCopy Manager

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdisprv](#) command.

-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** column.

-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 is displayed if all of the -t, -b and -o options are omitted.

OPERANDS

deviceName

Specify an AdvancedCopy Manager device name.

Refer to "[1.4 Managing Device on AdvancedCopy Manager](#)" for the devices to be specified.

Displays device information on all the transaction volumes if this operand and the associated options are omitted.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Display the usage statuses of the registered devices:

```
# /opt/FJSSwsts/bin/swstdevdisp -tb
Server Device          Size      Device-Mode  Mount-Point (Method) Backup-Engine
Work2  /dev/dsk/clt0d0s6    1.0 Gbyte Transaction /mnt/tran1 (ufs)   AdvancedCopy
Work2  /dev/dsk/clt1d0s6    1.0 Gbyte Transaction ---- (----)       AdvancedCopy
Work2  /dev/dsk/clt0d2s6    1.0 Gbyte Backup (used) ---- (----)       AdvancedCopy
Work2  /dev/dsk/clt0d2s3    1.0 Gbyte Backup (free) ---- (----)       AdvancedCopy
Work2  /dev/dsk/clt0d7s6    2.0 Gbyte Backup (free) ---- (----)       AdvancedCopy
Work2  /dev/dsk/clt0d7s3    2.0 Gbyte Backup (free) ---- (----)       AdvancedCopy
:
#
```

9.1.1.5 swstbkpolset (Backup Policy Setting Command)

NAME

swstbkpolset - sets a backup policy

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSSwsts/bin/swstbkpolset [-i intervalDays] [-s saveNumber] [deviceName]
```

To Execute on Management Server

```
/opt/FJSSwsts/bin/swstbkpolset [-h serverName] [-i intervalDays] [-s saveNumber] [deviceName]
```

DESCRIPTION

This command sets a backup policy (the number of preservation generations and the number of interval days).

The number of preservation generations is the number of generations of backup data that should be kept at one time. The number of interval days means the number of days from the execution of one backup until the next backup should be performed.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdspsrv](#) command.

-i *intervalDays*

Specifies the number of interval days. You can specify a value between 1 and 366. If this option is omitted during initial registration, 30 is set as the default. If this option is omitted during update, the existing value is inherited.

-s saveNumber

Specifies the number of preservation generations. You can specify a value between 1 and 31. If this option is omitted during initial registration, 2 is set as the default. If this option is omitted during update, the existing value is inherited.

OPERANDS

deviceName

Specify an AdvancedCopy Manager device name.

Refer to "[1.4 Managing Device on AdvancedCopy Manager](#)" for the devices to be specified.

If the device name is omitted, the backup policy is set for all the transaction volumes.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- To set 3 as the save number of a transaction volume as a backup policy:

```
# /opt/FJSSVswsts/bin/swstbkpolset -s 3 /dev/dsk/c1t0d5s3  
/dev/dsk/c1t0d5s3 swstbkpolset completed  
#
```

NOTES

- Even if the number of interval days is specified, AdvancedCopy Manager does not automatically create a backup.
- When you specify a backup policy, there must be as many backup volumes in operation according to the specified backup policy. To view information on the number of backup volumes required to perform the backup operation of a regular transaction volume, refer to "[3.2.4 Preparing Backup Volume](#)".
- You can change the number of preservation generations as long as the "history information count at the time" is equal to or less than the "number of preservation generations to be specified".
- Ordinary transaction volumes must be backed up to ordinary backup volumes.
If the transaction volume is a PRIMECLUSTER GD logical volume, the backup destination must be a backup volume of a PRIMECLUSTER GD logical volume.
- If SDX objects of PRIMECLUSTER GD are to be used, refer to information regarding "[Configuration and Conditions of Unsupported SDX Objects](#)".
- No backup policy can be set if:
 - There are not as many backup volumes as the specified number of preservation generations.
 - An option for more than the maximum number of preservation generations is specified.
 - An option for more than the maximum interval days is specified.
 - The replication management function is executing a process for the target transaction volume.
 - The specified device is an LU (disk unit).
- This command can only be executed while none of the following commands are running:
 - [swstsvrset](#)
 - [swstdevinfo](#)
 - [swstsrsem](#)
 - [swstresback](#)
 - [swstresrst](#)

9.1.1.6 swstbkpoldisp (Backup Policy Display Command)

NAME

swstbkpoldisp - displays a backup policy

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsts/bin/swstbkpoldisp [deviceName]
```

To Execute on Management Server

```
/opt/FJSVswsts/bin/swstbkpoldisp [-h serverName] [deviceName]
```

DESCRIPTION

This command displays a backup policy that has been set for a transaction volume. The following information is displayed.

Title	Description
Server	Displays a Managed Server name.
Device	Displays the AdvancedCopy Manager device name .
Interval-Days	Displays a number of interval days.
Save-Number	Displays a number of preservation generations.
Mount-Point (Method)	Displays the mount point of a device. The file system type of the mount point is displayed inside the parentheses.

OPTIONS

`-h serverName`

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdisprv](#) command.

OPERANDS

`deviceName`

Specify an AdvancedCopy Manager device name.

Refer to "[1.4 Managing Device on AdvancedCopy Manager](#)" for the devices to be specified.

If the device name is omitted, a backup policy for a transaction volume with a registered backup policy is displayed.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Display a backup policy of a transaction volume:

```
# /opt/FJSVswsts/bin/swstbkpoldisp /dev/dsk/clt0d5s3
Server Device          Interval-Days Save-Number Mount-Point (Method)
Work2 /dev/dsk/clt0d5s3 30              2           /stg(ufs)
#
```

9.1.1.7 swstbkpodel (Backup Policy Deletion Command)

NAME

swstbkpodel - deletes the backup policy

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsts/bin/swstbkpodel deviceName
```

To Execute on Management Server

```
/opt/FJSVswsts/bin/swstbkpodel [-h serverName] deviceName
```

DESCRIPTION

An existing backup policy can be deleted using this command.

Delete the backup history information of the transaction volume prior to the execution of this command.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdisprv](#) command.

OPERANDS

deviceName

Specify an AdvancedCopy Manager device name.

Refer to "[1.4 Managing Device on AdvancedCopy Manager](#)" for the devices to be specified.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- To delete the backup policies of a transaction volume:

```
# /opt/FJSVswsts/bin/swstbkpodel /dev/dsk/c1t0d5s3  
/dev/dsk/c1t0d5s3 swstbkpodel completed  
#
```

NOTES

- A backup policy cannot be deleted if:
 - The transaction volume to be deleted has backup history information.
 - A transaction volume whose backup policies are to be deleted has already started backup synchronous processing.
 - The replication management function is executing a process for the target transaction volume.
- This command can only be executed while none of the following commands are running:
 - [swstsvrset](#)
 - [swstdevinforeset](#)
 - [swstsrsematch](#)

- [swstresback](#)
- [swstresrst](#)

9.1.2 Operation Commands

This section describes operation commands for backup management.

9.1.2.1 swstbackup (Backup Execution Command)

NAME

swstbackup - backup

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsts/bin/swstbackup [-suspend|-T] deviceName [-Xdevmap deviceMapFile]
```

To Execute on Management Server

```
/opt/FJSVswsts/bin/swstbackup [-h serverName] [-suspend|-T] deviceName [-Xdevmap deviceMapFile]
```

DESCRIPTION

When using this command, the operation changes according to different states of the Advanced Copy that is being executed.

- When synchronous processing (EC) is not performed
Snapshot processing (OPC) is performed and the processing of information, such as the TimeStamp and backup volume of a copy target, is registered into backup history information.
- When synchronous processing (EC) is performed
The state of synchronous processing is checked and, in the case of an equivalent maintenance state, backup synchronous processing is suspended. Information such as time and backup volume of a copy target is simultaneously registered into backup history information. A backup of this form is called synchronous high-speed backup. Note that it aborts when an equivalent maintenance state is not reached. To use the "[Chapter 4 Backup Operation Using Suspend/Resume Function](#)" for synchronous high-speed backup, execute the backup command with `-suspend` option specified, and suspend backup synchronous processing.

In the above cases, a backup is established, and access to the backup volume is available immediately after the execution command.

AdvancedCopy Manager executes backup pre-processing and post-processing scripts before and after backup.

For information on customizing these scripts, refer to "[Appendix A Pre-processing and Post-processing of Backup and Restoration](#)".

OPTIONS

`-h serverName`

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdisprv](#) command.

`-Xdevmap deviceMapFile`

Specifies a target backup volume to perform the snapshot fast backup. In the operand, specify a "[device map file](#)" in which a combination of a transaction volume and a target backup volume are described. If this option is not specified, the target backup volume is automatically selected by the backup function. To execute this command on a remote basis using the `-h` option, specify a device map file (of a Managed Server) using an absolute pathname. A file name specified in Device-Map-File must not include any national character.

`-suspend`

This results in synchronous high-speed backup using the "[Chapter 4 Backup Operation Using Suspend/Resume Function](#)".

-T

Specifies that the differential snapshot high-speed backup be executed.

This option cannot be specified during synchronous processing.

This option is valid only for copying within a box when the ETERNUS Disk storage system supports the QuickOPC function.

If this option is not specified, ordinary snapshot processing (ie, OPC without using the QuickOPC function) is started. This option must always be specified to perform a backup operation using differential snapshots.

OPERANDS

deviceName

Specify an AdvancedCopy Manager device name.

Refer to "[1.4 Managing Device on AdvancedCopy Manager](#)" for the devices to be specified.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Back up a transaction volume:

```
# /opt/FJSSVswsts/bin/swstbackup /dev/dsk/c1t1d0s6  
/dev/dsk/c1t1d0s6 swstbackup completed  
#
```

NOTES

- To preserve data integrity, this command performs the backup only after inhibiting access to both transaction and backup volumes from all other processes. Thus, drive letters, logical volume names or device labels of devices are unassigned. This command terminates with an error if a drive letter, logical volume name or device label of a transaction volume is assigned or if:
 - The mount point has directories to which other volumes are also mounted.
 - The volume has a file in use.
- To mount a transaction volume from a server other than the server on which this command is entered, unmount the transaction volume, using the authority of the backup administrator. If the transaction volume must not be unmounted, create backup copies with the following procedure:
 1. Execute the sync command to synchronize the transaction volume and file system.
 2. Execute this command.
 3. Execute the fsck command for the backup volume, and check the file system.
- Backup cannot be executed in the following situations:
 - No backup policy has been specified for the specified transaction volume.
 - The backup that was executed one generation earlier was unsuccessful.
 - Restore is being performed to the transaction volume that is to be backed up.
 - Synchronous backup processing is performed in a transaction volume to be processed, but the state of equivalence is not reached;
 - Restoration is being executed using backup history information that causes a generation overflow (this may occur when restoration is performed using the backup volume of the oldest history).
 - Any of the box information (box identifier, OLU number, EXTENT start position, and EXTENT size) of a transaction volume has been changed.
 - No unused backup volume was available for use as the copying destination.
 - The replication management function is executing a process for the target transaction volume.

- The replication management function is executing a process for the target backup volume.
 - The replication management function is executing a process for backup records exceeding the number of preservation generations.
 - The mount point of a transaction volume set at registration has been changed.
 - A backup volume defined as the copy destination is mounted.
 - A device containing SynfinityFILE resources consisting of multiple transaction volumes is mounted.
 - When the transaction volume or backup volume is a logical volume of VxVM, or the structure of the logical volume is changed to a structure not supported by AdvancedCopy Manager.
 - When the transaction volume or backup volume is a logical volume of VxVM, or the logical volume is deleted from VxVM.
 - When the -T option is specified during the operation in which SDX objects are backed up in units of logical volumes
 - When either the transaction volume or the backup volume is changed to TFOV.
- You cannot perform backup if:
- The backup volume specified by the device map file cannot be paired with the transaction volume. This occurs if:
 - The specified backup volume is in use.
 - The transaction volume and the backup volume are not the same partition size.
 - The device map file contains an error (refer to "3.4.9.1 Describing Device Map File").
 - The device map file name includes a national character.
 - When an ordinary volume is specified as the transaction volume and a PRIMECLUSTER GD logical volume as the backup volume.
 - When a PRIMECLUSTER GD logical volume is specified as the transaction volume and an ordinary volume as the backup volume.
- At this time, if existing backup history information satisfies the condition of the number of saved generations specified by the [swstbkpolset](#) command, then information on the oldest generation is automatically deleted from the backup history information during the backup process, and the backup volume is cleared. However, if -suspend or -T option is specified for the backup history information, then the backup volume with the Suspend or tracking state is not cleared and the backup volume remains in use.
- The transaction volume on which the backup command is executed, one other transaction volume, and the suspended or tracking backup volume cannot be specified as backup destinations in the device map file.
- In the case of a SynfinityFILE file system whose transaction volume consists of more than one device (multi-partition), backup cannot be performed the unmount/mount tasks cannot be performed inside the pre-processing and post-processing scripts respectively.
- In this case, manually unmount the file system before backup is performed and re-mount it after the backup has completed.

- When the SynfinityFILE and AdvancedCopy Manager configurations are as follows:

```
# sfxinfo /dev/dsk/c1t3d10s1
ID      special                                size  mount  Type
0       /dev/dsk/c1t3d10s1(800019)              25986 /mnt   META
0       /dev/dsk/c1t3d10s1(800019)              105751 /mnt   DATA
1       /dev/dsk/c1t3d15s3(800028)              5120  /mnt   LOG
2       /dev/dsk/c1t3d16s3(800030)              131736 /mnt   DATA

# /opt/FJSVswsts/bin/swstdevdisp -t
Server Device          Size          Device-Mode Mount-Point (Method) Backup-Engine
SV1     /dev/dsk/c1t3d10s1  128.0 Mbyte Transaction /mnt (ufs)     AdvancedCopy
SV1     /dev/dsk/c1t3d15s3    5.0 Mbyte Transaction /mnt (ufs)     AdvancedCopy
SV1     /dev/dsk/c1t3d16s3  128.0 Mbyte Transaction /mnt (ufs)     AdvancedCopy

#
```

- Backup procedure

```
# umount /mnt
# /opt/FJSVswsts/bin/swstbackup /dev/dsk/c1t3d10s1
```

```

/dev/dsk/clt3d10s1 swstbackup completed
# /opt/FJSVswsts/bin/swstbackup /dev/dsk/clt3d15s3
/dev/dsk/clt3d15s3 swstbackup completed
# /opt/FJSVswsts/bin/swstbackup /dev/dsk/clt3d16s3
/dev/dsk/clt3d16s3 swstbackup completed
# mount -F sfxfs /dev/dsk/clt3d10s1 /mnt
#

```

- For details on how to use SDX objects of PRIMECLUSTER GD, refer to "[Configuration and Conditions of Unsupported SDX Objects](#)" and "[10.1.5.1 Notes on SDX Object Operations \(Backup Management\)](#)".
- The pre-processing and post-processing scripts are not executed when the mirror slice of the SDX object is TEMP.
- This command can be executed only while none of the following commands are running:
 - [swstsvrset](#)
 - [swstdevinfo](#)
 - [swstsrsemth](#)
 - [swstresback](#)
 - [swstresrst](#)

9.1.2.2 swstbackstat (Backup Execution Status Display Command)

NAME

swstbackstat - displays the completion status of a backup

SYNOPSIS

To be executed on Managed Server

```
/opt/FJSVswsts/bin/swstbackstat [deviceName]
```

To be executed on Management Server

```
/opt/FJSVswsts/bin/swstbackstat [-h serverName] [deviceName]
```

DESCRIPTION

This command displays the completion status of a backup that is being performed using the [swstbackup](#) command.

The following information is displayed.

Title	Description
Server	Displays a Managed Server name.
Transaction-Disk	Displays the AdvancedCopy Manager device name of a transaction volume.
Backup-Disk	Displays a backup volume name. If no backup has been collected, "----" is displayed.
Status	Displays a backup completion status using one of the following character strings: "----" : No backup has been collected (ie, no history exists). "succeeded" : Backup has been completed. "executing" : Copying using OPC is in progress. "failed" : Copying using OPC has been interrupted due to an error. "halt" : Copying using OPC is halt status. "gds-error" : Copying with PRIMECLUSTER GD ended abnormally. "?????" : Failed to get copy status.

Title	Description
Execute	Displays how much has been copied in percentage terms if "executing" is displayed in the Status column. Displays "----" unless "executing" is displayed.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdispsrv](#) command.

OPERANDS

deviceName

Specifies the "1.4 Managing Device on AdvancedCopy Manager" corresponding to a transaction volume.

If this operand is omitted, the backup execution statuses of all the transaction volumes are displayed.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Display the actual copying execution status:

```
# /opt/FJSVswsts/bin/swstbackstat
Server Transaction-Disk Backup-Disk Status Execute
Work1 /dev/dsk/clt0d0s3 /dev/dsk/clt0d2s3 succeeded ----
Work1 /dev/dsk/clt0d0s4 /dev/dsk/clt0d2s4 executing 70%
Work1 /dev/dsk/clt0d0s5 /dev/dsk/clt0d2s5 executing 30%
:
#
```

NOTES

- Do not use the OPC stop function of ETERNUS Web GUI while the snapshot fast backup is in progress. If you do use the OPC stop function of ETERNUS Web GUI, it displays "succeeded" whether or not the backup has actually been successful.
- If "failed", "halt", or "?????" is displayed in the Status column, a hardware error, etc. may have occurred. Refer to [8.3 Troubleshooting: Hardware Error During Backup](#) for information about this.
- Backup execution status display cannot be executed in the following circumstances:
 - When the transaction volume or backup volume is a logical volume of VxVM, the structure of the logical volume is changed to a structure not supported by AdvancedCopy Manager.
 - When the transaction volume or backup volume is a logical volume of VxVM, the logical volume is deleted from VxVM.
- If "gds-error" is displayed in the Status column, a problem with PRIMECLUSTER GD is assumed. Remedy the cause of the PRIMECLUSTER GD problem, and take one of the following steps:
 - If backup synchronous processing is in progress, execute the [swstcancelsync](#) command to cancel the backup synchronous processing.
 - If snapshot backup is in progress, execute the [swsthistdel](#) command to delete the backup history information.
 - If restoration is in progress, execute the [swstsrsemitch](#) command and perform restoration again.

9.1.2.3 swstrestore (Restore Execution Command)

NAME

swstrestore - restore

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsts/bin/swstrestore [-g generationNo | -v versionNo] [-r restoreDeviceName] [-Xgds-softcopy] deviceName
```

To Execute on Management Server

```
/opt/FJSVswsts/bin/swstrestore [-h serverName] [-g generationNo | -v versionNo] [-r restoreDeviceName] [-Xgds-softcopy] deviceName
```

DESCRIPTION

The data that exists in backup history information is restored using OPC.

When this command is executed, AdvancedCopy Manager performs the restoration pre-processing and post-processing scripts before and after copying data using OPC. For more information on customizing these scripts, refer to "[Appendix A Pre-processing and Post-processing of Backup and Restoration](#)".

You can restore data to a device that is not a transaction volume. When you do so, the device must not be registered as a transaction volume and the device must have the same amount of space available as the specified transaction volume.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdisprv](#) command.

-g *generationNo*

Specifies the relative generation number of data to be restored.

Check the relative generation number using the [swsthistsdisp](#) command.

-v *versionNo*

Specifies the absolute generation number of data to be restored.

Check the absolute generation number using the [swsthistsdisp](#) command.

-r *restoreDeviceName*

Specifies a restore destination device name to restore data to any device other than a transaction volume.

Cannot be specified for a PRIMECLUSTER GD logical volume.

-Xgds-softcopy

Performs restoration by software-copy for restoration in linkage with PRIMECLUSTER GD Snapshot. If this option is not specified, restoration is performed with OPC.

Note: If neither the -g or -v option is specified, the latest information from the information history management is restored.

OPERANDS

deviceName

Specifies the AdvancedCopy Manager device name corresponding to a transaction volume.

Refer to "[1.4 Managing Device on AdvancedCopy Manager](#)" for the devices to be specified.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Restore data with relative generation number 2. For information on the relative generation number, refer to data displayed using the `swsthistdisp` command:

```
# /opt/FJSVswsts/bin/swstrestore -g 2 /dev/dsk/clt0d0s6
/dev/dsk/clt0d0s6 swstrestore completed
#
```

- If the "Chapter 4 Backup Operation Using Suspend/Resume Function" is used for the backup operation on the transaction volume, the suspended synchronous processing is deleted and the restore command is executed. For information on the suspended synchronous processing, refer to the data displayed by the `swstsyncstat` command. Execute the `swstcancelsync` command to cancel the synchronous processing:

```
# /opt/FJSVswsts/bin/swstsyncstat /dev/dsk/clt0d0s6
Server Transaction-Disk Backup-Disk Status Execute
Work2 /dev/dsk/clt0d0s6 /dev/dsk/clt0d2s5 suspend ----
Work2 /dev/dsk/clt0d0s6 /dev/dsk/clt0d2s6 executing 75%
# /opt/FJSVswsts/bin/swstcancelsync -all /dev/dsk/clt0d0s6
/dev/dsk/clt0d0s6 swstcancelsync completed
# /opt/FJSVswsts/bin/swstsyncstat /dev/dsk/clt0d0s6
Server Transaction-Disk Backup-Disk Status Execute
Work2 /dev/dsk/clt0d0s6 ---- ---- ----
# /opt/FJSVswsts/bin/swstrestore /dev/dsk/clt0d0s6
/dev/dsk/clt0d0s6 swstrestore completed
#
```

NOTES

- On a transaction volume on which a file system has been constructed, restore a specific file as follows:
 1. Mount a backup volume. Check the backup volume by executing the Backup History List view/`swsthistdisp` command.
 2. Copy the file to be restored by executing, for example, the `cp` command.
 3. Unmount the backup volume.

For the SynfinityFILE file system, follow the procedure below:

1. Execute the following commands for a backup volume:

```
# sfxadm <RAW-device-name-of-the-backup-volume>
# fsck -F sfxfs -o f <RAW-device-name-of-the-backup-volume>
```

2. Mount the backup volume:

```
# mount -F sfxfs <backup-volume-name> <mount-point-name>
```

3. Copy the file to be restored by executing, for example, the `cp` command:

4. Unmount the backup volume:

```
# umount <backup-volume-name> or <mount-point-name>
```

- 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.
- Restoration or recovery cannot be executed in the following situations:
 - When both relative and absolute generation numbers have been specified.
 - When a device registered as a transaction volume has been specified in a parameter with the `-r` option.
 - When a device registered as a backup volume has been specified in a parameter with the `-r` option, the device has been registered in the backup registration information of a certain transaction volume. If you restore data to a backup volume without registered backup data, and the backup volume is used to back up a certain transaction volume, then data is overwritten. It is recommended that a backup volume should NEVER be specified as a restore destination device.

- When you perform restoration immediately after starting backup where the actual copying using OPC is still in progress, (except when the combination of transaction volume and backup volume is the same for backup and restoration). A restoration should always be performed after the actual copying using OPC is completed. Alternatively, on a transaction volume on which a file system has been constructed, perform restoration according to the method described in the Point "In a transaction volume on which a file system has been constructed, restore a specific file" above.
- A transaction volume currently in use with synchronous backup processing is specified.
- When any of the box information (box identifier, OLU number, EXTENT start position, and EXTENT size) of a transaction volume to be restored has been changed.
- The replication management function is executing a process for the target transaction volume.
- The replication management function is executing a process for the device specified in the -r option.
- The mount point set at registration of a transaction volume has been changed.
- A PRIMECLUSTER GD logical volume is specified as the restoration destination volume of an ordinary transaction volume.
- When the transaction volume or backup volume is a logical volume of VxVM, the structure of the logical volume is changed to a structure not supported by AdvancedCopy Manager.
- When the transaction volume or backup volume is a logical volume of VxVM, the logical volume is deleted from VxVM.
- A slice including VTOC is specified for the -r option.
- The -r option is specified for recovery of a PRIMECLUSTER GD logical volume.
- If the SynfinityFILE file system is already mounted before restoration of devices of the SynfinityFILE file system consisting of multiple partitions, unmount the SynfinityFILE file system, and then restore all devices that make up the SynfinityFILE file system.
- The -Xgds-softcopy option is specified for a device that is not a PRIMECLUSTER GD logical volume.
- A PRIMECLUSTER GD logical volume is specified as the restoration destination volume of an ordinary transaction volume.
- When an SDX object has been specified in a parameter with the -r option.
- When either of the transaction volume or the backup volume is changed to TFOV.
- The suspended or tracking backup volume cannot be used as the restore destination volume.
- The restore command cannot be executed on a transaction volume during synchronous processing (eg, during copying by EC, under the equivalency maintained status, in the Suspend state).
If the following errors occur during restore processing, then execute the [swstcancelsync](#) command (with the -all option specified) to cancel all of the synchronous processing of the transaction volume, and execute the restore command:
 - "swst0639 This Command cannot be run. This is because a copy is currently running"
 - "swst0634 OPC cannot be run. Error code=26 (SCSI command busy)"
- If the transaction volume is an SDX object and a non-single volume, you cannot perform restoration using this command. If so, perform recovery as follows:
 1. The backup volume is confirmed by backup history list view or [swsthistsdisp](#) command.
 2. copies onto the volume which is restore in the dd command
- In the case of a SynfinityFILE file system that consists of more than one device (multipartition), backup cannot be performed by using pre-processing and post-processing scripts to execute unmount/mount.
In the case of a SynfinityFILE file system that consists of more than one device (multipartition), unmount the file system before restoration is performed and mount it after that restoration is completed.
- When the SynfinityFILE and AdvancedCopy Manager configurations are as follows:

```
# sfxinfo /dev/dsk/c1t3d10s1
ID      special                                size  mount  Type
0       /dev/dsk/c1t3d10s1(800019)             25986 /mnt   META
0       /dev/dsk/c1t3d10s1(800019)             105751 /mnt   DATA
1       /dev/dsk/c1t3d15s3(800028)              5120  /mnt   LOG
```

```

2      /dev/dsk/clt3d16s3(800030)      131736 /mnt      DATA

# /opt/FJSVswsts/bin/swstdevdisp -t
Server Device          Size          Device-Mode Mount-Point (Method) Backup-Engine
SV1      /dev/dsk/clt3d10s1  128.0 Mbyte Transaction /mnt (ufs)      AdvancedCopy
SV1      /dev/dsk/clt3d15s3   5.0 Mbyte Transaction /mnt (ufs)      AdvancedCopy
SV1      /dev/dsk/clt3d16s3  128.0 Mbyte Transaction /mnt (ufs)      AdvancedCopy

```

```
#
```

- Restore procedure

```

# umount /mnt
# /opt/FJSVswsts/bin/swstrestore /dev/dsk/clt3d10s1
/dev/dsk/clt3d10s1 swstrestore completed
# /opt/FJSVswsts/bin/swstrestore /dev/dsk/clt3d15s3
/dev/dsk/clt3d15s3 swstrestore completed
# /opt/FJSVswsts/bin/swstrestore /dev/dsk/clt3d16s2
/dev/dsk/clt3d16s3 swstrestore completed
# mount -F sfxfs /dev/dsk/clt3d10s1 /mnt
#

```

- When the SynfinityFILE file system is restored, the restore destination device may not be able to be mounted. In this event, use the following procedure to mount the restore destination device:

1. Execute the following on the restore destination device:

```
# sfxadm <raw-device-name-of-restore-destination-device>
```

2. Mount the restore destination device as follows:

```
# mount -F sfxfs <restore-destination-device-name mount-point-name>
```

- For details on how to use SDX objects of PRIMECLUSTER GD, refer to "[Configuration and Conditions of Unsupported SDX Objects](#)" and "[10.1.5.1 Notes on SDX Object Operations \(Backup Management\)](#)".

- This command can only be executed while none of the following commands are running:

- [swstsvrset](#)
- [swstdevinfo](#)
- [swstsrsem](#)
- [swstresback](#)
- [swstresrst](#)

9.1.2.4 swstcancelrest (Restore Cancellation Command)

NAME

swstcancelrest - cancels OPC-based restoration

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsts/bin/swstcancelrest [-g generationNo | -v versionNo] [-r restoreDeviceName] [-emergency] deviceName
```

To Execute on Management Server

```
/opt/FJSVswsts/bin/swstcancelrest [-h serverName][ -g generationNo | -v versionNo][ -r restoreDeviceName] [-emergency] deviceName
```

DESCRIPTION

This command cancels OPC-based restoration.

A hardware error is displayed when the [swstrestore](#) command is executed while data is being copied using OPC. Fix the cause of the hardware error and repeat restoration using the [swstrestore](#) command.

Check whether a fault occurred using the [swstreststat](#) command.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdispsrv](#) command.

-g *generationNo*

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 [swstreststat](#) command.

-v *versionNo*

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 [swstreststat](#) command.

-r *restoreDeviceName*

Specifies the device name if data is restored to any other device than a transaction volume.

Check the device name using the [swstreststat](#) command.

-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 is cancelled if neither the **-g** or **-v** option is specified.

OPERANDS

deviceName

Specify an AdvancedCopy Manager device name.

Refer to "[1.4 Managing Device on AdvancedCopy Manager](#)" for the devices to be specified.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Cancel restoration of the transaction volume:

```
# /opt/FJSVswsts/bin/swstcancelrest -g 1 /dev/sdb6
/dev/sdb6 swstcancelrest completed.
#
```

- Restoration of the latest history information on the transaction volume is cancelled in emergency operation mode:

```
# /opt/FJSVswsts/bin/swstcancelrest -emergency /dev/sdb6
/dev/sdb6 swstcancelrest completed.
#
```

NOTES

- If you cancel OPC-based restoration using this command, the transaction volume is in an incomplete state and can no longer be used because part of the data has already been overwritten by the copy process. Perform restoration again using the [swstrestore](#) command.

If the transaction volume is a PRIMECLUSTER GD logical volume and restoration is cancelled, the state of the transaction volume becomes INVALID. In such cases, refer to the *PRIMECLUSTER Global Disk Services User's Guide (Solaris(TM) Operating System)*, and restore the logical volume.

- You cannot cancel restoration if:
 - The history specified in the -g or -v option does not exist.
 - Restoration is not performed on the device specified in the -r option.
 - Restoration is not performed from the specified history. In this case, processing is terminated with an information message "swst0303 Restoration for specified restoration device is not performed". The end status after this message is a normal end.
 - The replication management function is executing a process for the target transaction volume
- However, OPC sessions are neither confirmed nor cancelled by execution in emergency operation mode. If an OPC session is in the error suspended state, then unmanaged sessions remain. In this case, ETERNUS Web GUI must be used to cancel these sessions.
- If a logical volume of VxVM is deleted or its structure is changed to a structure not supported by AdvancedCopy Manager, restoration cannot be cancelled. In such cases, take action according to either item a) or b).
 - a. Make sure that Advanced Copy is not in progress for the replication volume to be deleted, and execute this command with the -emergency option specified. If Advanced Copy is in progress, stop Advanced Copy processing by using ETERNUS Disk storage system or take action according to item b).
 - b. Re-create the logical volume or return the logical volume structure to its original state, and then execute this command.
- This command can only be executed while none of the following commands are running:
 - [swstsvrset](#)
 - [swstdevinforeset](#)
 - [swstsrsematch](#)
 - [swstresback](#)
 - [swstresrst](#)

9.1.2.5 swstreststat (Restore Execution Status Display Command)

NAME

swstreststat - displays the execution status of restoration

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSSVswsts/bin/swstreststat [-g generationNo | -v versionNo] [-Xgds-softcopy] [deviceName]
```

To Execute on Management Server

```
/opt/FJSSVswsts/bin/swstreststat [-h serverName] [-g generationNo | -v versionNo] [-Xgds-softcopy] [deviceName]
```

DESCRIPTION

This command displays the execution status of restoration.
The following information is displayed.

Title	Description
Server	Displays a Managed Server name.
Transaction-Disk	Displays the AdvancedCopy Manager device name of a transaction volume.
Generation	Displays the relative generation number of restored backup data. "----" is 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. "----" is 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. "----" is displayed if restoration is not in progress.
Restore-Device	Displays a restoration target device name unless data is copied and restored to a transaction volume. "----" is displayed if restoration is not in progress, or if data is restored to a transaction volume.
Status	Displays the execution status: "----" : No copying using OPC or software-copy by PRIMECLUSTER GD is in progress. "executing" : Copying using OPC or software-copy by PRIMECLUSTER GD is in progress. "succeeded" : Copying using OPC or software-copy by PRIMECLUSTER GD has been completed. "failed" : Copying using OPC has been interrupted due to an error. "halt" : Copying using OPC is halt status. "gds-error" : Copying with PRIMECLUSTER GD ended abnormally. "?????" : Failed to get copy status.
Execute	Displays how much has been copied in percentage terms if "executing" is displayed in the Status column. Displays "----" unless "executing" is displayed.

OPTIONS

-h serverName

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdisprv](#) command.

-g generationNo

Specifies the relative generation number of data whose restore execution status is to be displayed.

-v versionNo

Specifies the absolute generation number of data whose restore execution status is to be displayed.

-Xgds-softcopy

Displays the restore execution status of a PRIMECLUSTER GD logical volume.

Use this option to check the copy status when restoration using the software-copy function of PRIMECLUSTER GD is in progress.

When the option is specified, the Status column display when copying is completed is "----", which is the same as before copying started.

Note: The restore execution statuses of all the histories are displayed if both the -g and -v options are omitted.

OPERANDS

deviceName

Specify an AdvancedCopy Manager device name.

Refer to "[1.4 Managing Device on AdvancedCopy Manager](#)" for the devices to be specified.

If this operand is omitted, the restore execution statuses of all the transaction volumes are displayed.

EXIT STATUS

- =0: Completed successfully
- >0: Terminated abnormally

EXAMPLES

- Display the restore execution statuses of transaction volumes:

```
# /opt/FJSVswsts/bin/swstreststat
Server Transaction-Disk Generation Version Backup-Disk Restore-Device Status Execute
Work2 /dev/dsk/clt0d0s3 1 10 /dev/dsk/clt0d2s3 /dev/dsk/clt0d1s3 executing 75%
Work2 /dev/dsk/clt0d0s4 ---- ---- ---- ---- ----
Work2 /dev/dsk/clt0d0s5 2 12 /dev/dsk/clt0d2s5 ---- succeeded ---
Work2 /dev/dsk/clt0d0s6 3 13 /dev/dsk/clt0d2s6 /dev/dsk/clt0d1s6 executing 75%
:
#
```

NOTES

- Do not use the OPC stop function of ETERNUS Web GUI while the snapshot fast backup is in progress. If you use the OPC stop function of ETERNUS Web GUI, this command displays "succeeded" whether or not the restoration has actually been successful.
- If "failed", "halt", or "?????" is displayed in the Status column, a hardware error, etc. may have occurred. Refer to [8.3 Troubleshooting: Hardware Error During Backup](#) for information about this.
- If no generation number has been specified in the -v or -g option, the restore execution statuses of all the histories are displayed. If, at this time, restoration is not in progress from any of the histories, "----" is displayed in all the information columns other than "Server" and "Transaction-Disk".
- In the following cases, restore execution status display processing cannot be executed:
 - When the transaction volume or backup volume is a logical volume of VxVM, the structure of the logical volume is changed to a structure not supported by AdvancedCopy Manager.
 - When the transaction volume or backup volume is a logical volume of VxVM, the logical volume is deleted from VxVM.
- If "gds-error" is the status displayed in the Status column, a problem with PRIMECLUSTER GD is assumed. Correct the cause of the PRIMECLUSTER GD problem, and take action as follows:
 - If backup synchronous processing is in progress, execute the [swstcancelsync](#) command to cancel the backup synchronous processing.
 - If snapshot backup is in progress, execute the [swsthistdel](#) command.
 - If restoration is in progress, execute the [swstsrsemitch](#) command and perform restoration again.

9.1.2.6 swsthistdisp (History Information Display Command)

NAME

swsthistdisp - displays the backup history information

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsts/bin/swsthistdisp [deviceName]
```

To Execute on Management Server

```
/opt/FJSVswsts/bin/swsthistdisp [-h serverName] [deviceName]
```

DESCRIPTION

This command displays history information in relation to data that has been backed up.

The following information is displayed.

Title	Description
Server	Displays a Managed Server name.
Device	Displays the AdvancedCopy Manager device name .
Mount-Point	Displays the mount point of a device. The file system type of the mount point is displayed inside the parentheses.
Generation	Displays a relative generation number. This item is not displayed if no backup history information exists.
Version	Displays an absolute generation number. This item is not displayed if no backup history information exists.
Backup-Date	Displays a TimeStamp reflecting when backup has been completed. This item is not displayed if no backup history information exists.
Backup-Disk	Displays the name of a backup volume to which data has been backed up. This item is not displayed if no backup history information exists.
Status	Displays an actual copying execution status using one of the following character strings: "succeeded" : Copying using OPC has been completed. "executing" : Copying using OPC is in progress. "failed" : Copying using OPC has been interrupted due to an error. "halt" : Copying using OPC is halt status. "gds-error" : Copying with PRIMECLUSTER GD ended abnormally. "?????" : Failed to get copy status. This item is not displayed if no backup history information exists.
Execute	Displays how much has been copied in percentage terms if "executing" is displayed in the Status column. Displays "----" unless "executing" is displayed. This item is not displayed if no backup history information exists.

OPTIONS

`-h serverName`

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdispsrv](#) command.

OPERANDS

`deviceName`

Specify an AdvancedCopy Manager device name.

Refer to "[1.4 Managing Device on AdvancedCopy Manager](#)" for the devices to be specified.

If this operand is omitted, the backup history information of all the transaction volumes is displayed.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Display the backup history information of all the transaction volumes:

```
# /opt/FJSVswsts/bin/swsthistdisp
Server=job1 Device=/dev/dsk/c1t0d0s6 Mount-Point=/mnt/Tran1 (ufs)
```

Generation	Version	Backup-Date	Backup-Disk	Status	Execute
1	10	2000/11/12 22:00	/dev/dsk/c1t0d2s6	succeeded	----
2	9	2000/11/11 22:00	/dev/dsk/c1t0d4s6	succeeded	----
:					
#					

NOTES

- If "failed", "halt", or "?????" is displayed in the Status column, a hardware error, etc. may have occurred. Refer to ["8.3 Troubleshooting: Hardware Error During Backup"](#) for information about this.
- In the following cases, history information display processing cannot be executed:
 - When the transaction volume or backup volume is a logical volume of VxVM, the structure of the logical volume is changed to a structure not supported by AdvancedCopy Manager.
 - When the transaction volume or backup volume is a logical volume of VxVM, the logical volume is deleted from VxVM.
- If "gds-error" is the status displayed in the Status column, a problem with PRIMECLUSTER GD is assumed. Correct the cause of the PRIMECLUSTER GD problem, and take action as follows:
 - If backup synchronous processing is in progress, execute the [swstcancelsync](#) command to cancel the backup synchronous processing.
 - If snapshot backup is in progress, execute the [swsthistdel](#) command.
 - If restoration is in progress, execute the [swstrsemitch](#) command and perform restoration again.

9.1.2.7 swsthistdel (History Information Deletion Command)

NAME

swsthistdel - deletes backup history information

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsts/bin/swsthistdel -g generationNo | -v versionNo | -z [-emergency] deviceName
```

To Execute on Management Server

```
/opt/FJSVswsts/bin/swsthistdel [-h serverName] -g generationNo | -v versionNo | -z [-emergency] deviceName
```

DESCRIPTION

This command deletes backup history information.
The backup volume of the deleted history is released and becomes an empty backup volume.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.
When executing on the target Managed server, it is not necessary to specify this option.
For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdisprv](#) command.

-g *generationNo*

Specifies the relative generation number of backup history information to be deleted.
This option cannot be specified at the same time as the -v or -z option.
Check the relative generation number using the [swsthistdisp](#) command.

-v versionNo

Specifies the absolute generation number of backup history information to be deleted.
This option cannot be specified at the same time as the *-g* or *-z* option.
Check the absolute generation number using the [swsthistdisp](#) command.

-z

Specifies that all the backup history information is to be deleted.
This option cannot be specified at the same time as the *-g* or *-v* option.

-emergency

Emergency operation mode is available for operations.
In such cases, OPC sessions are neither confirmed nor cancelled.

OPERANDS

deviceName

Specify an AdvancedCopy Manager device name.
Refer to "[1.4 Managing Device on AdvancedCopy Manager](#)" for the devices to be specified.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Delete a history with absolute generation number 10 of a transaction volume:

```
# /opt/FJSVswsts/bin/swsthistdel -v 10 /dev/dsk/c1t0d0s6  
/dev/dsk/c1t0d0s6 swsthistdel completed  
#
```

- The history information of relative generation number 1 of the transaction volume is deleted in the emergency operation mode:

```
# /opt/FJSVswsts/bin/swsthistdel -g 1 -emergency /dev/dsk/c1t0d0s6  
/dev/dsk/c1t0d0s6 swsthistdel completed  
#
```

NOTES

- You cannot delete backup history information if:
 - The backup history information specified in the *-g* or *-v* option does not exist.
 - The specified backup history information is currently being used in a restoration.
 - The replication management function is executing a process for the target transaction volume.
 - The replication management function is executing a process for the target backup volume.
- However, OPC sessions are neither confirmed nor cancelled by execution in emergency operation mode. If an OPC session is in the error suspended state, then unmanaged sessions remain. In this case, ETERNUS Web GUI must be used to cancel these sessions.
- Deleting the suspended history information does not cancel the synchronous processing (in an EC session).
Deleting the tracking history information does not cancel the tracking processing (in an OPC session).
- In synchronous high-speed backup with *-suspend* option specified, the backup volume remains in use even after the history information is deleted.
In differential snapshot high-speed backup with *-T* option specified, the backup volume remains in use even after the history information is deleted.
- If the history information is deleted while the OPC physical copy has not been completed in a backup operation mode other than the differential snapshot high-speed backup, physical copying of the deleted history is interrupted. In this case, the data in the backup volume is incomplete (ie, copy has stopped part way through). The backup volume is put in the unused state upon deletion

of the history information and becomes ready for use in the next backup.

In the differential snapshot high-speed backup mode, the physical copying of the deleted history is continued even after the history information on the incomplete OPC physical copy is deleted. Even after deletion of the history information, the backup volume is in the busy state and becomes ready for use in the next backup of the relevant transaction volume.

- If a logical volume of VxVM is deleted or its structure is changed to a structure not supported by AdvancedCopy Manager, restoration cannot be cancelled. In such cases, take action according to either item a) or b).
 - a. Make sure that Advanced Copy is not in progress for the replication volume to be deleted, and execute this command with the -emergency option specified. If Advanced Copy is in progress, stop Advanced Copy processing by using the ETERNUS Disk storage system).
 - b. Re-create the logical volume or return the logical volume structure to its original state, and then execute this command.
- For details on how to use SDX objects of PRIMECLUSTER GD, refer to "[Configuration and Conditions of Unsupported SDX Objects](#)" and "[10.1.5.1 Notes on SDX Object Operations \(Backup Management\)](#)".
- This command can be executed only while none of the following commands are running:
 - `swstsvrset`
 - `swstdevinfo`
 - `swstsrsem`
 - `swstresback`
 - `swstresrst`

9.1.2.8 swststartsync (Backup Synchronous Processing Start Command)

NAME

swststartsync - starts backup synchronous processing

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSSVswsts/bin/swststartsync [-Xgds-softcopy | -Xgds-selectcopy] deviceName [-Xdevmap deviceMapFile]
```

To Execute on Management Server

```
/opt/FJSSVswsts/bin/swststartsync [-Xgds-softcopy | -Xgds-selectcopy] [-h serverName] deviceName [-Xdevmap deviceMapFile]
```

DESCRIPTION

This command starts backup synchronous processing (ie, it starts disk-to-disk copy using EC).

Even if it is already in use, the transaction volume where the synchronous processing started as well as the suspended backup volume can be used as backup destinations.

Except for the transaction volume where the synchronous processing started and the suspended backup volume, transaction volumes cannot be specified as backup destinations in the device map file.

For a PRIMECLUSTER GD transaction volume, synchronous processing by software-copy is also possible.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdispsrv](#) command.

`-Xdevmap deviceMapFile`

Specifies a target backup volume to perform the backup synchronous processing. In the operand, specify a [device map file](#) in which a combination of a transaction volume and a target backup volume is described. If this option is not specified, a target backup volume is automatically selected. To execute this command on a remote basis using the `-h` option, specify a device map file of a Managed Server using an absolute pathname. A file name specified in *Device-Map-File* must not include any national characters.

`-Xgds-softcopy`

Specifies the start of backup synchronous processing of an SDX object and use of the software-copy function instead of the Advanced Copy function.

If you specify this option to start differential copying when the Advanced Copy function is being used, Advanced Copy processing stops and software-copy processing starts.

You can specify the option when the function in linkage with PRIMECLUSTER GD Snapshot is being used.

`-Xgds-selectcopy`

Specifies that PRIMECLUSTER GD selects the copy processing method of an SDX object.

You can specify the option when the function in linkage with PRIMECLUSTER GD Snapshot is being used.

OPERANDS

deviceName

Specify an AdvancedCopy Manager device name.

Refer to "[1.4 Managing Device on AdvancedCopy Manager](#)" for the devices to be specified.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- To start the backup synchronous processing for a transaction volume:

```
# /opt/FJSVswsts/bin/swststartsync /dev/dsk/clt0d0s6  
/dev/dsk/clt0d0s6 swststartsync completed.  
#
```

- Backup synchronous processing of a logical volume (`/dev/sfdsk/CLS01/dsk/VOL01`) of an SDX object is started. The copy processing method is software-copy:

```
# /opt/FJSVswsts/bin/swststartsync -Xgds-softcopy /dev/sfdsk/CLS01/dsk/VOL01  
/dev/sfdsk/CLS01/dsk/VOL01 swststartsync completed.  
#
```

NOTES

- When you execute the backup synchronous processing start command for a transaction volume to which the backup synchronous processing is already being executed, the message "swst0301 Backup synchronous processing is already in progress." is output.
- You cannot start backup synchronous processing in the following cases. In any of the cases below, perform the action according to the [System Administrator Corrective Action] message.
 - When specifying a device that does not support the EC function. (no linkage with PRIMECLUSTER GD Snapshot)
 - When the backup volumes required for starting the backup synchronous processing cannot be obtained.
 - When specifying a device that is being restored by OPC.
 - When mounting the backup volume to which data is copied.
 - When the transaction volume or backup volume is a logical volume of VxVM, the structure of the logical volume is changed to a structure not supported by AdvancedCopy Manager.
 - When the transaction volume or backup volume is a logical volume of VxVM, the logical volume is deleted from VxVM.

- The -Xgds-softcopy option is specified for a device that is not a PRIMECLUSTER GD logical volume. Alternatively, the -Xgds-selectcopy option is specified for a device that is not a PRIMECLUSTER GD logical volume.
- The replication management function is executing a process for the target transaction volume.
- When either of the transaction volume or the backup volume is changed to TFOV.
- If a suspended backup volume is the transaction volume at the start of the synchronous processing, then the suspended backup volume is selected as the backup destination with the highest priority, and the synchronous processing restarts (ie, differential copy) starts. Otherwise, the synchronous processing (ie, full copy) starts.
- Backup synchronous processing with the output destination backup volume specified cannot be started in the following circumstances:
 - The backup volume specified in the device map file cannot be used as the backup volume to be paired with the transaction volume. This means that either of the following conditions is not satisfied:
 - The specified backup volume is in the unused state.
 - The transaction volume and backup volume have the same partition size.
 - The target transaction volume is incorrectly described in the device map file (Refer to "[3.4.9.1 Describing Device Map File](#)").
 - Any special or foreign characters, operators or reserved words are used in the device map file names.
 - An ordinary volume is specified as the transaction volume and a PRIMECLUSTER GD logical volume as the backup volume.
 - A PRIMECLUSTER GD logical volume is specified as the transaction volume and an ordinary volume as the backup volume.
- For details on how to use SDX objects of PRIMECLUSTER GD, refer to "[Configuration and Conditions of Unsupported SDX Objects](#)" and "[10.1.5.1 Notes on SDX Object Operations \(Backup Management\)](#)".
- This command cannot be executed when any of the following commands are running:
 - `swstsvrset`
 - `swstdevinforeset`
 - `swstsrsematch`
 - `swstresback`
 - `swstresrst`

9.1.2.9 swstcancelsync (Backup Synchronous Processing Cancel Command)

NAME

swstcancelsync - cancels the backup synchronous processing

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsts/bin/swstcancelsync [-emergency] [-bd Backup-Device-Name | -all] Device-Name
```

To Execute on Management Server

```
/opt/FJSVswsts/bin/swstcancelsync [-h serverName] [-emergency] [-bd backupDeviceName | -all] deviceName
```

DESCRIPTION

The backup synchronous processing can be cancelled during copying by EC or under the equivalency maintained status.

The synchronous processing can be cancelled with the backup volume during copying by EC, under the equivalency maintained status, or while the Suspend state is set. Specifying cancellation of all of the synchronous processing cancels all of the backup synchronous processing including those in the Suspended state.

If the system detects an error while backup synchronous processing is being executed, copying in progress at that time terminates abnormally. In this case, use this command to cancel the backup synchronous processing. Use the [swstsyncstat](#) command to check for the error.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdispsrv](#) command.

-emergency

Emergency operation mode is available for operations.

In such cases, EC sessions are neither confirmed nor cancelled.

-bd *backupDeviceName*

The synchronous processing on a specific backup volume can be cancelled.

Specify the copy destination backup volume name. In the operand, specify the [AdvancedCopy Manager device name](#) for the specified device.

Note that the backup volume name and the -all option cannot be specified together.

-all

All of the synchronous processing of the transaction volume can be specified to be cancelled. Note that this specification and the -bd option cannot be used together.

OPERANDS

deviceName

Specify an AdvancedCopy Manager device name.

Refer to "[1.4 Managing Device on AdvancedCopy Manager](#)" for the devices to be specified.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Cancelling backup synchronous processing to a transaction volume:

```
# /opt/FJSVswsts/bin/swstcancelsync /dev/dsk/clt0d0s6
/dev/dsk/clt0d0s6 swstcancelsync completed.
#
```

- Cancelling backup synchronous processing of transaction volume in the emergency operation mode:

```
# /opt/FJSVswsts/bin/swstcancelsync -emergency /dev/dsk/clt0d0s6
/dev/dsk/clt0d0s6 swstcancelsync completed.
#
```

- Cancelling backup synchronous processing of the backup volume from the transaction volume:

```
# /opt/FJSVswsts/bin/swstsyncstat /dev/dsk/clt0d0s6
Server Transaction-Disk Backup-Disk Status Execute
Work2 /dev/dsk/clt0d0s6 /dev/dsk/clt0d2s5 suspend ---
Work2 /dev/dsk/clt0d0s6 /dev/dsk/clt0d2s6 executing 75%
# /opt/FJSVswsts/bin/swstcancelsync /dev/dsk/clt0d0s6 -bd /dev/dsk/clt0d2s6
/dev/dsk/clt0d0s6 swstcancelsync completed.
# /opt/FJSVswsts/bin/swstsyncstat /dev/dsk/clt0d0s6
Server Transaction-Disk Backup-Disk Status Execute
```



```
Work2 /dev/dsk/clt0d0s6 /dev/dsk/clt0d2s5 suspend ----
#
```

- Cancelling all of the backup synchronous processing of the transaction volume:

```
# /opt/FJSVswsts/bin/swstsyncstat /dev/dsk/clt0d0s6
Server Transaction-Disk Backup-Disk Status Execute
Work2 /dev/dsk/clt0d0s6 /dev/dsk/clt0d2s5 suspend ----
Work2 /dev/dsk/clt0d0s6 /dev/dsk/clt0d2s6 executing 75%
# /opt/FJSVswsts/bin/swstcancelsync /dev/dsk/clt0d0s6 -all
/dev/dsk/clt0d0s6 swstcancelsync completed.
# /opt/FJSVswsts/bin/swstsyncstat /dev/dsk/clt0d0s6
Server Transaction-Disk Backup-Disk Status Execute
Work2 /dev/dsk/clt0d0s6 ---- ---- ----
#
```

NOTES

- However, EC sessions are neither confirmed nor cancelled by execution in the emergency operation mode. If an EC session is in the error suspended state or the halt state, then unmanaged sessions remain. In this event, ETERNUS Web GUI must be used to cancel these sessions.
- You cannot cancel synchronous backup processing in the following case:
 - When the transaction volume or backup volume is a logical volume of VxVM, the structure of the logical volume is changed to a structure not supported by AdvancedCopy Manager.
 - When the transaction volume or backup volume is a logical volume of VxVM, the logical volume is deleted from VxVM.
 - The replication management function is executing a process for the target transaction volume.
- If a logical volume of VxVM is deleted or its structure is changed to a structure not supported by AdvancedCopy Manager, restoration cannot be cancelled. In such cases, take action according to either item a) or b).
 - Make sure that Advanced Copy is not in progress for the replication volume to be deleted, and execute this command with the -emergency option specified. If Advanced Copy is in progress, stop Advanced Copy processing by using ETERNUS Disk storage system).
 - Re-create the logical volume or return the logical volume structure to its original state, and then execute this command.
- For details on how to use SDX objects of PRIMECLUSTER GD, refer to "[Configuration and Conditions of Unsupported SDX Objects](#)" and "[10.1.5.1 Notes on SDX Object Operations \(Backup Management\)](#)".
- This command cannot be executed when any of the following commands are running:
 - [swstsvrset](#)
 - [swstdevinforeset](#)
 - [swstsrsematch](#)
 - [swstresback](#)
 - [swstresrst](#)

9.1.2.10 swstsyncstat (Backup Synchronous Processing Progress Display Command)

NAME

swstsyncstat - displays the execution status of the backup synchronous processing

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsts/bin/swstsyncstat [deviceName]
```

To Execute on Management Server

```
/opt/FJSSVswsts/bin/swstsyncstat [-h serverName] [deviceName]
```

DESCRIPTION

The execution status of the backup synchronous processing is displayed (during copying, under the equivalency maintained status, or in the suspended status).

If one transaction volume has multiple EC sessions, then the statuses of all of these EC sessions are displayed.

The table below lists the meaning of each title.

Title	Description
Server	Displays the name of the Managed Server.
Transaction-Disk	Displays the AdvancedCopy Manager device name of a transaction volume.
Backup-Disk	Displays the name of the backup volume. When backup synchronous processing is not being executed, the system displays "----".
Status	Displays the status of backup synchronous processing: "----" : Backup synchronous processing is not being executed. "executing" : Backup synchronous processing is being executed but equivalency maintained status has not been established. "equivalent" : Equivalency maintained status is established. "suspend" : Suspend status is established, or software-copy by PRIMECLUSTER GD is temporary stopped status. "failed" : Backup synchronous processing was abnormally terminated. "nosession" : No session exists (eg, due to inconsistent resource information). "halt" : Backup synchronous processing is halt status. "gds-error" : Copying with PRIMECLUSTER GD ended abnormally. "?????" : Failed to get copy status.
Execute	Displays the progress of the backup synchronous processing as a percentage (0 to 100%). "----" is displayed when the status is Suspend and at times when the backup synchronous processing is not running.

OPTIONS

-h serverName

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdisprv](#) command.

OPERANDS

deviceName

Specify an AdvancedCopy Manager device name.

Refer to "[1.4 Managing Device on AdvancedCopy Manager](#)" for the devices to be specified.

When you omit the device name, the system displays the status of backup synchronous processing for all of the transaction volumes.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- To display the status of backup synchronous processing to a transaction volume:

```
# /opt/FJVSswsts/bin/swstsyncstat /dev/sdb6
Server Transaction-Disk Backup-Disk Status Execute
Work1 /dev/sdb6 /dev/sdd5 suspend ----
Work1 /dev/sdb6 /dev/sdd6 executing 75%
#
```

- To display the status of backup synchronous processing for all of the transaction volumes:

```
# /opt/FJVSswsts/bin/swstsyncstat
Server Transaction-Disk Backup-Disk Status Execute
Work1 /dev/dsk/clt0d0s3 /dev/dsk/clt0d2s1 suspend ----
Work1 /dev/dsk/clt0d0s3 /dev/dsk/clt0d2s3 executing 75%
Work1 /dev/dsk/clt0d0s3 /dev/dsk/clt0d3s3 suspend ----
Work1 /dev/dsk/clt0d1s3 /dev/dsk/clt0d3s5 equivalent 100%
Work1 /dev/dsk/clt0d0s6 /dev/dsk/clt0d2s6 failed ----
Work1 /dev/dsk/clt0d1s4 ---- ---- ----
#
```

NOTES

- If "failed", "halt", or "?????" is indicated as the status of the backup synchronous processing (in the Status column), a hardware failure, etc. has possibly occurred. Refer to ["8.3 Troubleshooting: Hardware Error During Backup"](#) for more information about this.
- If "nosession" is indicated as the status of the backup synchronous processing (in the Status column), resource information having inconsistency is a possible cause. Use the [swstrsemitch](#) command to remove the inconsistency from the resource information.
- In the following cases, backup synchronous processing progress display command cannot be executed:
 - When the transaction volume or backup volume is a logical volume of VxVM, the structure of the logical volume is changed to a structure not supported by AdvancedCopy Manager.
 - When the transaction volume or backup volume is a logical volume of VxVM, the logical volume is deleted from VxVM.
- If "gds-error" is the status displayed in the Status column, a problem with PRIMECLUSTER GD is assumed. Correct the cause of the PRIMECLUSTER GD problem, and take action as follows:
 - If backup synchronous processing is in progress, execute the [swstcancelsync](#) command to cancel the backup synchronous processing.
 - If snapshot backup is in progress, execute the [swsthistdel](#) command.
 - If restoration is in progress, execute the [swstrsemitch](#) command and perform restoration again.

9.1.2.11 swstexecstat (Execution Status Display Command)

NAME

swstexecstat - displays the execution status of other commands executed

SYNOPSIS

To Execute on Managed Server

```
/opt/FJVSswsts/bin/swstexecstat [-j] [deviceName]
```

To Execute on Management Server

```
/opt/FJVSswsts/bin/swstexecstat [-j] [-h serverName] [deviceName]
```

DESCRIPTION

This command displays the execution status of other commands executed on a transaction volume. The execution status of OPC/EC is displayed if OPC/EC is being executed.

The following information is displayed.

Title	Description
Server	Displays a Managed Server name.
Device	Displays the AdvancedCopy Manager device name .
Last-Backup-Date	Displays the time when the last backup execution command was accepted.
Interval	Displays how many days have passed after the specified number of interval days since the last backup date has been exceeded. This item is displayed as "DELAY=number-of-days-passing". "OK" is displayed if the number of interval days has not been exceeded.
Status	<p>Displays whether any of the following commands is being executed on a transaction volume. If so, "EXEC" is displayed. Otherwise, "IDLE" is displayed.</p> <ul style="list-style-type: none"> - swstbkpolset - swstbkpoldel - swstbackup - swstrestore - swsthistdel - swststartsync - swstcancelsync - swstcancelrest - swstcanceltrk
Mount-Point (Method)	Displays the mount point of a device. The file system type of the mount point is displayed inside the parentheses.
Execute	<p>Displays the name of the command being executed or the progress of copying.</p> <p>"----" : Neither a command nor OPC/EC is being executed.</p> <p>"swststartsync" : The backup synchronous processing start command is being processed.</p> <p>"swstcancelsync" : The backup synchronous processing cancel command is being processed.</p> <p>"swstbackup" : The backup execution command is being processed.</p> <p>"swstrestore" : The restore execution command is being processed.</p> <p>"swstcancelrest" : The restore cancellation command is being processed.</p> <p>"swstbkpolset" : The backup policy setting command is being processed.</p> <p>"swstbkpoldel" : The backup policy deletion command is being processed.</p> <p>"swsthistdel" : The history information deletion command is being processed.</p> <p>"swstcanceltrk" : The tracking cancel command is being processed.</p> <p>"sync(xxx%)" : The progress of backup synchronous processing by EC or the software-copy function of PRIMECLUSTER GD.</p> <p>"sync(failed)" : EC ended abnormally.</p> <p>"sync(halt)" : EC is halt status.</p> <p>"snapshot(xxx%)" : The progress of OPC backup.</p> <p>"snapshot(failed)" : OPC backup ended abnormally.</p> <p>"snapshot(halt)" : OPC backup is halt status.</p> <p>"restore(xxx%)" : The progress of restoration.</p> <p>"restore(failed)" : OPC restoration ended abnormally.</p> <p>"restore(halt)" : OPC restoration is halt status.</p> <p>"gds-error": PRIMECLUSTER GD copying ended abnormally.</p> <p>"?????" : Failed to get copy status.</p>

Title	Description
	During synchronous high-speed backup with -suspend specified, no operation ("----") is indicated in the Execute column.

If multiple processes are performed on transaction volumes (for example, when the backup synchronous processing start command is issued while a real copy of the snapshot backup is being handled), the display in the Execute column conforms to the following rules:

- The display of a command name takes priority if both copying and command executions are in progress.
- When a disk copy of the snapshot backup and a disk copy of the backup synchronous processing are being handled simultaneously, the progress of the disk copy of the backup synchronous processing is displayed. A disk copy for restoration is not performed during disk copy of the snapshot backup or backup synchronous processing due to a restriction of ETERNUS Disk storage system.

The following information is displayed when -j option is specified.

Title	Description
Server	Displays a Managed Server name.
Device	Displays the AdvancedCopy Manager device name .
Execute-Date	Displays the command execution date and time (TimeStamp).
Status	<p>Displays whether any of the following commands are being executed on a transaction volume. If so, "EXEC" is displayed. Otherwise, "IDLE" is displayed.</p> <ul style="list-style-type: none"> - swstbkpolset - swstbkpodel - swstbackup - swstrestore - swsthistdel - swststartsync - swstcancelsync - swstcancelrest - swstcanceltrk
Mount-Point (Method)	Displays the mount point of a device. The file system type of the mount point is displayed inside the parentheses.
Execute	<p>Displays the name of the command being executed or the progress of copying.</p> <p>"----" : Neither a command nor OPC/EC is being executed. "swststartsync" : The backup synchronous processing start command is being processed. "swstcancelsync" : The backup synchronous processing cancel command is being processed. "swstbackup" : The backup execution command is being processed. "swstrestore" : The restore execution command is being processed. "swstcancelrest" : The restore cancellation command is being processed. "swstbkpolset" : The backup policy setting command is being processed. "swstbkpodel" : The backup policy deletion command is being processed. "swsthistdel" : The history information deletion command is being processed. "swstcanceltrk" : The tracking cancel command is being processed. "sync(xxx%)" : The progress of backup synchronous processing by EC or the software-copy function of PRIMECLUSTER GD. "sync(failed)" : EC ended abnormally. "snapshot(xxx%)" : The progress of OPC backup. "snapshot(failed)" : OPC backup ended abnormally. "restore(xxx%)" : The progress of restoration. "restore(failed)" : OPC restoration ended abnormally.</p>

Title	Description
	<p>"restore(halt)" : OPC restoration is halt status.</p> <p>"gds-error": PRIMECLUSTER GD copying ended abnormally.</p> <p>"?????" : Failed to get copy status.</p> <p>During synchronous high-speed backup with -suspend specified, no operation ("----") is indicated in the Execute column.</p>

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdisprv](#) command.

-j

Specifies to display information in another format.

OPERANDS

deviceName

Specifies the [AdvancedCopy Manager device name](#) corresponding to a transaction volume.

If this operand is omitted, the execution statuses of all the transaction volumes are displayed.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Display the execution statuses of transaction volumes:

```
# /opt/FJSVswsts/bin/swstexecstat
Server Device          Last-Backup-Date Interval Status Mount-Point (Method) Execute
Work1 /dev/dsk/clt0d0s3 2000/12/11 12:20 OK IDLE /usr1 (ufs) ----
Work1 /dev/dsk/clt0d0s4 2000/12/10 12:20 DELAY=1 IDLE /usr2 (ufs) sync(22%)
Work1 /dev/dsk/clt0d0s5 2000/12/09 12:20 DELAY=2 IDLE /usr3 (ufs) snapshot(45%)
:
#
```

- An example of using the -j option is as follows:

```
# /opt/FJSVswsts/bin/swstexecstat -j
Server Device          Execute-Date      Status Mount-Point (Method) Execute
Work1 /dev/dsk/clt0d0s3 2000/12/11 12:20 IDLE /usr1 (ufs) ----
Work1 /dev/dsk/clt0d0s4 2000/12/10 12:20 IDLE /usr2 (ufs) sync(22%)
Work1 /dev/dsk/clt0d0s5 2000/12/09 12:20 IDLE /usr3 (ufs) snapshot(45%)
:
#
```

NOTES

- If "sync (failed)", "snapshot(failed)", "restore(failed)", "restore(halt)", or "?????" is displayed in the **Execute** column, a hardware error may have occurred. Refer to ["8.3 Troubleshooting: Hardware Error During Backup"](#) for more information about this.
- In the following cases, execution status display command cannot be executed:
 - When the transaction volume or backup volume is a logical volume of VxVM, the structure of the logical volume is changed to a structure not supported by AdvancedCopy Manager.

- When the transaction volume or backup volume is a logical volume of VxVM, the logical volume is deleted from VxVM.
- If "gds-error" is the status displayed in the Status column, a problem with PRIMECLUSTER GD is assumed. Correct the cause of the PRIMECLUSTER GD problem, and take action as follows:
 - If backup synchronous processing is in progress, execute the [swstcancelsync](#) command to cancel the backup synchronous processing.
 - If snapshot backup is in progress, execute the [swsthistdel](#) command.
 - If restoration is in progress, execute the [swstsrsemitch](#) command and perform restoration again.

9.1.2.12 swstcanceltrk (Tracking Cancel Command)

NAME

swstcanceltrk - stops tracking processing

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsts/bin/swstcanceltrk [-copy | -emergency] deviceName (-bd backupDeviceName | -all)
```

To Execute on Management Server

```
/opt/FJSVswsts/bin/swstcanceltrk [-h serverName] [-copy | -emergency] deviceName (-bd backupDeviceName | -all)
```

DESCRIPTION

This command stops tracking processing.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdisprv](#) command.

-copy

Specifies that physical copying, if active, should also be stopped (the OPC session should be cancelled).

This option cannot be specified for a backup volume that has already contained a backup history. For a volume whose copy destination volume has contained a backup history, delete the history information in advance using the [swsthistdel](#) command.

-emergency

Specifies that the command should operate in the emergency operation mode.

The command does not check tracking processing and does not stop tracking processing or the OPC session.

-bd *backupDeviceName*

Stops tracking processing for the specified backup volume. Specify the target backup volume name as an operand. This option cannot be specified at the same time as the -all option.

-all

Stops the entire tracking processing specified for the transaction volumes. This option cannot be specified at the same time as the -bd option.

OPERANDS

deviceName

Specify an AdvancedCopy Manager device name.

Refer to "1.4 Managing Device on AdvancedCopy Manager" for the devices to be specified.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- The following example stops the tracking processing that was specified for the backup volume from the transaction volume:

```
# /opt/FJSVswsts/bin/swstcanceltrk /dev/dsk/clt0d0s6 -bd /dev/dsk/clt0d2s6
/dev/dsk/clt0d0s6 swstcanceltrk completed.
#
```

- The following example stops tracking processing and physical copying that were specified for the transaction volume:

```
# /opt/FJSVswsts/bin/swsttrkstat /dev/dsk/clt0d1s6
Server Transaction-Disk Backup-Disk Status Execute Update
job2 /dev/dsk/clt0d1s6 /dev/dsk/clt0d2s6 executing 33% -----
# /opt/FJSVswsts/bin/swstcanceltrk -copy /dev/dsk/clt0d1s6 -bd /dev/dsk/clt0d2s6
/dev/dsk/clt0d0s6 swstcanceltrk completed.
#
```

NOTES

- If the command is executed in emergency operation mode, no OPC session is cancelled. A session that is not managed remains in operation if the command is entered during tracking. The session must be cancelled using ETERNUS Web GUI or PRIMECLUSTER GD.
- Tracking processing cannot be stopped under the following condition. In this case, take action according to "System Administrator Corrective Action" provided for the message output.
 - When tracking processing is not performed for the specified copy source and destination volumes.
- This command cannot be executed while one of the following commands is being executed:
 - [swstsvrset](#)
 - [swstdevinfo](#)
 - [swstsrsem](#)
 - [swstresback](#)
 - [swstresrst](#)

9.1.2.13 swsttrkstat (Tracking Status Display Command)

NAME

swsttrkstat - displays the status of tracking processing

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsts/bin/swsttrkstat [deviceName]
```

To Execute on Management Server

```
/opt/FJSVswsts/bin/swsttrkstat [-h serverName] [deviceName]
```


DESCRIPTION

This command displays the status of tracking processing.

When multiple types of synchronous processing are running on one transaction volume, the command displays the statuses of all types of tracking processing.

The displayed data is explained below.

Title	Description
Server	Displays a Managed Server name.
Transaction-Disk	Displays the AdvancedCopy Manager device name of a transaction volume.
Backup-Disk	Displays a backup volume name. Displays "----" when the volume is not in the tracking status.
Status	Displays the execution status. "----" : Not in the process of tracking "executing" : In the process of physical copying and tracking "tracking" : In the process of tracking "failed" : Physical copying or tracking terminated abnormally "nosession" : No session exists. (eg, resource information is inconsistent.) "?????" : Failed to get copy status.
Execute	Displays the progress of copying as a percentage when "executing" is displayed in the Status column. "----" is displayed when "tracking" is displayed in the Status column.
Update	Displays the amount of updated data after logical copying as a percentage when "tracking" is displayed in the Status column. "----" is displayed when "executing" is displayed in the Status column.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdispsrv](#) command.

OPERANDS

deviceName

Specify an AdvancedCopy Manager device name.

Refer to "[1.4 Managing Device on AdvancedCopy Manager](#)" for the devices to be specified.

If this operand is omitted, the command displays the statuses of tracking processing for all transaction volumes.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- The following example displays the status of tracking processing of a transaction volume:

```
# /opt/FJSVswsts/bin/swsttrkstat /dev/dsk/c1t0d0s6
Server Transaction-Disk Backup-Disk Status Execute Update
job1 /dev/dsk/c1t0d0s6 /dev/dsk/c1t0d2s6 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/dsk/clt0d0s3 /dev/dsk/clt0d2s3 nosession ---- ----
job1 /dev/dsk/clt0d0s4 /dev/dsk/clt0d2s4 executing 75% ----
job1 /dev/dsk/clt0d0s5 /dev/dsk/clt0d2s5 failed ---- ----
job1 /dev/dsk/clt0d0s6 /dev/dsk/clt0d2s6 ?????? ---- ----
job1 /dev/dsk/clt0d0s7 /dev/dsk/clt0d2s7 tracking ---- 12%
job1 /dev/dsk/clt0d1s4 ---- ---- ----
:
```

NOTES

- When the execution status (ie, the value displayed in the "Status" column) is "failed" or "?????", a hardware problem, etc. has occurred. Refer to [8.3 Troubleshooting: Hardware Error During Backup](#) for more information about this.
- When the execution status (ie, the value displayed in the "Status" column) is "nosession," the resource information is probably inconsistent. Use the [swstsrsemth](#) command to remove any inconsistencies from the resource information.
- The execution status of the tracking process cannot be displayed under the following conditions:
 - When the transaction volume or backup volume is a logical volume of VxVM, the structure of the logical volume is changed to a structure not supported by AdvancedCopy Manager.
 - When the transaction volume or backup volume is a logical volume of VxVM, the logical volume is deleted from VxVM.

9.1.3 Maintenance Commands

This section describes the backup management maintenance commands.

9.1.3.1 swstsrsemth (Resource Match Command)

NAME

swstsrsemth - recovers resource information

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsts/bin/swstsrsemth [-x]
```

To Execute on Management Server

```
/opt/FJSVswsts/bin/swstsrsemth [-x] [-h serverName]
```

DESCRIPTION

This command recovers resource information if a system failure, etc. has made the resource information inconsistent.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdisprv](#) command.

-x

Achieves consistency with a repository. Specify this option if a backup management files and a repository become inconsistent.

EXIT STATUS

- =0: Completed successfully
- >0: Terminated abnormally

EXAMPLES

- Recover resource information:

```
# /opt/FJSVswsts/bin/swstsrsemtdch
swstsrsemtdch completed
#
```

NOTES

- This command can only be executed if none of the following commands are running:
 - [swstsvrset](#)
 - [swstbkpolset](#)
 - [swstbkpoldel](#)
 - [swstbackup](#)
 - [swstrestore](#)
 - [swstcancelrest](#)
 - [swststartsync](#)
 - [swstcancelsync](#)
 - [swstdevinforeset](#)
 - [swsthistdel](#)
 - [swstcanceltrk](#)
 - [swstresback](#)
 - [swstresrst](#)
 - [swstsrsemtdch](#)
- Resource match processing cannot be executed in the following circumstances:
 - When the transaction volume or backup volume is a logical volume of VxVM, the structure of the logical volume is changed to a structure not supported by AdvancedCopy Manager.
 - When the transaction volume or backup volume is a logical volume of VxVM, the logical volume is deleted from VxVM.
 - The [swstrestore](#) command using software-copy with PRIMECLUSTER GD is aborted, such as for a forced termination, and software-copy processing has not yet completed.

9.1.3.2 swstresback (Resource Backup Command)

NAME

swstresback - backup of backup management file

SYNOPSIS

```
/opt/FJSVswsts/bin/swstresback backupDirectory
```

DESCRIPTION

A backup management file is backed up to the specified directory.

- Backup management file
Management information, such as Managed Server configuration information

The information backed up by this command can be restored by the [swstresrst](#) command.

OPERANDS

backupDirectory

Specifies the name of a directory to which the backup management file is saved.

EXIT STATUS

- =0: Completed successfully
- >0: Terminated abnormally

EXAMPLES

- Backup of the backup management file:

```
# /opt/FJJSVswsts/bin/swstresback /home/backup  
swstresback completed  
#
```

NOTES

- You cannot perform backup if:
 - The specified directory does not exist.
 - Copying of data to the specified directory failed because of insufficient free space, etc.
- This command can be executed only while none of the following commands are running:
 - [swstsvrset](#)
 - [swstbkpolset](#)
 - [swstbkpoldel](#)
 - [swstbackup](#)
 - [swstrestore](#)
 - [swstcancelrest](#)
 - [swststartsync](#)
 - [swstcancelsync](#)
 - [swstdevinforeset](#)
 - [swsthistdel](#)
 - [swstcanceltrk](#)
 - [swstresback](#)
 - [swstresrst](#)
 - [swstsrsemth](#)

9.1.3.3 swstresrst (Resource Restore Command)

NAME

swstresrst - restore of backup management file

SYNOPSIS

```
/opt/FJSVswsts/bin/swstresrst backupDirectory
```

DESCRIPTION

A backup management file is restored from the specified directory.

- Managed resources information
Management information, such as Managed Server configuration information.

OPERANDS

backupDirectory

The backup location directory name directed by the [swstresback](#) command is specified.

EXIT STATUS

- =0: Completed successfully
- >0: Terminated abnormally

EXAMPLES

- Restore the backup management file:

```
# /opt/FJSVswsts/bin/swstresrst /home/backup  
swstresrst completed  
#
```

NOTES

- Before the backup management file is restored by this command, the [swstsvrset](#) command must be executed.
- You cannot perform restoration if:
 - The specified directory does not exist.
 - No backup management files exist in the specified directory.
 - Copying of data from the specified directory for restoration failed because of insufficient free space, or for some other reason.
- This command can be executed only while none of the following commands are running:
 - [swstsvrset](#)
 - [swstbkpolset](#)
 - [swstbkpoldel](#)
 - [swstbackup](#)
 - [swstrestore](#)
 - [swstcancelrest](#)
 - [swststartsync](#)
 - [swstcancelsync](#)
 - [swstdevinforeset](#)
 - [swsthistdel](#)
 - [swstcanceltrk](#)
 - [swstresback](#)
 - [swstresrst](#)

9.2 Configuration Management Commands

This section explains the operation of the configuration management command. This manages the information on a server device which performs backup and replication, or is a candidate for being used for backup or replication.



Note

Operating Environment for Command Execution

The table below lists the requirements for the operating environment in which to execute the different commands.

Table 9.4 Operating Environment for Executing Configuration Management Commands

Command	Operating Environment of AdvancedCopy Manager
stgcmmodnode	To execute this command on a Management Server (Management Server transaction), the Management Server (Management Server transaction) must be running.
stgxfwcmaddsrv	The Management Server (Management Server transaction) and target Managed Server (Managed Server transaction) must be running.
stgxfwcmmodsrv	The Management Server (Management Server transaction) and target Managed Server (Managed Server transaction) must be running.
stgxfwcmdelsrv	The Management Server (Management Server transaction) must be running.
stgxfwcmsetdev	The Management Server (Management Server transaction) and target Managed Server (Managed Server transaction) must be running.
stgxfwcmdeldev	The Management Server (Management Server transaction) and target Managed Server (Managed Server transaction) must be running.
stgxfwcmdispsrv	The Management Server (Management Server transaction) must be running.
stgxfwcmdispdev	The Management Server (Management Server transaction) must be running.
stgxfwcmdisppat	The Management Server (Management Server transaction) must be running.
stgenvdisp	This command is executable when both Management Server (Management Server transaction) and Managed Server (Managed Server transaction) are not running. Note that when the cluster service (named differently in respective clustered systems) is not running in cluster operation, this command cannot display certain information.
stgmgrinfoaset	The package-backup-of-the-management-information execution server must be running.
stgmgrinfobkup	The Management Server (Management Server transaction), package-backup-of-the-management-information execution server, and package-backup-of-the-management-information target servers must be running.
stgmgrinfoasdisp	The package-backup-of-the-management-information execution server must be running.
stgxfwcmsetsdx	For cluster operation, the Managed Server transaction must be running.
stgxfwcmsetmode	For cluster operation, the Managed Server transaction must be running.
stgrepcnv	This command is executable when both Management Server (Management Server transaction) and Managed Server (Managed Server transaction) are not running.
stgforcedelsrv	The Management Server (Management Server transaction) must be running. This command is executable when target Managed Server is not running.



Note

In Cluster Operation

In cluster operation, execute a command by taking the following steps:

1. If necessary, specify the logical node name to the transaction in the environment variable SWSTGNODE. Specify the environment variable as follows:

```
# SWSTGNODE=logicalNodeName
# export SWSTGNODE
```

2. Execute the command.

The table below lists the commands available in cluster operation, indicates whether environment variables need to be configured, and shows the nodes on which commands are executable. For notes and cautions on command execution, refer to the Remarks column.

Table 9.5 Commands When Cluster Operation Is Used

Command	Environment Variable	Command Execution Node	Remarks
stgcmmmodnode	Required	Transaction operating node	-
stgxfwcmaddsrv	Not required	Transaction operating node	Executable only on Management Server transaction
stgxfwcmmodsrv	Not required	Transaction operating node	Executable only on Management Server transaction
stgxfwcmdelsrv	Not required	Transaction operating node	Executable only on Management Server transaction
stgxfwcmsetdev	Not required	Transaction operating node	Executable only on Management Server transaction
stgxfwcmdeldev	Not required	Transaction operating node	Executable only on Management Server transaction
stgxfwcmdispsrv	Not required	Transaction operating node	Executable only on Management Server transaction
stgxfwcmdispdev	Not required	Transaction operating node	Executable only on Management Server transaction
stgxfwcmdisppat	Not required	Transaction operating node	Executable only on Management Server transaction
stgenvdisp	Not required	Transaction operating node	-
stgmgrinfobkup	Required	Transaction operating node	-
stgmgrinfodisp	Required	Transaction operating node	-
stgmgrinfoset	Required	Transaction operating node	-
stgxfwcmsetsdx	Required	Transaction operating node	-
stgxfwcmsetmode	Required	Transaction operating node	-
stgrepcnv	Not required	Transaction operating node	Executable only on Management Server transaction
stgforcedelsrv	Not required	Transaction operating node	Executable only on Management Server transaction

9.2.1 stgcmmodnode (Management Server Information Change Command)

NAME

stgcmmodnode - changes or displays Management Server information

SYNOPSIS

```
/opt/FJSVswstf/bin/stgcmmodnode [-i ipAddress] [-p portNumber] [-n serverName]
```

DESCRIPTION

This command has the following two functions:

- Displays Management Server information specified on the Managed Server on which this command is currently executed
- Changes Management Server information.

This command changes the following information on a Management Server:

- IP address of a Management Server
- PORT number used by a communication daemon that runs on a Management Server
- Server name of a Management Server

Refer to the following for the change procedure which uses this command:

- ["8.5.1 Changing IP Address of Management Server"](#)
- ["8.5.3 Changing Port Number of Management Server or Managed Server"](#)
- ["8.5.4 Changing Server Name of Management Server or Managed Server"](#)

OPTIONS

None

Displays the server information currently being managed.

-i ipAddress

Specifies the IP address of a Management Server, to change the IP address of the Management Server.

-p portNumber

Specifies the port number of a Management Server, to change the port number of the Management Server.

-n serverName

Specifies the server name of a Management Server to change the server name of the Management Server. This option cannot be specified if this command is executed on a Management Server.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Display server information:

```
# /opt/FJSVswstf/bin/stgcmmodnode
swstf2413 IP Address = 192.0.2.10
swstf2413 PORT Number = 1226
swstf2413 Storage Management Server's Name = fujitsu-storage-management-server
#
```


- Change the IP address of a Management Server to 192.0.2.11:

```
# /opt/FJSVswstf/bin/stgcmmodnode -i 192.0.2.11
swstf2401 Storage Management Server registry updated.
#
```

NOTES

- This command can be executed only by root users.
- For the port number of a Management Server, you can specify a value between 1024 and 65535.
- With this command, you must specify the same value as the one specified on a Management Server. An error occurs in the operation if an invalid value is specified.
- To execute this command, the AdvancedCopy Manager daemons must be stopped. After successful execution, you must restart the AdvancedCopy Manager daemons.
- This command is not available in an IPv6 environment. When changing IP address information in an IPv6 environment, use the [stgxfwcmmodsrv](#) command.
-

9.2.2 stgxfwcmaddsrv (Server Information Addition Command)

NAME

stgxfwcmaddsrv - adds Managed Server information

SYNOPSIS

```
/opt/FJSVswstf/bin/stgxfwcmaddsrv [-n serverName] -i ipAddress -p portNumber
```

DESCRIPTION

This command adds Managed Server information to be managed by AdvancedCopy Manager. This command can only be executed from the Management Server.

OPTIONS

-n serverName

Specifies the name of the Managed Server to be added. The name must be less than 255 bytes and must be a unique server name so that AdvancedCopy Manager can manage the server. The usable characters are 1-byte alphanumeric characters, minus sign (-), underscore (_), number sign (#), and period (.), with the first character being a 1-byte alphanumeric character. A 2-byte character causes an error.

Note that the server name is case-sensitive. When specifying the server name in other commands or Web Console, enter the server name exactly as entered in this command, including the case.

Non-cluster operation: Host name must be recognized by the operating system

Cluster operation: AdvancedCopy Manager logical node name must be specified during cluster setup

-i ipAddress

Specifies the IP address of an added Managed Server. IPv4 address or IPv6 address can be specified.

- IPv4 address
Format that separates decimal digits with periods ".". (Values according to the IPv4 rules)
- IPv6 address
Format that separates hexadecimal digits with colons ":".

Specify a physical IP address if the Managed Server is managed as a physical server. Specify a logical IP address if the IP address is made redundant and managed (as a logical server), the same as under the cluster environment.

-p *portNumber*

Specify the port number that is set on the Managed Server side to communicate with the Managed Server to be added. (Value from 1,024 to 65,535)

Under the cluster environment, specify the port number that is allocated to the service name *stgxfws_?* (Logical node name of AdvancedCopy Manager specified during cluster setup).

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Use IPv4 addresses to add servers:

```
# /opt/FJSVswstf/bin/stgxfwcmaddsrv -i 192.0.2.10 -p 1226
stgxfwcmaddsrv completed
#
```

- Use IPv6 addresses to add servers:

```
# /opt/FJSVswstf/bin/stgxfwcmaddsrv -i 2001:db8::1234:0:0:1 -p 1226
stgxfwcmaddsrv completed
#
```

- Add a logical server so that the cluster setup is performed by specifying a server name:
(Logical IP address: 192.0.2.20; port number allocated to *stgxfws_logicalNodeName*: 1227)

```
# /opt/FJSVswstf/bin/ stgxfwcmaddsrv -n Server1 -i 192.0.2.20 -p 1227
stgxfwcmaddsrv completed
#
```

NOTES

- This command can be executed only by root users.
- Change to the value of 800 characters or less before executing the command when the value of the environment variable PATH is more than 800 characters.
- Do not specify broadcast IP address to the IP address of the Managed Server to be added.
- Do not specify a server which another Management Server is managing as a Managed Server.
- Because a Management Server and a logical server for a Management Server transaction are added during database setup and cluster setup, they are not added by this command.
When processing takes a long time, a period (.) is displayed until it is complete.
- This command cannot be executed while the following commands are being executed:
 - [stgxfwcmaddsrv](#)
 - [stgxfwcmmodsrv](#)
 - [stgxfwcmdelsrv](#)
 - [stgxfwcmsetdev](#)
 - [stgxfwcmdeldev](#)
- This command cannot be executed while the following operations are being performed in the initial window by using Web Console:
 - Register Server
 - Change Server Details
 - Delete Server

- Add/Reload Device
- Delete Device

9.2.3 stgxfwcmmodsrv (Server Information Change Command)

NAME

stgxfwcmmodsrv - changes the following server information

SYNOPSIS

```
/opt/FJSVswstf/bin/stgxfwcmmodsrv -n serverName [-s newServerName] [-i newIpAddress] [-p newPortNumber] | [-f serverInformationChangeInstructionFilename [-y]]
```

DESCRIPTION

This command changes the following server information managed by AdvancedCopy Manager. It is necessary to execute this command on the Management Server.

- A Managed Server IP address
- A Managed Server port number
- The server name of a Management Server
- The server name of a Managed Server

Server information can also be changed as a batch while maintaining the consistency of multiple items of information between multiple servers in the managed system by specifying the server information change instruction file.

The following server information can be changed:

- Management Server IP address
- Managed Server IP address
- Management Server port number
- Managed Server port number
- Server name of the Management Server
- Server names of Managed Servers

OPTIONS

-n serverName

Specify a server name that is currently being managed.

Specify a server name displayed by the [stgxfwcmdispsrv](#) command exactly as displayed, including the case.

-s newServerName

Specify a new server name in less than 255 bytes. The usable characters are: 1-byte alphanumeric characters, minus sign (-), underscore (_), number sign (#), and period (.), with the first character being a 1-byte alphanumeric character. (A 2-byte character causes an error.)

Note that the new server name is case-sensitive. When specifying the server name in other commands or Web Console, enter the server name exactly as entered in this command, including the case.

Specify at the same time as the *-n* option.

-i newIpAddress

Specifies changed IP addresses.

IPv4 address or IPv6 address can be specified.

- IPv4 address

Format that separates decimal digits with periods ".". (Values according to the IPv4 rules)

- IPv6 address
Format that separates hexadecimal digits with colons ":".

Specify a physical IP address if the Managed Server is managed as a physical server. Specify a logical IP address if the IP address is made redundant and managed (as a logical server), as in a cluster environment.
Specify at the same time as the -n option.

-p *newPortNumber*

Specify a port number required for post-change communication. The values are from 1024 to 65535)
Under the cluster environment, specify the port number that is allocated to service name stgxfws_? (Logical node name of AdvancedCopy Manager specified during cluster setup)
For non-cluster operations, specify the port number assigned to service name stgxfws.
Specify at the same time as the -n option.

-f *serverInformationChangeInstructionFilename*

Specify the absolute path of the server information change instruction filename.
This option can only be specified at the same time as than the -y option.
Format for the server information change instruction file

- In the first line, enter the server name of a currently Managed Server, the information class to be changed, and the new value, delimited by spaces or tabs.
- The information classes to be changed are as follows:
S: Server name
I: IP address
P: Port number
- For IP address, specify the following values according to the server type:

For Management Server

Specify either one or both of IPv4 address and IPv6 address.

When specifying both IP addresses, separate them with commas "," (IPv4 and IPv6 in that order).

If both IP addresses are set but either of the IP addresses is to be changed, both IP addresses should be written as follows.
Any IP address that is not written is deleted:

```
Manager      I  192.0.2.10, 2001:db8::1234:0:0:1
```

For Managed Server

Specify either one of IPv4 address and IPv6 address.

- If multiple information classes, such as the IP address and the port number, need to be changed at the same server, enter the IP address change instruction, then enter the port number change instruction on the second line.

```
Manager      I  192.0.2.10
Server1      I  192.0.2.20
Server1      P  9888
Server2      S  DstServer
```

-y

Specify this option to continue processing if the management system contains multiple servers and an error occurs at a specific server.

Remove the cause of the error at the server where the error occurred, and perform the required change procedure separately for that server.

Specify at the same time as the -f option.

EXIT STATUS

- =0: Completed successfully
- >0: Terminated abnormally

EXAMPLES

- The IP address of a Managed Server (with server name: Server1) is changed to 192.0.2.40:

```
# /opt/FJSVswstf/bin/stgxfwcmmodsrv -n Server1 -i 192.0.2.40
stgxfwcmmodsrv completed
#
```

- The server name of a Managed Server (with server name: Server1) is changed to Agent5:

```
# /opt/FJSVswstf/bin/stgxfwcmmodsrv -n Server1 -s Agent5
stgxfwcmmodsrv completed
#
```

- The server information change instruction file is specified to change server information as a batch:

```
# /opt/FJSVswstf/bin/stgxfwcmmodsrv -f /acm/modfile
stgxfwcmmodsrv completed
#
```

- The server information change instruction file is specified to change server information as a batch, even if an error occurs at a specific server:

```
# /opt/FJSVswstf/bin/stgxfwcmmodsrv -f /acm/modfile -y

swstf2459 Resource adjustment command terminated abnormally. Server name=[Server1]
swsrp0625 This command cannot be executed because execution of other AdvancedCopy Manager commands
is in progress.

stgxfwcmmodsrv completed, excluding the server where the error occurs
#
```

NOTES

- Since the `stgcmmodnode` command is used to change the following server information, this information can only be changed if the server information change instruction file is specified:
 - The IP address of a Management Server
 - The port number of a Management Server
- When changing server information by this command, the definition information contained in the machine and AdvancedCopy Manager should be changed beforehand. Refer to "[8.5 Changing Operating Environment](#)" for information on executing the processing.
- This command can be executed only by root users.
- Change to the value of 800 characters or less before executing the command when the value of the environment variable PATH is more than 800 characters.
- When processing takes a long time, a period (.) is displayed until it is complete.
- The following notes apply to changing the Management Server information (except when the `-f` option is specified):
 - Use the `stgcmmodnode` command to change the IP address and port number, then restart the various AdvancedCopy Manager services.
 - After Management Server information is changed, execute the `stgcmmodnode` command at the Managed Servers to set the changed Management Server information.
- The following notes apply to changing the Managed Server information:
 - Stop the various AdvancedCopy Manager services before changing the IP address. Then, change the IP address in the software that provides the system or logical IP address and restart the AdvancedCopy Manager daemon.
 - Stop the various AdvancedCopy Manager services before changing the port number. Then, change the port number in the system and restart the AdvancedCopy Manager daemon.

- The following notes apply when the -f option is specified:
 - The server name or IP address that is to be changed cannot be in use by another server that is currently managed.
 - The server name to be changed is case-sensitive. Specify a new server name displayed by the `stgxfwcmdispsrv` command exactly as displayed, including the case.
 - The new server name is case-sensitive. When specifying the server name in other commands or Web Console, enter the server name exactly as entered in this command, including the case.
 - The Management Server that executes this command and the AdvancedCopy Manager services and Managed Server transactions at all Managed Servers managed by Management Server transactions must be running.
 - Backup management commands and replication management commands must not be executed at the Management Server that executes this command nor in the AdvancedCopy Manager services and Managed Server transactions at all Managed Servers managed by Management Server transactions.
If commands are executed, they may terminate abnormally. If the server information change command terminates abnormally, perform recovery as indicated in the error message. If backup management or replication management commands are accidentally executed during execution of the server information change command, and if the backup management and replication management commands terminate abnormally, re-execute the backup management or replication management commands after normal termination of the server information change command.
- The following notes apply when the -y option is specified:
 - If this option is specified, processing is forcibly continued if a communication or other error occurs for some servers. Therefore, consistency of the overall system cannot be maintained. Usually, do not specify this option.
Specify this option only when absolutely necessary due to problems at some servers.
 - Take the action indicated in the error message at the server which output the message. Normal operation is not possible until this action is taken, so take action promptly.
- This command cannot be executed while the following commands are being executed:
 - `stgxfwcmaddsrv`
 - `stgxfwcmmodsrv`
 - `stgxfwcmdelsrv`
 - `stgxfwcmsetdev`
 - `stgxfwcmdeldev`
- This command cannot be executed while the following operations are being performed in the initial window by using Web Console:
 - Register Server
 - Change Server Details
 - Delete Server
 - Add/Reload Device
 - Delete Device

9.2.4 stgxfwcmdelsrv (Server Information Deletion Command)

NAME

`stgxfwcmdelsrv` - deletes the Managed Server information

SYNOPSIS

```
/opt/FJSVswstf/bin/stgxfwcmdelsrv -n serverName
```

DESCRIPTION

This command deletes the Managed Server information that AdvancedCopy Manager is currently managing. It is necessary to execute this command on the Management Server.

OPTIONS

`-n serverName`

Specify the Managed Server name to be deleted.

Specify a server name displayed by the [stgxfwcmdispsrv](#) command exactly as displayed, including the case.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- A Managed Server (server name: Server1) information is deleted:

```
# /opt/FJSVswstf/bin/stgxfwcmdelsrv -n Server1
stgxfwcmdelsrv completed
#
```

NOTES

- This command can be executed only by root users.
- The Management Server information cannot be deleted.
- When processing takes a long time, a period (.) is displayed until it is complete.
- This command cannot be executed while the following commands are being executed:
 - [stgxfwcmaddsrv](#)
 - [stgxfwcmmodsrv](#)
 - [stgxfwcmdelsrv](#)
 - [stgxfwcmsetdev](#)
 - [stgxfwcmdeldev](#)
- This command cannot be executed when the following operations are being performed in the initial window by using Web Console:
 - Register Server
 - Change Server Details
 - Delete Server
 - Add/Reload Device
 - Delete Device

9.2.5 stgxfwcmsetdev (Device Information Collection/Reflection Command)

NAME

`stgxfwcmsetdev` - collects and reflects the device information

SYNOPSIS

```
/opt/FJSVswstf/bin/stgxfwcmsetdev -n serverName [-d deviceName] [-f fileName] [-y]
```

DESCRIPTION

This command collects the latest device information detected by the server and displays the difference between this and the information managed by AdvancedCopy Manager. After this, it displays the specified device difference information in the management information for AdvancedCopy Manager. It is necessary to execute this command on the Management Server.

The device information currently being managed and the device information detected are compared, stored, deleted, updated, and used in the following cases:

- When operation is started using AdvancedCopy Manager
- When device information being managed is changed

OPTIONS

-n serverName

Specify the server name whose device information is to be collected or reflected.

Specify a server name displayed by the [stgxfwcmdispsrv](#) command exactly as displayed, including the case.

-d deviceName

Specify the device name whose information is to be collected or reflected.

If this specification is omitted, all devices on the specified server are assumed to be specified.

-f fileName

Use the absolute path to specify the file name to which the difference between the device information managed currently and the device information collected is output. Only the device difference information is output to the file, and the reflection processing for the subsequent management information is not executed.

This option cannot be specified with the *-y* option.

-y

The difference between the device information managed by AdvancedCopy Manager and the latest device information detected is reflected unconditionally in the management information, without an inquiry being made about that difference.

This option cannot be specified with the *-f* option.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- The information for each device of a server is collected or reflected:

```
# /opt/FJSVswstf/bin/stgxfwcmsetdev -n Server1
>Device detected
  1: /dev/rdisk/c1t0d0s6 (/aaa)
  2: /dev/rdisk/c1t0d0s7
>Specify devices to be added. [* or ?-? or ?,? or q]

>Device not detected
  1: /dev/rdisk/c2t1d0s1
  2: /dev/rdisk/c2t1d0s2 (/bbb)
  3: /dev/rdisk/c3t0d0s7

>Specify a device requested to be deleted. [* or ?-? or ?,? or q]
*1,3
>Information for the following device is changed. Information will be updated.
  1: /dev/rdisk/c1t0d0s5
>Do you want to continue the processing? [y or n] y
stgxfwcmsetdev completed
#
```


- The information for a specific device managed by the server is collected or reflected:

```
# /opt/FJSVswstf/bin/stgxfwcmsetdev -n Server1 -d /dev/rdisk/c4t0d10s0
> The following device information has been changed. This information is updated.
  1:/dev/rdisk/c4t0d10s0 (/zzz)
> Do you want to continue the processing? [y or n] y
stgxfwcmsetdev completed
#
```

NOTES

- The following message is displayed if the number of devices is high:

```
- hit Enter for more or s to select -
```

Press the 'Enter' key to continue to display more devices. To interrupt the display and specify either an addition or a deletion, press 's' key.

- Use the following methods to request addition or deletion:

- Enter "*" to specify all.
- Enter numbers separated by a hyphen to specify a range. (Example: 1-3)
- Enter numbers separated by a comma to specify separate numbers. (Example: 2, 4)
- A combination of range and separate number specifications can be entered. (Example: 1-3, 5)
- Enter no data to specify no devices to add or delete (only press Enter in this case).
- Enter "q" to stop the processing.

- This command can be executed only by root users.

- Obtaining disk information for HP-UX servers with GPT disks using AdvancedCopy Manager is not possible.

- When processing takes a long time, a period (.) is displayed until it is complete.

- The following message may be displayed when a device requested to be added is specified. It is an indication of an attempt to register a device that is in the box and may have invalid settings. Make sure that this problem is fixed before continuing the processing (This message is also displayed when the -y option is specified.)

```
>The setting of license or the table size is not done, or the table size is
insufficient, to the cabinet where the partition that makes up the following
device exists.
>Continue? [y or n]
```

- The following message may be displayed when a device has been requested to be added.

```
> The partitions that comprise the following device are already managed.
> Do you want to continue the processing? [y or n]
```

This message may be displayed in the following case.

Check the following before continuing the processing: whether the device is subject to duplicated registration. This message is also displayed when the -y option is specified.

- In the case two or more logical servers running on the same physical server are managed on the clustered system operating environment.
 - In the case a device with the other name that indicates a physical disk in ETERNUS Disk storage system is reflected into the management information of AdvancedCopy Manager.
- The amount of time required to process this command is proportional to the total number of devices defined on the specified server when the device information is collected. Therefore, when the number of devices is large, execute the command at a time of low

CPU and I/O loads.

Note that it takes about 0.5 second for each device (partition) when there is no load on the system.

- This command cannot be executed while the following commands are being executed:
 - [stgxfwcmaddsrv](#)
 - [stgxfwcmmodsrv](#)
 - [stgxfwcmdelsrv](#)
 - [stgxfwcmsetdev](#)
 - [stgxfwcmdeldev](#)
- This command cannot be executed by using Web Console when the following operations are being performed:
 - Register Server
 - Change Server Details
 - Delete Server
 - Add/Reload Device
 - Delete Device

9.2.6 stgxfwcmdeldev (Device Information Deletion Command)

NAME

stgxfwcmdeldev - deletes the device information

SYNOPSIS

```
/opt/FJSVswstf/bin/stgxfwcmdeldev -n serverName -d deviceName
```

DESCRIPTION

This command deletes the device information currently being managed by AdvancedCopy Manager. It is necessary to execute this command in the Management Server.

The [stgxfwcmsetdev](#) command compares the device information currently being managed and the device information detected before it executes the processing. However, this command deletes the device information irrespective of whether the device information currently being managed has been detected.

OPTIONS

-n *serverName*

Specify the name of the server that is managing the device to be deleted.

Specify a server name displayed by the [stgxfwcmdispsrv](#) command exactly as displayed, including the case.

-d *deviceName*

Specify the name of the device to be deleted.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- The device information for a Managed Server is deleted:

```
# /opt/FJJSVswstf/bin/stgxfwcmdeldev -n Server1 -d /dev/rdisk/c4t0d12s7
stgxfwcmdeldev completed
#
```

NOTES

- This command can be executed only by root users.
- When processing takes a long time, a period (.) is displayed until it is complete.
- This command cannot be executed while the following commands are being executed:
 - [stgxfwcmaddsrv](#)
 - [stgxfwcmmodsrv](#)
 - [stgxfwcmdelsrv](#)
 - [stgxfwcmsetdev](#)
 - [stgxfwcmdeldev](#)
- This command cannot be executed by using Web Console if the following operations are being performed:
 - Register Server
 - Change Server Details
 - Delete Server
 - Add/Reload Device
 - Delete Device

9.2.7 stgxfwcmdispsrv (Server Information Display Command)

NAME

stgxfwcmdispsrv - displays the server information

SYNOPSIS

```
/opt/FJJSVswstf/bin/stgxfwcmdispsrv [-n serverName]
```

DESCRIPTION

This command displays the server information managed by AdvancedCopy Manager. It is necessary to execute this command on the Management Server.

The contents to be displayed are as follows:

Title	Description
Server	A server name is displayed.
Server-Type	A server type is displayed. "Management-Server" : Management Server "Server" : Managed Server
Tape	"---" is always displayed.
OS-Type	An operating system type is displayed.
OS-Version	An operating system version is displayed.

Title	Description
Version	The version of this product is displayed.
IP-Address	<p>The IP address for communication with the server is displayed.</p> <ul style="list-style-type: none"> - For Management Server <ul style="list-style-type: none"> - If the number of registered IP addresses is one: Registered information is displayed. - If the number of registered IP addresses is two: Registered IP addresses are displayed, IPv4 address and IPv6 address in that order separated with commas. IPv6 address is expressed in abbreviated notation. - For Managed Server Registered IP address is displayed.
Port-Number	The port number for communication with the server is displayed.

OPTIONS

-n *serverName*

Specify the server name whose information is to be displayed.

Specify a server the same name entered at the [stgxfwcmaddsrv](#) or [stgxfwcmmodsrv](#) command exactly, including the case.

If this option is omitted, the information for all servers managed by AdvancedCopy Manager is displayed.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- The information for all servers is displayed:

```
# /opt/FJSVswstf/bin/stgxfwcmdispsrv
Server  Server-Type      Tape OS-Type OS-Version Version IP-Address      Port-Number
Manager Management-Server --- Solaris 11      V16.0  192.0.2.10      1226
Server1 Server              --- Solaris 10      V16.0  192.0.2.20      1226
Server2 Server              --- Solaris 10      V16.0  2001:db8::1234:0:0:1 1226
#
```

- The information for a Managed Server(server name: Server1) is displayed:

```
# /opt/FJSVswstf/bin/stgxfwcmdispsrv -n Server1
Server  Server-Type      Tape OS-Type OS-Version Version IP-Address      Port-Number
Server1 Server              --- Solaris 10      V16.0  192.0.2.20      1226
#
```

- The information for a Management Server(server name: Manager) is displayed:

```
# /opt/FJSVswstf/bin/stgxfwcmdispsrv -n Manager
Server  Server-Type      Tape OS-Type OS-Version Version IP-Address      Port-Number
Manager Management-Server --- Solaris 11      V16.0  192.0.2.10,2001:db8::1234:0:0:1 1226
#
```

NOTES

- In "OS-Version", the OS version that is displayed refers to the following time:
In the case of the Management Server: the time when the database was created.

In the case of the Managed Server: the time when the Managed Server was added.

Even if the OS is upgraded after this, the displayed OS version retains the version at the database creation time and the time the Managed Server was added.

- This command can be executed only by root users.
- When processing takes a long time, a period (.) is displayed until it is complete.

9.2.8 stgxfwcmdispdev (Device Information Display Command)

NAME

stgxfwcmdispdev - displays the Device information

SYNOPSIS

```
/opt/FJSVswstf/bin/stgxfwcmdispdev -n serverName [-d deviceName [-sm|-grp]]
```

DESCRIPTION

This command displays information on a device being managed by AdvancedCopy Manager. The list of devices which use the same copy area as the specified device, the specified device itself, and the devices which exist in the same logical group can be displayed. It is necessary to execute this command on the Management Server.

The following contents are displayed. Note that "SDX" or "SDX object" in the table are SynfinityDISK or PRIMECLUSTER GD terms.

Title	Description
S	<p>The status of each device is indicated. The following is displayed according to the device status:</p> <ul style="list-style-type: none"> * : Device not set up (The AdvancedCopy license setting may be undefined, copy table size may be unspecified, or the copy table may be too small in the box containing this device.) D : Device deletion error U : Device update error
Device	<p>Solaris server:</p> <ul style="list-style-type: none"> A RAW device name is displayed. <p>Windows server:</p> <ul style="list-style-type: none"> A device name managed internally by AdvancedCopy Manager is displayed. The name is displayed in the g?d? and g?d?p? format. g? : Disk array number (0 if the device is not installed in the ETERNUS Disk storage system) d? : Disk number(*1) p? : Logical disk number(*2) <p>(*1)This is a number given uniquely within the ETERNUS Disk storage system. This number differs from the disk number in Disk Management.</p> <p>(*2)This is the sequential number of a logical disk on the physical disk. The logical disk indicates the following contents in Disk Management. The number is given by counting up from the left to the right in Disk Management.</p> <ul style="list-style-type: none"> - In the case of a basic partition, the logical disk is that entire partition. - In the case of an extended partition, the logical disk is a logical drive in that partition. <p>In the case of an SDX object, the name is displayed in the disk class name or mirror volume name g?d?p? format.</p> <ul style="list-style-type: none"> Disk class name : Disk class name of SDX object Mirror volume name : Mirror volume name of SDX object <p>HP-UX server:</p> <ul style="list-style-type: none"> A RAW device name or a volume group name under LVM is displayed.

Title	Description
	<p>Linux server: When a RAW device has been created, its name is displayed; otherwise, a block device name is displayed.</p> <p>AIX server: A RAW device name or a volume group name under LVM is displayed.</p>
Block-Device	<p>Solaris server: A block device name is displayed.</p> <p>Windows server: The name is displayed in the disk?p? format. disk? : Disk number displayed by disk administrator p? : Logical disk number (*2)</p> <p>HP-UX server: A block device name or a volume group name under LVM is displayed.</p> <p>Linux server: A block device name is displayed.</p> <p>AIX server: A block device name or a volume group name under LVM is displayed.</p>
Mount-Point	<p>Solaris, HP-UX, AIX and Linux server: A mount point name is displayed.</p> <p>Windows server: A drive character is displayed.</p>
File-System	A file system name is displayed.
Device-Type	A device type is displayed. Normal/SDX(slice)/LVM/VxVM/SDX(volume)
Size(MB)	The capacity (in megabytes) allocated to a device is displayed.
RDB-Name	<p>Windows and Linux server: The Symfoware RDB system name is displayed.</p> <p>Solaris, HP-UX and AIX server: "---" is always displayed.</p>
Log-Group-Name	<p>Windows and Linux server: The Symfoware log group name is displayed.</p> <p>Solaris, HP-UX and AIX server: "---" is always displayed.</p>
DB-Space-Name	<p>Windows and Linux server: The Symfoware database space name is displayed.</p> <p>Solaris, HP-UX and AIX server: "---" is always displayed.</p>

OPTIONS

-n *serverName*

Specify the server name that manages the device whose information is to be displayed.

-d *deviceName*

Specify the device name whose information is to be displayed.

If this option is omitted, the information for all devices managed by the specified server is displayed.

-sm

With the -d option also specified, this option specifies that a list of the devices using the same copy area as the specified device be displayed.

-grp

The list of the devices which exist in the same logical group as a specified device is displayed by specifying it simultaneously with the -d option. The device type of the specified device can be specified with the d option only for LVM/VxVM.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- The information for a device of a server is displayed:

```
# /opt/FJJSVswstf/bin/stgxfwcmdispdev -n Server1 -d /dev/rdsk/c4t0d10s0
S Device          Block-Device          Mount-Point File-System Device-Type Size(MB)
  /dev/rdsk/c4t0d10s0 /dev/dsk/c4t0d10s0 /zzz          ufs          Normal      1028
#
```

- The information for each device of a server is displayed:

```
# /opt/FJJSVswstf/bin/stgxfwcmdispdev -n Server1
S Device          Block-Device          Mount-Point File-System Device-Type Size(MB)
  /dev/sfdsk/class/rdsk/vol1 /dev/sfdsk/class/dsk/vol1 ---          ---          SDX(Volume) 25000
* /dev/sfdsk/class/rdsk/vol2 /dev/sfdsk/class/dsk/vol2 /work        ufs          SDX(Volume) 1250
U /dev/rdsk/c1t0d0s0 /dev/dsk/c1t0d0s0    /storagedata ufs          Normal      5200
D /dev/rdsk/c1t0d0s1 /dev/dsk/c1t0d0s1    ---          ---          Normal      15400
#
```

- Device of the server is displayed and device information which uses the same copy area is displayed:

```
# /opt/FJJSVswstf/bin/stgxfwcmdispdev -n Server1 -d /dev/vx/rdsk/dg1/vol1 -sm
Device          Mount-Point
/dev/vx/rdsk/dg1/vol1 /storagedata
/dev/vx/rdsk/dg1/vol2 ---
/dev/vx/rdsk/dg1/vol3 /zzz
#
```

- Device information which exists in the same logical group as device of the server is displayed:

```
# /opt/FJJSVswstf/bin/stgxfwcmdispdev -n Server1 -d /dev/vx/rdsk/dg1/vol1 -grp
Device          Mount-Point
/dev/vx/rdmp/c1t0d0s2 ---
/dev/vx/rdmp/c1t0d1s2 ---
/dev/vx/rdsk/dg1/vol1 /storagedata
/dev/vx/rdsk/dg1/vol2 ---
/dev/vx/rdsk/dg1/vol3 /zzz
/dev/vx/rdsk/dg1/vol4 ---
#
```

NOTES

- The command displays the above details based on information registered in AdvancedCopy Manager. Therefore, when the configuration has been changed and there is a device not registered in a logical group, and it is registered afterwards, the correct information is not displayed. In this case, re-attempt the acquisition of device information about the registration of the device that was not previously registered, and the device which does the "Acquisition/reflection of information on all devices" processing.
- "RDB-Name", "Log-Group-Name" and "DB-Space-Name" are displayed under the following conditions:
 - The Managed Server's OS is Linux or Windows.
 - The Symfaware Server Advanced Backup Controller V10.0 or later is installed in the Managed Server.
 - The target device to be displayed exists in the ETERNUS Disk storage system.
 - When collecting the device information, the target RDB operating environment is running.
- This command can be executed only by root users.
- When processing takes a long time, a period (.) is displayed until it is complete.

9.2.9 stgxfwcmdisppat (Partition Information Display Command)

NAME

stgxfwcmdisppat - displays the information for the partitions

SYNOPSIS

```
/opt/FJSVswstf/bin/stgxfwcmdisppat -n serverName -d deviceName
```

DESCRIPTION

This command displays the information for the partitions that comprise a device managed by AdvancedCopy Manager. It is necessary to execute this command in the Management Server.

The following contents are displayed. Note that "SDX" or "SDX object" in the table are SynfinityDISK or PRIMECLUSTER GD terms.

Title	Description
Partition	<p>Solaris server:</p> <p>A RAW device name is displayed. In the case of an SDX object, the name is displayed in the RAW device name c?t?d? format.</p> <p>Windows server:</p> <p>When the device is installed in the ETERNUS Disk storage system, the same contents as the stgxfwcmdispdev command are displayed. When the device is not installed in the ETERNUS Disk storage system, the name is displayed in the following formats.</p> <p>MBR Disks:</p> <p>S?p?</p> <p>S? : Signature (8-digit hexadecimal number) p? : Logical disk number (*1)</p> <p>GPT Disks:</p> <p>G? : DiskId of disk (32-bit hexadecimal character string) p? : Logical disk number (*1)</p> <p>(*1) This is the sequence number of a logical disk on the physical disk. The logical disk indicates the following contents in Disk Management. The number is given by counting up from the left to the right in Disk Management.</p> <ul style="list-style-type: none">- In the case of a basic partition, the logical disk is the entire partition.- In the case of an extended partition, the logical disk is a logical drive in that partition

Title	Description
	<p>HP-UX server: A RAW device name or a volume group name under LVM is displayed.</p> <p>Linux server: When a RAW device has been created, its name is displayed. Otherwise, a block device name is displayed. In the case of an SDX object (i.e., a slice), it displays in the form of RAW device name:sd?. Refer to "1.4.2 SDX Objects of PRIMECLUSTER GD/SynfinityDISK" for details.</p> <ul style="list-style-type: none"> - RAW device name: The udev device name indicated by the sdxinfo DEVNAME value - RAW device name: The sdxinfo DEVNAME value <p>AIX server: A RAW device name or a volume group name under LVM is displayed.</p>
Identifier	<p>Solaris server: A block device name is displayed. In the case of an SDX object, the name is displayed in the block device name c?t?d? format.</p> <p>Windows server: A device name managed internally by AdvancedCopy Manager is displayed. The same contents as Device of the stgxfwcmdispdev command are displayed.</p> <p>HP-UX server: A block device name or a volume group name under LVM is displayed.</p> <p>Linux server: A block device name is displayed. In the case of an SDX object (i.e., a slice), it displays in the form of a block device name:sd?. Refer to "1.4.2 SDX Objects of PRIMECLUSTER GD/SynfinityDISK" for details.</p> <ul style="list-style-type: none"> - RAW device name: The udev device name indicated by the sdxinfo DEVNAME value - RAW device name: The sdxinfo DEVNAME value <p>AIX server: A block device name or a volume group name under LVM is displayed.</p>
Physical-Disk	<p>Solaris, HP-UX, Linux, or AIX server: "---" is always displayed.</p> <p>Windows server: A physical device name is displayed. The same contents as 'Block-Device' of the stgxfwcmdispdev command are displayed.</p>
Disk-Type	<p>When the device is installed in the ETERNUS Disk storage system, "GR" is displayed. When the device is not installed in the ETERNUS Disk storage system, "---" is displayed.</p> <p>Solaris or Linux server: When the device is an SDX object in units of logical volumes, "---" is displayed.</p>
LU-Number	<p>When the device is installed in the ETERNUS Disk storage system, a logical unit number is displayed. When the device is not installed in the ETERNUS Disk storage system, "---" is displayed.</p> <p>Solaris or Linux server: When the device is an SDX object in units of logical volumes, "---" is displayed.</p>

OPTIONS

`-n serverName`

Specify the server name that manages a device whose information is to be displayed.

Specify a server name displayed by the `stgxfwcmdispsrv` command exactly as displayed, including the case.

`-d deviceName`

Specify a device name whose information is to be displayed.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- The partition information for a device of a server is displayed:

```
# /opt/FJSVswstf/bin/stgxfwcmdisppat -n Server1 -d /dev/rdisk/c4t0d10s0
Partition          Identifier          Physical-Disk  Disk-Type  LU-Number
/dev/rdisk/c4t0d10s0 /dev/dsk/c4t0d10s0 ---          GR          10
#
```

NOTES

- This command can be executed only by root users.
- When processing takes a long time, a period (.) is displayed until it is complete.

9.2.10 stgenvdisp (Environment Information Display Command)

NAME

`stgenvdisp` - displays the values of the environment settings

SYNOPSIS

```
/opt/FJSVswstf/bin/stgenvdisp [-n serverName]
```

DESCRIPTION

This command displays the environmental information set in AdvancedCopy Manager in the server where the command is executed. The contents of the displayed information are described below.

Title	Description
Server type	The server type is displayed as follows: Management server : Management Server or Management Server transaction Server : Managed Server or Managed Server transaction
Cluster system	The type of clustered system is displayed as follows: MSCS : Clustered system on Windows environment GENERAL : Clustered system on Solaris, Linux, or HP-UX environment
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.

Title	Description
Service name	The transaction name is displayed.
Cluster name	The cluster name is displayed.
Mount point	The mount point for AdvancedCopy Manager shared data is displayed.
Disk resource	The name of the shared disk resource used for AdvancedCopy Manager shared data is displayed.
Drive	The drive letter used for AdvancedCopy Manager shared data is displayed.
Network name	The network name used for AdvancedCopy Manager is displayed.
External connection	The external connection network name is displayed.
Subnet mask	The subnet mask is displayed.
IP address resource	IP address resource information is displayed.
IP by cluster setup	The IPv4 format logical IP address specified in the cluster setting command is displayed. However, note that the displayed IP address is not the active logical IP address if the IP address in either "8.5.6.1 Changing IP Address of Management Server Transaction" or "8.5.6.2 Changing IP Address of Managed Server Transaction" is changed during operation.
IPv6 by cluster setup	The IPv6 format logical IP address specified in the cluster setting command is displayed. However, note that the displayed IP address is not the active logical IP address if the IP address in either "8.5.6.1 Changing IP Address of Management Server Transaction" or "8.5.6.2 Changing IP Address of Managed Server Transaction" is changed during operation.
Manager IP address	<p>The IP address of the Management Server is displayed. Displayed information varies with the Server type.</p> <ul style="list-style-type: none"> - For "Management server" <p>The IP address of the Management Server is displayed. If both IPv4 address and IPv6 address are set, they are displayed in two lines. If the Management Server is a cluster configuration, the physical IP address is displayed.</p> - For "Managed server" <p>The IP address of the Management Server to which this Managed Server is registered is displayed. IPv4 address or IPv6 address is displayed. However, if the Managed Server is not registered on the Management Server (operation of adding servers is not done), "???" is displayed in the value field.</p> <p>IPv6 address is expressed in abbreviated notation.</p>
Communication	<p>The communication service port number (the service name in the services file: stgxfws) is displayed. In a cluster environment, the transaction communication service port number (the service name in the services file: stgxfws_logicalNodeName) is displayed.</p>
Program install directory	Displays the "Program directory" specified during installation.
Fixed configuration directory	Displays the "Environment directory" specified during installation.
Modifying configuration directory	Displays the "Work directory" specified during installation.
Version	Displays the version and level of AdvancedCopy Manager.
Language	<p>Displays the AdvancedCopy Manager language type.</p> <p>JAPANESE : Japanese version ENGLISH : English version</p>

OPTIONS

-n *serverName*

Specifies the logical node name (logical node name specified in the cluster setup command) of the transaction.

In Cluster Environment

If this option is specified, information on only the specified transactions is output. Otherwise, information on all transactions is output.

In Non-cluster Environment

This option cannot be specified. If the option is specified, then message swstf6002 is output and processing ends. Otherwise, information is output.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Information in a non-clustered system is displayed:

[For Management Server using IPv4 address]

```
# /opt/FJSVswstf/bin/stgenvdisp
Server type           : Management server
Manager IP address    : 192.0.2.10
Communication         : 1226
-----
Version              : V16.6
Language             : ENGLISH
Program install directory : /opt
Fixed configuration directory : /etc/opt
Modifying configuration directory : /var/opt
#
```

[For Management Server using IPv4 address and IPv6 address]

```
# /opt/FJSVswstf/bin/stgenvdisp
Server type           : Management server
Manager IP address    : 192.0.2.10
Manager IP address    : 2001:db8::1234:0:0:1
Communication         : 1226
-----
Version              : V16.6
Language             : ENGLISH
Program install directory : /opt
Fixed configuration directory : /etc/opt
Modifying configuration directory : /var/opt
#
```

[For Managed Server using IPv4 address]

```
# /opt/FJSVswstf/bin/stgenvdisp
Server type           : Server
Manager IP address    : 192.0.2.10
Communication         : 1226
-----
Version              : V16.6
Language             : ENGLISH
Program install directory : /opt
Fixed configuration directory : /etc/opt
Modifying configuration directory : /var/opt
#
```

[For Managed Server using IPv6 address]

```
# /opt/FJSVswstf/bin/stgenvdisp
Server type      : Server
Manager IP address : 2001:db8::1234:0:0:1
Communication    : 1226
```

```
-----
Version          : V16.6
Language         : ENGLISH
Program install directory : /opt
Fixed configuration directory : /etc/opt
Modifying configuration directory : /var/opt
```

#

- Information on all transactions in a cluster environment is displayed.

[For Managed Server using IPv4 address]

```
# /opt/FJSVswstf/bin/stgenvdisp
Server type      : Management server
Cluster system   : GENERAL
Node type        : Primary
Node name        : nodemgr
Service name     : manager_service
Mount point      : /acmmgrdisk
IP by cluster setup : 192.0.2.10
Manager IP address : 192.0.2.10
Communication    : 11000
```

```
-----
Server type      : Server
Node name        : Local-Service
Manager IP address : 192.0.2.10
Communication    : 1226
```

```
-----
Version          : V16.6
Language         : ENGLISH
Program install directory : /opt
Fixed configuration directory : /etc/opt
Modifying configuration directory : /var/opt
```

#

[For Management Server using IPv4 address and IPv6 address]

```
# /opt/FJSVswstf/bin/stgenvdisp
Server type      : Management server
Cluster system   : GENERAL
Node type        : Primary
Node name        : nodemgr
Service name     : manager_service
Mount point      : /acmmgrdisk
IP by cluster setup : 192.0.2.10
IPv6 by cluster setup : 2001:db8::1234:0:0:1
Manager IP address : 192.0.2.10
Manager IP address : 2001:db8::1234:0:0:1
Communication    : 11000
```

```
-----
Server type      : Server
Node name        : Local-Service
Manager IP address : 192.0.2.10
Communication    : 1226
```

```
-----
Version          : V16.6
Language         : ENGLISH
Program install directory : /opt
```

```
Fixed configuration directory      : /etc/opt
Modifying configuration directory  : /var/opt
```

```
#
```

[For Managed Server using IPv4 address]

```
# /opt/FJSVswstf/bin/stgenvdisp
Server type      : Server
Cluster system   : GENERAL
Node type        : Primary
Node name        : nodeagt1
Service name     : agent_servicel
Mount point      : /acmagtldisk
Manager IP address : 192.0.2.10
Communication    : 11000
-----
Server type      : Server
Cluster system   : GENERAL
Node type        : Secondary
Node name        : nodeagt2
Service name     : agent_service2
Mount point      : /acmagt2disk
Manager IP address : 192.0.2.10
Communication    : 12000
-----
Server type      : Server
Node name        : Local-Service
Manager IP address : 192.0.2.10
Communication    : 1226
-----
Version          : V16.6
Language         : ENGLISH
Program install directory : /opt
Fixed configuration directory : /etc/opt
Modifying configuration directory : /var/opt
```

```
#
```

[For Managed Server using IPv6 address]

```
# /opt/FJSVswstf/bin/stgenvdisp
Server type      : Server
Cluster system   : GENERAL
Node type        : Primary
Node name        : nodeagt1
Service name     : agent_servicel
Mount point      : /acmagtldisk
Manager IP address : 2001:db8::1234:0:0:1
Communication    : 11000
-----
Server type      : Server
Cluster system   : GENERAL
Node type        : Secondary
Node name        : nodeagt2
Service name     : agent_service2
Mount point      : /acmagt2disk
Manager IP address : 2001:db8::1234:0:0:1
Communication    : 12000
-----
Server type      : Server
Node name        : Local-Service
Manager IP address : 2001:db8::1234:0:0:1
Communication    : 1226
-----
Version          : V16.6
```

```
Language           : ENGLISH
Program install directory : /opt
Fixed configuration directory : /etc/opt
Modifying configuration directory : /var/opt
#
```

NOTES

- To output cluster operation information, execute the command from an active node. Executing the command from a standby node displays only partial information. To output cluster operation information while a transaction is stopped, either set the shared disk containing the AdvancedCopy Manager shared data of the transaction to the online state, or mount the shared disk and then execute the command.
- When this command is executed in a clustered system, the environment variable SWSTGNODE need not be set.
- This command can be executed only by root users.
- "???" is displayed for any information that cannot be displayed because of an error.
- Some items may not be output, depending on the server type (whether Management Server or Managed Server) or cluster type.

9.2.11 stgmgrinfoset (Management Information Package Backup Server Registration/Deletion Command)

NAME

stgmgrinfoset - registers, deletes or displays the management information

SYNOPSIS

```
/opt/FJISVswstf/bin/stgmgrinfoset [-del | -d directoryName -c managementServerEncoding]
```

DESCRIPTION

This command can be used to register, delete or display information about the management information package backup execution server(the server on which the [stgmgrinfobkup](#) command is executed).

- This command registers the management information package backup execution server.
- The command deletes the management information package backup execution server.
- The command displays information about the management information package backup execution server.

OPTIONS

None

Displays information about the management information package backup execution server.

-del

Deletes the management information package backup execution server.

-d *directoryName*

Specify the absolute path to the directory where the backup data is to be stored.
Only alphanumeric characters and underscores can be used in the directory name.

-c *managementServerEncoding*

Specify the encoding designated when the Management Server was installed. Specify ASCII.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Register the management information package backup execution server:

(assume that the path to the directory where backup data is to be stored is /stgbkupdata, and the encoding used on the Management Server is ASCII):

```
# /opt/FJSVswstf/bin/stgmgrinfoet -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/stgmgrinfoet -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/stgmgrinfoet
Directory      : /stgbkupdata
Manager code   : ASCII
#
```

NOTES

- This command can be executed only by root users.
- To change the backup data storage directory, first use the -del option to delete the current directory, then register a new backup data storage directory.
- Register only one Managed Server or Management Server as the management information package backup execution server.
- Set the name of the logical node in the environment variable SWSTGNODE when the Managed Server registered as the management information package backup execution server is in a clustered system.
- Create the directory specified in the -d option in advance.
- In an environment where a Managed Server with IPv4 address and a Managed server with IPv6 address coexist, register the Management Server as a management information package backup execution server.

9.2.12 stgmgrinfobkup (Management Information Package Backup Command)

NAME

stgmgrinfobkup - backup the management information

SYNOPSIS

```
/opt/FJSVswstf/bin/stgmgrinfobkup [-n serverName] [-t backupType]
```

DESCRIPTION

This command backup the management information stored on the Managed Server and the Management Server.

OPTIONS

-n *serverName*

Creates backups of the management information stored on the specified server.

If this option is not specified, the management information on all Managed Server (including the Management Server) is backed up.

-t backupType

Specifies the types of data that are to be backed up. Specify one of the values listed below. Do not specify more than one value in the -t option:

RegiData : Registry

MgrInfoData : Policy file for management information package backup

BkupData : Backup management

ReplData : Replication management

If this option is not specified, backups are created as defined in the policy file. For details on the policy file, refer to the section "[8.1.2.1 Preliminaries](#)".

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Back up the management information on all the Managed Server as defined in the policy file:

```
# /opt/FJSVswstf/bin/stgmgrinfobkup
swstf3014 Backup of management information is started. Start time=04/07/02 11:00:00, Server
name=acm01
swstf3020 Backup of management information was ended normally. End time=04/07/02 11:02:00, Server
name=acm01
swstf3014 Backup of management information is started. Start time=04/07/02 11:02:00, Server
name=acm02
swstf3020 Backup of management information was ended normally. End time=04/07/02 11:05:00, Server
name=acm02
swstf3021 The normal end of the backup of management information was carried out by all servers.
#
```

- Back up the management information on Managed Server as defined in the policy file:

```
# /opt/FJSVswstf/bin/stgmgrinfobkup -n acm01
swstf3014 Backup of management information is started. Start time=04/07/02 11:00:00, Server
name=acm01
swstf3020 Backup of management information was ended normally. End time=04/07/02 11:02:00, Server
name=acm01
swstf3021 The normal end of the backup of management information was carried out by all servers.
#
```

- Back up the registries on all of the Managed Server:

```
# /opt/FJSVswstf/bin/stgmgrinfobkup -t RegiData
swstf3014 Backup of management information is started. Start time=04/07/02 11:00:00, Server
name=acm01
swstf3020 Backup of management information was ended normally. End time=04/07/02 11:02:00, Server
name=acm01
swstf3014 Backup of management information is started. Start time=04/07/02 11:02:00, Server
name=acm02
swstf3020 Backup of management information was ended normally. End time=04/07/02 11:05:00, Server
name=acm02
swstf3021 The normal end of the backup of management information was carried out by all servers.
#
```

- Back up the backup management located on Managed Server:

```
# /opt/FJSVswstf/bin/stgmgrinfobkup -n acm01 -t BkupData
swstf3014 Backup of management information is started. Start time=04/07/02 11:00:00, Server
name=acm01
swstf3020 Backup of management information was ended normally. End time=04/07/02 11:02:00, Server
name=acm01
```

```
swstf3021 The normal end of the backup of management information was carried out by all servers.  
#
```

NOTES

- This command can be executed only by root users.
- Be sure to check the information (ie, a message indicating completion or an error) that is output by this command.
- This command can be executed only on the management information package backup execution server.
- Set the name of the logical node in the environment variable SWSTGNODE when the management information package backup execution server is in a clustered system.
- Do not execute other AdvancedCopy Manager commands or the Web Console operations while this command is running. Conversely, do not execute this command while another AdvancedCopy Manager command or a Web Console operation is running.
- Do not execute this command again if it is already running.
- If this command was executed under one of the following circumstances from a Managed Server on which AdvancedCopy Manager Version 14.2 or earlier is installed, backup of authentication data is not implemented.
 - When AuthData was specified in the -t option
 - When the value "YES" was specified for the AuthData key in the Management Server section of the policy file for batch backup of management data
- To save backup data to a tape or similar external medium, save all files and directories under the backup data storage directory.
- Registry backups are made only when RegiData is specified in the -t option. A registry backup is not 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.
- In an environment where a Managed Server with IPv4 address and a Managed Server with IPv6 address coexist, register the Management Server as a management information package backup execution server to execute management information package backup from the Management Server.
- If the version of Management Server is lower than that of Managed Server, back up the management information on Management Server.

9.2.13 stgmgrinfodisp (Management Information Package Backup Status Display Command)

NAME

stgmgrinfodisp - displays the status of execution of the stgmgrinfobkup

SYNOPSIS

```
/opt/FJSVswstf/bin/stgmgrinfodisp [-n serverName]
```

DESCRIPTION

This command displays the status of execution of the [stgmgrinfobkup](#) command.

The types of information listed below are output.

If no backups are created when the management information package backup command is executed or if the management information package backup command ends in error, then no data is output.

Item	Description
Server	Server name

Item	Description
RegiData	Date and time at which the management information package backup command created a backup of a registry (ie, the date and time that the management information package backup command was executed) In yy/mm/dd hh:mm:ss format
BkupData	Date and time at which the management information package backup command created a backup of the backup management (ie, the date and time that the management information package backup command was executed) yy/mm/dd hh:mm:ss
ReplData	Date and time at which the management information package backup command created a backup of replication management (ie, the date and time that the management information package backup command was executed) yy/mm/dd hh:mm:ss
MgrInfoData	Date and time at which the management information package backup command created a backup of the policy file for management information package backup (ie, the date and time that the management information package backup command was executed) yy/mm/dd hh:mm:ss

OPTIONS

-n *serverName*

Specifies the name of the server where the backup status is to be output.

If this option is not specified, the status of backups on the Managed Server for which the management information package backup command was executed is displayed.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Output the information about all the Managed Server:

```
# /opt/FJSVswstf/bin/stgmgrinfodisp
Server:      acm01
  RegiData:  04/08/01 09:00:00
  ReplData:  04/08/02 16:32:22
Server:      acm02
  RegiData:  04/08/01 09:00:00
  BkupData:  04/08/03 16:00:00
#
```

- Output the information about Managed Server named acm01:

```
# /opt/FJSVswstf/bin/stgmgrinfodisp -n acm01
Server:      acm01
  RegiData:  04/08/01 09:00:00
  ReplData:  04/08/02 16:32:22
#
```

NOTES

- This command can be executed only by root users.
- This command can be executed only on the management information package backup execution server.

- When the management information package backup execution server is in a clustered system, set the name of the logical node in the environment variable SWSTGNODE.

9.2.14 stgxfwcmsetsdx (SDX Object Management Unit Configuration Command)

NAME

stgxfwcmsetsdx - sets SDX object management unit

SYNOPSIS

```
/opt/FJSVswstf/bin/stgxfwcmsetsdx [-slice | -volume]
```

DESCRIPTION

The unit of operation of the SDX object treated with AdvancedCopy Manager (Refer to "3.8.1 Units of SDX Object Operations") is set. Moreover, the unit of operation of a present SDX object is displayed.

The display form is 'Unit of management (SWSTGNODE = <configuration value>).'

- Unit of management
- Slice in case of each slice
- Volume in case of each logical volume

One of these values are used value-environment variable SWSTGNODE is set. The value is Null when environment variable SWSTGNODE is not set.

OPTIONS

none

The unit of present of the SDX object management is displayed.

-slice

The unit of management of the SDX object treated with AdvancedCopy Manager is set to the unit of the slice.

-volume

The unit of management of each logical volume of the SDX object treated with AdvancedCopy Manager is set.

EXIT STATUS

- =0: Completed successfully
- >0: Terminated abnormally

EXAMPLES

- The unit of present of the SDX object management in non-cluster environment is displayed:

```
# /opt/FJSVswstf/bin/stgxfwcmsetsdx
Volume (SWSTGNODE=null)
#
```

- The unit of the SDX object of the Management Server service (The logical node name: aserv) is set to the unit of the slice:

```
# SWSTGNODE=aserv; export SWSTGNODE
# /opt/FJSVswstf/bin/stgxfwcmsetsdx -slice
stgxfwcmsetsdx completed (SWSTGNODE=aserv)
#
```

NOTES

- This command can be executed only by root users.

- When the cluster is operated, each Managed Server service is set.

9.2.15 stgxfwcmsetmode (Information Collection Mode Configuration Command)

NAME

stgxfwcmsetmode - sets the collection mode for information

SYNOPSIS

```
/opt/FJSVswstf/bin/stgxfwcmsetmode -virt {none | zone | disp}
```

DESCRIPTION

For operation in a virtual environment, this command sets the Virtual Environment Support Function. It also displays the currently set contents of the Virtual Environment Support Function.

When this command is completed successfully, "(SWSTGNODE=*configuration variable*)" is displayed at the end of the line after the command. The following string is displayed in "*configuration variable*":

- If environment variable SWSTGNODE is set: value of environment variable SWSTGNODE
- If environment variable SWSTGNODE is not set: null

OPTIONS

-virt {none | zone | disp}

Specifies whether to use the Virtual Environment Support Function. The following values can be specified:

none : The Virtual Environment Support Function is not used.

zone : The Oracle Solaris Zones Support Function is used.

disp : Check the current setup status of Virtual Environment Support Function.

Information Output When "disp" Is Specified

virt : none The Virtual Environment Support Function is not used.

virt : zone The Oracle Solaris Zones Support Function is used.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- To specify that the Oracle Solaris Zones Support Function is used:

```
# /opt/FJSVswstf/bin/stgxfwcmsetmode -virt zone
stgxfwcmsetmode completed (SWSTGNODE=nodeagt)
#
```

- To check the current setup status of Virtual Environment Support Function:

```
# /opt/FJSVswstf/bin/stgxfwcmsetmode -virt disp
virt : zone (SWSTGNODE=nodeagt)
#
```

9.2.16 stgrepcnv (Repository Data Migration Command)

NAME

stgrepcnv - migrates the repository data

SYNOPSIS

```
/opt/FJSVswstf/bin/stgrepocnv -d directoryName | fileName
```

DESCRIPTION

This command recovers the migration target repository data to the current version repository. It is necessary to execute this command in the Management Server.

OPTION

-d *directoryName* | *fileName*

When migrating from ETERNUS SF Manager Version 16.0 or earlier, specify the absolute path name of the directory where the migration target repository data is stored.

When migrating from ETERNUS SF Manager Version 16.1 or later, specify the absolute path name for the migration repository data backup file.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

```
/opt/FJSVswstf/bin/stgrepocnv -d /var/tmp/backdir  
.....  
stgrepocnv completed
```

NOTES

- This command can be executed only by root users.
- For Management Server transactions, execute this command at a currently operating server. The environment variable SWSTGNODE need not be set.

9.2.17 stgforcedelsrv (Managed Server Forcibly Deletion Command)

NAME

stgforcedelsrv - deletes the information of Managed Server forcibly

SYNOPSIS

```
/opt/FJSVswstf/bin/stgforcedelsrv -n managedServer -b beforeFile -a afterDir
```

DESCRIPTION

This command deletes the Managed Server information forcibly. This command must be executed on a Management Server.

As this command deletes the backup/replication settings of the Managed Server, it must be performed with caution. Before performing the command always refer to "[8.6 Action When Managed Server Information Cannot Be Deleted](#)".

OPTIONS

-n *managedServer*

Specify the server name of Managed Server to be deleted forcibly.

-b *beforeFile*

Specify the full path name of the file containing the database backup.

The command deletes the Managed Server information forcibly from the database in the file specified with the -b option and stores the database after deletion in the directory specified with -a option.

-a afterDir

Specify the full path name of the directory that stores the database after the Managed Server information has been deleted. Note the following important points:

- If the specified directory does not exist, it must be created before executing the command.
- Specify an empty directory, containing no subdirectories or files.
- Specify a directory that does not disappear after the server has rebooted.

EXIT STATUS

- =0: Completed successfully
- >0: Terminated abnormally

EXAMPLES

Delete the information of Managed Server (Server1). After deleting its information, the database is created under the "/var/tmp/afterDB" directory.

```
# /opt/FJSVswstf/bin/stgforcedelsrv -n Server1 -b /var/tmp/beforeDB/db_file.txt -a /var/tmp/afterDB
Begin the forcing deletion? (y/n) : y
stgforcedelsrv completed
#
```

NOTES

- This command can be executed only by root users.
- The Management Server information cannot be deleted.
- Do not perform AdvancedCopy Manager operations (backup, replication, modifications to the operating environment, etc.) while executing this command.
- If the command terminates abnormally, temporary files may remain under the directory specified with -a option. When the command terminates abnormally, delete all the files under the directory specified with the -a option before executing the command again.
- Several instances of the command cannot be executed simultaneously.

9.3 Replication Management Commands

This section describes operations using Replication Management Commands.

Note

Only a root user can execute all the commands of replication management. When command execution is carried out by other users, a message "swsrp0601 This Command cannot be run. This is because the user is not a root user." is displayed, and the command aborts.

Note

Operating Environment for Command Execution

The table below lists the requirements for the operating environment in which to execute the different commands.

Table 9.6 Operating Environment When Executing Replication Management Commands

Command	Operating Environment of AdvancedCopy Manager
swsrpcancel	To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating. When the Managed Server runs on Windows and is in EC operation, Windows services required

Command	Operating Environment of AdvancedCopy Manager
	by the copy destination Managed Server must be running (except when the -m or -t option is specified). For inter-server replication in the equivalency maintain state, the associated server must be operating (except when the -m option is specified).
swsrpchsyc	To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating.
swsrpdelvol	To execute this command on a Managed Server, the daemons or services of the Management Server must be operating. To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating. For inter-server replication, the daemons or services of the Managed Server in communication must be operating. (except when the -d option is specified).
swsrprebuffstat	To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating.
swsrprebuffset	To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating.
swsrpmake	To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating. When the Managed Server runs on Windows and is engaged in replication on single server and EC operation, Windows services required by the Managed Server must be running. For inter-server replication, the daemons or services of the Managed Server in communication must be operating. (except when the -d option is specified).
swsrprecoverres	To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating. When the Managed Server runs on Windows, Windows services required by the Managed Server must be running. To specify the -r option, Windows services required by the Management Server must be running.
swsrprevsync	To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating.
swsrpsetvol	To execute this command on a Managed Server, Windows services required by the Management Server must be running. To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating. For inter-server replication, the daemons or services of the Managed Server in communication must be operating.
swsrpstartsync	To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating. When the Managed Server runs on Windows, Windows services required by the Managed Server must be running. For inter-server replication, the associated server must be running (except a case where the -m option is specified).
swsrpstat	To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed Server must be operating.
swsrpvolinfo	To execute this command on a Management Server with the -h option specified, the daemons or services of the Management Server and target Managed 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 environment variable SWSTGNODE. Specify the environment variable as follows:

```
# SWSTGNODE=logicalNodeName
# export SWSTGNODE
```

2. Execute the command.

The table below lists the commands available in cluster operation, indicates whether environment variables must be configured and shows the nodes on which commands are executable. For notes and cautions on command execution, refer to the Remarks column.

Table 9.7 Commands When Cluster Operation Is Used

Command	Environment Variable	Command Execution Node	Remarks
swsrpsetvol	Required	Transaction operation node	-
swsrpvoinfo	Required	Transaction operation node	-
swsrpdelvol	Required	Transaction operation node	-
swsrprebuffstat	Required	Transaction operation node	-
swsrprebuffset	Required	Transaction operation node	-
swsrpstartsync	Required	Transaction operation node	-
swsrpmake	Required	Transaction operation node	-
swsrpstat	Required	Transaction operation node	-
swsrpcancel	Required	Transaction operation node	-
swsrpchsync	Required	Transaction operation node	-
swsrprevsync	Required	Transaction operation node	-
swsrprecoverres	Required	Transaction operation node	-

9.3.1 Environment Definition Commands

This section describes environment definition commands for replication management.

9.3.1.1 swsrpsetvol (Replication Volume Information Setting Command)

NAME

swsrpsetvol - sets the replication source volumes and replication destination volumes

SYNOPSIS

To Execute on Managed Server

For Replication on Single Server

```
/opt/FJSSVswsrp/bin/swsrpsetvol [-n] [-u] [-Xvtoc] [-Xgroup groupName]
originalVolumeName replicaVolumeName
```

For Server-to-server Replication

```
/opt/FJSVswsrp/bin/swsrpsetvol [-n] [-u] [-Xvtoc] -o (ORG|REP|BOTH) [-Xgroup groupName]
originalVolumeName replicaVolumeName
```

To Execute on Management Server

For Replication on Single Server

```
/opt/FJSVswsrp/bin/swsrpsetvol [-h serverName] [-n] [-u] [-Xvtoc] [-Xgroup groupName]
originalVolumeName replicaVolumeName
```

For Server-to-server Replication

```
/opt/FJSVswsrp/bin/swsrpsetvol [-h serverName] [-n] [-u] -o (ORG|REP|BOTH) [-Xvtoc] [-Xgroup
groupName] originalVolumeName replicaVolumeName
```

DESCRIPTION

This command sets the replication source volumes and replication destination volumes that are used in a replication operation. Execute this command to declare that the replication operation is performed using the specified replication volumes.

When using the Storage Cluster Continuous Copy Function, configure the replication source and replication destination volume information that are commonly used by the Primary Storage and Secondary Storage.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdisprv](#) command.

-u

Specifies the direction of copying.

If this option is specified, copying is performed only in the direction from replication source volume to a replication destination volume. Copying from a replication destination volume to a replication source volume is prohibited.

If this option is omitted, copying between the replication source and replication destination volumes is bidirectional.

-o {ORG|REP|BOTH}

When performing the server-to-server replication, specify the operation server(s) authorized to manage replication.

Specify one of the following for this operand:

ORG : Sets the replication source server as the operation server.

REP : Sets the replication destination server as the operation server.

BOTH : Sets both the copy-source server and copy-destination server as operation servers.

When BOTH has been specified, all commands can be executed from either the replication source server or the replication destination server.

When ORG or REP was specified, the non-operation server cannot issue the following commands:

- [swsrpdelvol](#)
- [swsrpstartsync](#)
- [swsrpmake](#)
- [swsrpcancel](#)
- [swsrpchsync](#)
- [swsrprevsync](#)

However, in the following cases, the non-operation server can issue the above commands.

- When specifying the -d option with the [swsrpdelvol](#) command.

- When specifying the -c option with the [swsrpcancel](#) command.
- When specifying the -b option with the [swsrpcancel](#) command.

-n

If this option is omitted, replication volume information cannot be set when the sizes of the replication source volume and replication destination volume are different.

When this option is specified, the sizes of the replication source volume and replication destination volume are not checked. Use this option when volume sizes are different (for example, for replication between different OSs).

When the volume sizes of the replication source volume and replication destination volume are different, copying is performed as follows:

The address at which copying starts is the first sector of the replication destination volume.

The copy data size is the size of the smaller volume.

Therefore, ensure that the size of the copy destination volume is the same size or larger than the size of the copy source volume.

This option cannot be used for SDX object logical volumes (PRIMECLUSTER GD Snapshot linkage function).

-Xgroup *groupName*

Specifies that the copy source and copy destination volumes are registered in a group.

If the specified group does not exist, a new group is created. Otherwise, the specified copy source and copy destination volumes are added to an existing group.

The characters that can be used for the group name are single-byte alphanumeric characters, the minus sign (-), underscore (_), number sign (#) and period (.). The first character must be a single-byte alphanumeric character. Up to 64 characters can be used.

-Xvtoc

If this option is not specified, Solaris slices that include VTOC cannot be set as replication source volumes or replication destination volumes.

Specification of this option avoids checking to see whether or not replication source and replication destination volumes include VTOC.

This option is specified for the following replication source and replication destination volumes:

- Under Solaris, VTOC-labeled disk slices that contain VTOC
- Under Solaris, a VTOC-labeled disk second slice that indicates the entire disk
- Under Solaris, an EFI-labeled disk MPxIO device that indicates the entire disk

This option is ignored if it is set for volumes other than the above.

OPERANDS

originalVolumeName

Specifies a replication source volume name.

For the volume names of other servers, use the format "*Volume-Name@Managed-Server-Name*".

Specifies the [AdvancedCopy Manager device name](#) for the *Volume-Name*.

replicaVolumeName

Specifies a replication destination volume name.

For the volume names of other servers, use the format "*Volume-Name@Managed-Server-Name*".

Specifies the [AdvancedCopy Manager device name](#) for the *Volume-Name*.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Set replication volume information by specifying volume/dev/dsk/c1t2d0s3 of Managed Server SRC-SV as the replication source volume and volume /dev/dsk/c1t2d0s3 of Managed Server TARG-SV as the replication destination volume. In addition, specify one-

way copying for the copy direction and specify the replication source server (SRC-SV) as the operation server. The example below shows the command executed on SRC-SV:

```
# /opt/FJSVswsrp/bin/swsrpsetvol -u -o ORG /dev/dsk/c1t2d0s3 /dev/dsk/c1t2d0s3@TARG-SV
swsrpsetvol completed
#
```

- When the same settings are made from TARG-SV, the example is as follows (be careful with the volume name):

```
# /opt/FJSVswsrp/bin/swsrpsetvol -u -o ORG /dev/dsk/c1t2d0s3@SRC-SV /dev/dsk/c1t2d0s3
swsrpsetvol completed
#
```

NOTES

- Replication volume information cannot be set if:
 - Replication volume information has already been set.
 - The specified replication source volume and replication destination volume are already set as the replication destination volume and the replication source volume respectively.
 - The -n option is not specified when the sizes of the replication source volume and the replication destination volume are different.
 - Two volumes that have the same box information (ie, box identifier, OLU number, EXTENT start position, and EXTENT size) as the box information of the specified two volumes have already been set as replication volumes.
 - The replication destination volume has been registered as a backup volume by the backup management function.
 - The replication source volume has been registered as a backup volume by the backup management function and bi-directional copying is specified as a replication volume attribute.
 - The specified replication destination volume and replication source volume are volumes belonging to other servers (at least one volume must be a volume on the local server).
 - If the replication source volume and the replication destination volume are located on another ETERNUS Disk storage system, the remote copy function of the ETERNUS Disk storage system is not correctly set.
 - Communication with the Management Server fails.
 - Communication with a destination server fails for server-to-server replication.
 - The IP version registered on AdvancedCopy Manager differs between the replication source server and the replication destination server.
 - When the replication source volume and the replication destination volume are a replication management function and a logical volume of unsupported VxVM (For details on the logical volume that can be supported by the replication management, refer to "1.4.1 Normal Devices").
 - The license for the ETERNUS Disk storage system that contains the specified replication source/target volume is one of the following:
 - The license is not registered in the license management function.
 - The license is expired.
 - When the specified replication source and replication destination volumes are neither of the following combinations:
 - Both the replication source volume and the replication destination volume are TFOV.
 - Both the replication source volume and the replication destination volume are non-TFOV.
 - The specified replication source and destination volumes belong to different TFO groups.
 - The specified replication source and destination volumes are in the failover or in the failback state.
- Note the following points about creating copy source volumes and copy destination volumes under Solaris:
 - When creating multiple volumes (slices) on the same disk, volume areas must not overlap.

- If slices contain VTOC, check the notes in "Slices that contain VTOC" under "General notes".
- When you change physical information on volumes (ie, partition size, drive letter, and box information on ETERNUS Disk storage system device using such functions as an OS management tool (format command etc) or ETERNUS Web GUI provided by ETERNUS Disk storage system, PRIMECLUSTER GD, 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 [swsrpdelvol](#) command.
 2. Refer to "Add/Reload Device" in the *ETERNUS SF Web Console Guide* to reload the configuration. Otherwise the [stgxfwcmsetdev](#) command is executed in a Management Server.
 3. Reset the replication volume information using the [swsrpsetvol](#) command.
- When you use a replication destination volume by mounting after a replication volume is created in another partition, 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.
- Group configurations cannot be made in the group configuration conditions are not met.
- When SDX objects of PRIMECLUSTER GD are used: Refer to "[Configuration and Conditions of Unsupported SDX Objects](#)", and "[10.1.5.2 Notes on SDX Object Operations \(Replication Management\)](#)".
- This command cannot be executed while any of the following commands are running:
 - [swsrpsetvol](#)
 - [swsrpdelvol](#)
 - [swsrpstartsync](#)
 - [swsrpmake](#)
 - [swsrpstat](#)
 - [swsrpcancel](#)
 - [swsrpchsyc](#)
 - [swsrprevsync](#)
 - [swsrprecoverres](#)

9.3.1.2 swsrpvoinfo (Replication Volume Information Display Command)

NAME

swsrpvoinfo - displays the replication volume information

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsrp/bin/swsrpvoinfo [-L] [originalVolumeName]
```

To Execute on Management Server

```
/opt/FJSVswsrp/bin/swsrpvoinfo [-h serverName] [-L] [originalVolumeName]
```

DESCRIPTION

This command displays the replication volume information registered by the [swsrpsetvol](#) command. The following information is displayed.

Title	Description
Server	Displays a Managed Server name.

Title	Description
Original-Volume	Displays a replication source volume name.
Size	Displays the volume size. For a logical volume of VxVM, the size of the public area (slice) that makes up the volume is displayed instead of the volume size.
Replica-Volume	Displays a replication destination volume name.
Copy	Displays the copying direction. uni-direction : One-way copying from a replication source volume to a replication destination volume is allowed. bi-direction : Bidirectional copying is allowed.
Op-Server	Displays the operation server of a pair of replication volumes. original : The replication source server is the operation server. replica : The replication destination server is the operation server. both : Sets both the copy-source server and copy-destination server as operation servers.
Group (Only when the -L option is specified)	Displays the group name. "----" is displayed if the copy source and destination volumes are not registered in a group.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdisprv](#) command.

-L

Specifies that information is displayed using the extended format.

For the items displayed by this option, refer to the table in the description above.

The items that are displayed may be expanded to support new functions provided by future versions.

OPERANDS

originalVolumeName

Specifies a replication source volume name.

For the volume names of other servers, use the format "Volume-Name@Managed-Server-Name."

Specifies the [AdvancedCopy Manager device name](#) for the volume name.

Displays all the replication volume information for the specified replication source volume. If this operand is omitted, all the replication volume information existing on the Managed Server on which this command is executed is displayed.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Direct that all the replication volume information on a Management Server is to be displayed on Managed Server SRC-SV:

```
# /opt/FJSVswsrp/bin/swsrpvoinfo -h SRC-SV
Server Original-Volume          Size      Replica-Volume          Size      Copy
Op-Server
SRC-SV /dev/dsk/clt2d0s3@SRC-SV  4.0 Gbyte /dev/dsk/clt2d0s3@TARG-SV  4.0 Gbyte uni-direction
original
```

```

SRC-SV /dev/dsk/c1t2d1s3@SRC-SV      4.0 Gbyte /dev/dsk/c1t2d1s3@TARG-SV  5.0 Gbyte bi-direction
original
SRC-SV /dev/dsk/c1t2d2s3@SRC-SV      9.1 Gbyte /dev/dsk/c1t2d2s3@TARG-SV  9.1 Gbyte bi-direction
replica
SRC-SV /dev/dsk/c1t2d2s4@SRC-SV      7.5 Gbyte /dev/dsk/c1t2d2s4@TARG-SV  7.5 Gbyte bi-direction
both
#

```

NOTES

- Replication volume information cannot be displayed if:
 - The specified volume has not been set as a replication source volume
- This command cannot be executed while the following command is running:
 - [swsrprecoverres](#)

9.3.1.3 swsrpdelvol (Replication Volume Information Deletion Command)

NAME

swsrpdelvol - deletes the replication volume

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsrp/bin/swsrpdelvol [-e] [-d] originalVolumeName replicaVolumeName
```

To Execute on Management Server

```
/opt/FJSVswsrp/bin/swsrpdelvol [-h serverName] [-e] [-d] originalVolumeName replicaVolumeName
```

DESCRIPTION

This command deletes the replication volume information that has been set.

If copy processing is being performed on a replication volume that is to be deleted, execute the command after the copy processing has stopped.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

When performing the server-to-server replication, specify the name of the operation server configured with the [swsrpsetvol](#) command.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdispsv](#) command.

-e

Emergency mode is available for operations.

Specifies deleting replica volume information without checking the execution state of an Advanced Copy. Use this option only when it is clear that the Advanced Copy has not executed.

Deletes replication volume information without confirming the release of the master-proxy relationship of an SDX object when the function in linkage with PRIMECLUSTER GD Snapshot is used.

-d

Specifies that no communication should take place between the linked servers in server-to-server replication. 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 server-to-server replication.
The option can be used on the server not specified as the operation server.

OPERANDS

originalVolumeName

Specifies a replication source volume name.
For the volume names of other servers, use the format "*Volume-Name@Managed-Server-Name*".
Specifies the [AdvancedCopy Manager device name](#) for the *Volume-Name*.

replicaVolumeName

Specifies a replication destination volume name.
For the volume names of other servers, use the format "*Volume-Name@Managed-Server-Name*".
Specifies the [AdvancedCopy Manager device name](#) for the *Volume-Name*.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Delete replication volume information on Managed Server SRC-SV:

```
# /opt/FJVSwsrp/bin/swsrpdelvol /dev/dsk/c1t2d0s3 /dev/dsk/c1t2d0s3@TARG-SV  
swsrpdelvol completed  
#
```

NOTES

- Replication volume information cannot be deleted if:
 - The specified replication source volume and the replication destination volume are not set as replication volumes.
 - The copy processing (synchronization processing and snapshot processing) is performed between replication volumes.
 - For server-to-server replication, the Managed Server that executed the command is not the operation server of the replication volume.
 - Communication with a Management Server fails.
 - For server-to-server replication, communication with a destination server fails.
- However, OPC sessions are neither confirmed nor cancelled by execution in emergency operation mode. If an OPC session is in the error suspended state, then unmanaged sessions remain. In this event, ETERNUS Web GUI must be used to cancel these sessions.
- If a logical volume of VxVM 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 ETERNUS Disk storage system, 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.
- When SDX objects of PRIMECLUSTER GD are used: Refer to "[Configuration and Conditions of Unsupported SDX Objects](#)", and "[10.1.5.2 Notes on SDX Object Operations \(Replication Management\)](#)".
- This command cannot be executed while any of the following commands are running:
 - `swsrpsetvol`
 - `swsrpdelvol`
 - `swsrpstartsync`
 - `swsrpmake`

- [swsrpstat](#)
- [swsrpcancel](#)
- [swsrpchsyc](#)
- [swsrprevsync](#)
- [swsrprecoverres](#)

9.3.1.4 swsrprebuffstat (REC Buffer Information Display Command)

NAME

swsrprebuffstat - displays information for the REC buffer

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsrp/bin/swsrprebuffstat [-Xaltpath] [-L] volumeName
```

To Execute on Management Server

```
/opt/FJSVswsrp/bin/swsrprebuffstat [-h server-Name] [-Xaltpath] [-L] volumeName
```

DESCRIPTION

This command displays information for the REC buffer used by ETERNUS Disk storage system when REC is performed in Consistency mode.

The following information is displayed.

Title	Description
BoxID	The box identifier for the box connected to the server where the command was executed.
ID	The ID for the REC buffer.
Use	The attribute (usage) for the REC buffer. SEND : REC buffer for sending data RECV : REC buffer for receiving data
Rate	The usage rate of the REC buffer as a percentage. If "Use" Is SEND Displays the usage rate of the REC buffer as a percentage. This is never 0%, even if REC is not being performed in Consistency mode. If "Use" Is RECV "0%" is always displayed. "----" is displayed if the REC buffer cannot be used.
RemoteBoxID	The box identifier for the box performing a remote copy.
DiskBuff	Displays the REC Disk buffer status only when -L option is specified: - --- : REC buffer is used for receiving - active : REC Disk buffer is available for use - invalid : REC Disk buffer is not set - warning(code=xx) : REC Disk buffer is available for use but there is a disk fault - inactive(code=xx) : REC Disk buffer cannot be used

Title	Description
	<ul style="list-style-type: none"> - not_support : REC Disk buffer is not supported <p>The following codes are displayed at code=xx in the case of warning or inactive:</p> <ul style="list-style-type: none"> - 04 : Rebuild or other recovery operation is in progress (REC Disk buffer is available for use) - 05 : RAID that comprises the REC Disk buffer has a fault - 06 : The REC Disk buffer RAID group is being formatted - 07 : Copy destination does not support REC Disk buffer - 08 : A RAID that comprises the REC Disk buffer has no redundancy
DiskRate	<p>Displays the busy rate as a percentage if the REC Disk buffer state is "active" or "warning" only when -L option is specified. If the REC Disk buffer is not being used, 0% is displayed.</p> <p>In cases other than the above, "---" is displayed.</p>

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdispsrv](#) command.

-Xaltpath

When replication volume is applying per logic volume by the SDX object, an alternative path is specified as an operand in the form of the following.

/dev/sfdsk/<class name>/dsk/<volume name>:<DEVNAM value for sdxinfo>

-L

This option specifies to display in extended format.

For the items displayed by this option, refer to the table in the description above.

Note that additional items corresponding to new functions in future versions might be displayed in future when this option is specified.

OPERANDS

volumeName

Specifies the name of a volume that has been registered as a replication volume.

You can only specify volumes that are connected to the Managed Server where the command is executed.

When replication volume is applying per logic volume by the SDX object, it is necessary to use the -Xaltpath option and to specify an alternative path.

Specifies the [AdvancedCopy Manager device name](#) as the volume name.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Display information for the REC buffer used by the REC, and replication is performed within the Managed Server "Server1":

```
# /opt/FJSVswsrp/bin/swsrprecbuffstat /dev/FJSVmplb/dsk/mplb2048s3@Server1 -L
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use Rate RemoteBoxID DiskBuff DiskRate
1 SEND 25% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB active 50%
```

```
# /opt/FJSVswsrp/bin/swsrprecbuffstat /dev/FJSVmplb/dsk/mplb3048s3@Server1
BoxID = BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
ID Use Rate RemoteBoxID DiskBuff DiskRate
1 RECV 0% AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA --- ---
#
```

- Display information for the REC buffer used by the REC in an environment where replication is being performed between Managed Server "Server1" and Managed Server "Server2":

[Command executed on Managed Server "Server1"]

```
# /opt/FJSVswsrp/bin/swsrprecbuffstat /dev/FJSVmplb/dsk/mplb2048s3@Server1 -L
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use Rate RemoteBoxID DiskBuff DiskRate
1 SEND 25% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB active 0%
2 RECV 0% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB --- ---
#
```

[Command executed on Managed Server "Server2"]

```
# /opt/FJSVswsrp/bin/swsrprecbuffstat /dev/FJSVmplb/dsk/mplb2148s3@Server2 -L
BoxID = BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
ID Use Rate RemoteBoxID DiskBuff DiskRate
1 RECV 0% AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA --- ---
2 SEND 25% AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA active 0%
#
```

- Display information for the REC buffer used by the REC if a box where the copy source volume is located has been set up to make remote copies to multiple boxes, and replication is being performed between Managed Server "Server1" and Managed Server "Server2":

[Command executed on Managed Server "Server1"]

```
# /opt/FJSVswsrp/bin/swsrprecbuffstat /dev/FJSVmplb/dsk/mplb2048s3@Server1 -L
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use Rate RemoteBoxID DiskBuff DiskRate
1 SEND 25% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB warning(code=08) 0%
2 RECV 0% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB --- ---
3 SEND 12% CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC active 50%
4 RECV 0% CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC --- ---
5 SEND 12% DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD invalid --- ---
6 RECV 0% EEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEE --- ---
7 SEND 12% FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF inactive(code=05) --- ---
8 SEND 12% GGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGG active 0%
#
```

- Display information for the REC buffer used by the REC, and replication is being performed between Managed Server "Server1" and Managed Server "Server2":
(The SDX object is in a two-box mirror configuration.)

[Command executed on Managed Server "Server1"]

```
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class1/dsk/volume1:mplb2048
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use Rate RemoteBoxID
1 SEND 25% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
2 RECV 0% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class1/dsk/volume1:mplb3048
BoxID = CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
ID Use Rate RemoteBoxID
1 SEND 12% DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD
2 RECV 0% DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD
#
```

[Command executed on Managed Server "Server2"]

```

# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class2/dsk/volume1:mp1b2148
BoxID = BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
ID Use Rate RemoteBoxID
1 RECV 0% AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
2 SEND 12% AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class2/dsk/volume1:mp1b3148
BoxID = DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD
ID Use Rate RemoteBoxID
1 RECV 0% CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
2 SEND 12% CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
#

```

NOTES

- Information is displayed only for boxes that contain replication volumes for the server where the command was executed.
If inter-box replication is being performed between servers, then the command must be executed on both servers involved in the replication in order to check information for the REC buffers in both boxes.
- It takes a short time for changes to REC buffer settings to take effect, during which time the buffer may become unusable (the Rate column is displayed as "----"). In this case, wait about one minute and then execute the command again.
- This command cannot be executed while the following command is running:
 - [swsrprecoverres](#)
- The REC buffer information cannot be displayed by using TFOV.

9.3.1.5 swsrprecbuffset (REC Buffer Settings Change Command)

NAME

swsrprecbuffset - changes the settings of the REC buffer

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsrp/bin/swsrprecbuffset -Xbuffid buffer-id -Xbuffuse {SEND|RECV} [-Xaltpath] volumeName
```

To Execute on Management Server

```
/opt/FJSVswsrp/bin/swsrprecbuffset [-h serverName] -Xbuffid buffer-id -Xbuffuse {SEND|RECV} [-Xaltpath] volumeName
```

DESCRIPTION

This command changes the settings of the REC buffer used by ETERNUS Disk storage system when REC is performed in Consistency mode.

For information about the setting method of REC buffer, refer to "User's Guide Setup / Maintenance" of ETERNUS Web GUI.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdsrv](#) command.

-Xbuffid *buffer-id*

Specifies the ID for the REC buffer.

-Xbuffuse {SEND|RECV}

Specifies the usage for the REC buffer.

SEND : REC buffer for sending data
RECV : REC buffer for receiving data

-Xaltpath

When replication volume is a logical volume in the SDX object, an alternative path is specified as an operand in the following form.

/dev/sfdsk/<class name>/dsk/<volume name>:<DEVNAM value for sdxinfo>

OPERANDS

volumeName

Specifies the name of a volume that has been registered as a replication volume.

You can only specify volumes that are connected to the Managed Server where the command is executed.

When replication volume is applying per logic volume by the SDX object, it is necessary to use the -Xaltpath option and to specify an alternative path.

Specifies the [AdvancedCopy Manager device name](#) as the volume name.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Change the attributes for the REC buffer used by the REC, and replication is being performed within the Managed Server "Server1":

```
# /opt/FJSVswsrp/bin/swsrprebuffstat /dev/FJSVmplb/dsk/mplb2048s3@Server1
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use Rate RemoteBoxID
1 SEND 25% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
# /opt/FJSVswsrp/bin/swsrprebuffset -Xbuffid 1 -Xbuffuse RECV /dev/FJSVmplb/dsk/
mplb2048s3@Server1
swsrprebuffset completed
# /opt/FJSVswsrp/bin/swsrprebuffstat /dev/FJSVmplb/dsk/mplb2048s3@Server1
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use Rate RemoteBoxID
1 RECV 0% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
# /opt/FJSVswsrp/bin/swsrprebuffstat /dev/FJSVmplb/dsk/mplb3048s3@Server1
BoxID = BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
ID Use Rate RemoteBoxID
1 RECV 0% AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
# /opt/FJSVswsrp/bin/swsrprebuffset -Xbuffid 1 -Xbuffuse SEND /dev/FJSVmplb/dsk/
mplb3048s3@Server1
swsrprebuffset completed
# /opt/FJSVswsrp/bin/swsrprebuffstat /dev/FJSVmplb/dsk/mplb3048s3@Server1
BoxID = BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
ID Use Rate RemoteBoxID
1 SEND 12% AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
#
```

- Change the connected box for the REC buffer of the copy source volume , and replication is being performed between Managed Server "Server1" and Managed Server "Server2":

[Command executed on Managed Server "Server1"]

```
# /opt/FJSVswsrp/bin/swsrprebuffstat /dev/FJSVmplb/dsk/mplb2048s3@Server1
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use Rate RemoteBoxID
1 SEND 25% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
2 RECV 0% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
# /opt/FJSVswsrp/bin/swsrprebuffset -Xbuffid 1 -Xbuffuse SEND /dev/FJSVmplb/dsk/
mplb3048s3@Server1
```

```

swsrprecbuffset completed
# /opt/FJSVswsrp/bin/swsrprecbuffset -Xbuffid 2 -Xbuffuse RECV /dev/FJSVmplb/dsk/
mplb3048s3@Server1
swsrprecbuffset completed
# /opt/FJSVswsrp/bin/swsrprecbuffstat /dev/FJSVmplb/dsk/mplb3048s3@Server1
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use Rate RemoteBoxID
1 SEND 12% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
2 RECV 0% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
#

```

- Change the attributes for the REC buffer used by the REC, and replication is being performed between Managed Server "Server1" and Managed Server "Server2":

(The SDX object is in a two-box mirror configuration.)

[Command executed on Managed Server "Server1"]

```

# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class1/dsk/volume1:mplb2048
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use Rate RemoteBoxID
1 SEND 25% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class1/dsk/volume1:mplb3048
BoxID = CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
ID Use Rate RemoteBoxID
1 SEND 12% DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD
# /opt/FJSVswsrp/bin/swsrprecbuffset -Xbuffid 1 -Xbuffuse RECV -Xaltpath /dev/sfdsk/class1/dsk/
volume1:mplb2048
swsrprecbuffset completed
# /opt/FJSVswsrp/bin/swsrprecbuffset -Xbuffid 1 -Xbuffuse RECV -Xaltpath /dev/sfdsk/class1/dsk/
volume1:mplb3048
swsrprecbuffset completed
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class1/dsk/volume1:mplb2048
BoxID = AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
ID Use Rate RemoteBoxID
1 RECV 0% BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class1/dsk/volume1:mplb3048
BoxID = CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
ID Use Rate RemoteBoxID
1 RECV 0% DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD

```

[Command executed on Managed Server "Server2"]

```

# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class2/dsk/volume1:mplb2148
BoxID = BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
ID Use Rate RemoteBoxID
1 RECV 0% AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class2/dsk/volume1:mplb3148
BoxID = DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD
ID Use Rate RemoteBoxID
1 RECV 0% CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
# /opt/FJSVswsrp/bin/swsrprecbuffset -Xbuffid 1 -Xbuffuse SEND -Xaltpath /dev/sfdsk/class2/dsk/
volume1:mplb2148
swsrprecbuffset completed
# /opt/FJSVswsrp/bin/swsrprecbuffset -Xbuffid 1 -Xbuffuse SEND -Xaltpath /dev/sfdsk/class2/dsk/
volume1:mplb3148
swsrprecbuffset completed
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class2/dsk/volume1:mplb2148
BoxID = BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
ID Use Rate RemoteBoxID
1 SEND 12% AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
# /opt/FJSVswsrp/bin/swsrprecbuffstat -Xaltpath /dev/sfdsk/class2/dsk/volume1:mplb3148
BoxID = DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD
ID Use Rate RemoteBoxID

```

```
1 SEND 12% CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
#
```

NOTES

- The device name on a box to display information can be specified as `-Xaltpath`.
- Settings for dedicated REC buffers must be made in advance.
Change the settings while there are no Consistency mode sessions in either the ACTIVE or ERROR SUSPEND state.
- It takes a short time for changes to buffer settings to take effect, during which time the REC buffer becomes unusable. Do not perform REC operations during this time.
- This command cannot be executed while any of the following commands are running:
 - `swsrpsetvol`
 - `swsrpdelvol`
 - `swsrpstartsync`
 - `swsrpmake`
 - `swsrpstat`
 - `swsrpcancel`
 - `swsrpchsync`
 - `swsrprevsync`
 - `swsrprecoverres`
- The REC buffer settings cannot be changed by using TFOV.

9.3.2 Operation Commands

This section describes operation commands for replication management.

9.3.2.1 swsrpstartsync (Synchronous Processing Start Command)

NAME

`swsrpstartsync` - starts synchronous processing

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsrp/bin/swsrpstartsync [ [-y [-a | -v]] -k | -i] [-g] [-m] [-Xskip | -Xremain] | -Xgds-softcopy | -Xgds-selectcopy] [-t] [-Xda] [-Xgds-slice-skipchk] fromVolumeName toVolumeName
```

To Execute on Management Server

```
/opt/FJSVswsrp/bin/swsrpstartsync [-h Server-Name] [ [-y [-a | -v]] -k | -i] [-g] [-m] [-Xskip | -Xremain] | -Xgds-softcopy | -Xgds-selectcopy] [-t] [-Xda] [-Xgds-slice-skipchk] fromVolumeName toVolumeName
```

To Execute on Managed Server: Group Specification

```
/opt/FJSVswsrp/bin/swsrpstartsync [-y [-a | -v]] -k | -i] [-g] [-m] [-Xskip | -Xremain] [-t] [-Xda] [-Xgds-slice-skipchk] [-Xreverse] -Xgroup groupName
```

To Execute on Management Server: Group Specification

```
/opt/FJSVswsrp/bin/swsrpstartsync [-h serverName] [ [-y [-a | -v]] -k | -i] [-g] [-m] [-Xskip | -Xremain] [-t] [-Xda] [-Xgds-slice-skipchk] [-Xreverse] -Xgroup groupName
```

DESCRIPTION

Start the synchronous replication processing (copying between disks by EC and REC).

It copies 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 just copies updated data to the destination volume. Use this command at the start of operation or when you want to reflect updated data (differences) from the equivalency maintenance state.

For a PRIMECLUSTER GD transaction volume, synchronous processing using a software copy is also possible.

If the Storage Cluster Continuous Copy function is used, the synchronous replication process is started on both the Primary Storage and Secondary Storage.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

When performing the server-to-server replication, specify the name of the operation server configured with the [swsrpsetvol](#) command.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdispsrv](#) command.

-y

Specifies inter-box synchronization in synchronous mode. This option is valid only when total copy is started for inter-box replication.

Specify this option together with the **-a** option to start the split mode for inter-box synchronization in Automatic Split mode.

Specify this option together with the **-v** option to start the mode in Manual Split mode.

Specifying the **-y** option without the **-a** or **-v** option starts the mode in Automatic Split mode.

Not specifying the **-y**, **-k**, and **-i** options sets Through mode.

Synchronization in intra-box replication is always handled in synchronous mode.

-a

Specifies the start of split mode for inter-box synchronization in Automatic Split mode. This option is valid only when total copy is started for inter-box replication.

Specify this option together with the **-y** option.

You cannot use this option when operating SDX objects in units of logical volumes.

-v

Specifies the start of split mode for inter-box synchronization in Manual Split mode. This option is valid only when total copy is started for inter-box replication.

Specify this option together with the **-y** option.

You must specify this option when specifying synchronization mode (the **-y** option) on a box that does not support Automatic Split mode.

You cannot use this option when operating SDX objects in units of logical volumes.

-k

Specifies the setting of transfer mode for inter-box synchronization in Stack mode. This option is valid only when total copy is started for inter-box replication.

Not specifying the **-y**, **-k**, or **-i** option sets Through mode.

You cannot use this option when operating SDX objects in units of logical volumes.

-i

Specifies the setting of transfer mode for inter-box synchronization in Consistency mode. This option is valid only when total copy is started for inter-box replication.

Not specifying the **-y**, **-k**, or **-i** option sets Through mode.

You cannot use this option when operating SDX objects in units of logical volumes.

-g

Specifies the setting of recovery mode for inter-box synchronization in Manual Recovery mode. This option is valid only when total copy is started for inter-box replication.

Not specifying this option sets Automatic Recovery mode.

You cannot use this option when operating SDX objects in units of logical volumes.

-m

For server-to-server replication, this option specifies that communication processing with non-operational servers is not performed. When this option is specified, volume status check processing and the pre-processing for the replication destination volume of a non-operational server are not performed.

This option is valid only during server-to-server replication.

-t

Specifies that pre-processing for the replication destination volume is not performed.

Use this option if any of the following conditions apply:

- AdvancedCopy Manager does not perform pre-processing or post-processing for the replication destination volume, and independent pre-processes and post-processes are inserted before and after AdvancedCopy Manager commands. (*)
- Replication destination pre-processing and post-processing are judged to be unnecessary (for example, for databases constructed on RAW devices). (*)
- Synchronous processing (REC) between boxes is restarted after a forcible suspend.

*: If you specify this option to start or resume synchronous processing, specify the -t option also when executing the [swsrpmake](#) command.

-Xgds-softcopy

Specifies the start of synchronous processing of an SDX object and use of the software-copy function instead of the Advanced Copy function.

If you specify this option to start differential copying when the Advanced Copy function is being used, the Advanced Copy processing stops and software-copy processing starts.

You can specify the option when the function is used with PRIMECLUSTER GD Snapshot is being used.

The option is applicable only to SDX objects.

The option cannot be specified together with -y or -Xgds-selectcopy option.

-Xgds-selectcopy

Specifies the start of synchronous processing without a specified copy type, such as copying with the ETERNUS Disk storage system (OPC/EC) or software-copy. PRIMECLUSTER GD selects the type of copying to be started.

You can specify the option when the function in linkage with PRIMECLUSTER GD Snapshot is being used.

The option is applicable only to SDX objects.

The option cannot be specified together with -Xgds-softcopy option.

-Xgroup *groupName*

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).

This option is valid only when the -Xgroup option is specified.

-Xskip

Specifies that synchronous processing must be suspended using the initial copy skip function. This option is valid only when total copy is started for inter-box replication.

This option is not supported when PRIMECLUSTER GD Snapshot linkage is being used.

-Xremain

Enables data that has been updated to the copy destination during replication established status when synchronous processing is restarted.

This option is valid only when inter-box replication is restarted.

This option is used to reflect to the copy destination volume only those parts of the copy source volume that have been updated, after synchronous processing has been started using the `-Xskip` option and backup data has been copied from the tape media to the copy destination volume.

This option is not supported when PRIMECLUSTER GD Snapshot linkage is being used.

`-Xgds-slice-skipchk`

Skips the SDX object status check processing that is usually performed as part of the copy destination pre-processing if SDX objects are operated in slice units.

If this option is specified for a stopped volume, the `-t` option must also be specified.

`-Xda`

Disables update (WRITE) and access (READ) to the replication destination volume from the synchronous processing start until the replication has been completed.

This option is valid only when synchronous processing is started.

If the option is omitted, only update (WRITE) to the replication destination volume is disabled.

You cannot use this option when operating SDX objects in units of logical volumes.

OPERANDS

fromVolumeName

Specifies a replication source volume name.

Specifies the replication source volume or the replication destination volume set with the `swsrpsetvol` command.

Specifies the [AdvancedCopy Manager device name](#) for the volume name.

toVolumeName

Specifies a replication destination volume name.

Specifies the replication source volume or the replication destination volume set with the `swsrpsetvol` command.

Specifies the [AdvancedCopy Manager device name](#) for the volume name.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Start replication from `/dev/dsk/c1t2d0s3` of a Management Server to `/dev/dsk/c1t2d0s3@TARG-SV` of Managed Server SRC-SV:

```
# /opt/FJSVswsrp/bin/swsrpstartsync -h SRC-SV /dev/dsk/c1t2d0s3 /dev/dsk/c1t2d0s3@TARG-SV
FROM=/dev/dsk/c1t2d0s3@SRC-SV,TO=/dev/dsk/c1t2d0s3@TARG-SV swsrpstartsync completed
#
```

- Start replication from `/dev/dsk/c1t2d0s3` of a Management Server to `/dev/dsk/c1t2d0s3@TARG-SV` of a Managed Server SRC-SV with replication destination volume access disabled.

```
# /opt/FJSVswsrp/bin/swsrpstartsync -h SRC-SV -Xda /dev/dsk/c1t2d0s3 /dev/dsk/c1t2d0s3@TARG-SV
FROM=/dev/dsk/c1t2d0s3@SRC-SV,TO=/dev/dsk/c1t2d0s3@TARG-SV swsrpstartsync completed
#
```

- Start replication for group "GRP1":

```
# /opt/FJSVswsrp/bin/swsrpstartsync -Xgroup GRP1
GROUP=GRP1 swsrpstartsync completed
#
```

NOTES

- Replication cannot be started if:
 - The specified source volume and destination volume have not been set as replication volumes.

- A replication source volume is specified as the copy volume and a replication destination volume is specified as the copied material volume, and one-way copy is set as the copy attribute of the replication volume.
 - Volume information (box identifier, OLU, EXTENT start position, and EXTENT size) is changed after operation starts. In this case, the operation cannot continue for the corresponding volume. Stop the replication operation for the corresponding volume, and then delete and reset the replication volume information.
 - For server-to-server replication, the Managed Server on which the command is executed is not the operation server of the replication volume.
 - For server-to-server replication, communication with a destination server fails.
 - Replication cannot be started as a result of the preceding Advanced Copy function.
 - The -g option is specified, except for a remote copy.
 - When the original volume or replica volume is a logical volume of VxVM, the structure of the logical volume is changed to a structure not supported by AdvancedCopy Manager.
 - When the original volume or replica volume is a logical volume of VxVM, the logical volume is deleted from VxVM.
 - The -Xgds-softcopy or -Xgds-selectcopy option is specified when the function in linkage with PRIMECLUSTER GD Snapshot is not being used.
 - The -a, -v, -k, -i, or -g option is specified when the function in linkage with PRIMECLUSTER GD Snapshot is being used.
 - The backup management function is executing a process for a specified source/destination volume
 - Either the copy source volume or the copy destination volume is a Snap Data Volume.
 - If the -Xgroup option is specified, replication volumes are processed in order, but processing is terminated immediately if an error occurs.
 - When the replication to be performed satisfies either of the following two conditions, perform unmount of the copy target volume before executing a command.
 - When a replication is performed between different OSs
 - When the sizes of replication source volume and replication destination volume differ
 - When the file systems of replication source volume and replication destination volume differ
- In the above mentioned cases, if dismount have been performed, post-processing of the copy target volume at the time of duplicate creation results in an error.
- You cannot change synchronization mode when synchronization is resumed.
When synchronization is resumed, you must specify the -y, -a, -v, -k, -i, and -g options as follows:
 - Specify only the options that match the mode of synchronization to be resumed, or
 - Do not specify any of the options.
 - The processing of this command varies according to the status of the synchronization processing of the specified replication source and replication destination volumes.

Table 9.8 Differences in Processing Corresponding to Synchronization Status

Status of Synchronization Processing	Processing
Dismount status	Start of total copy
Total or incremental copy status	Informational message (swsrp2401) is output and the command ends normally
Equivalency maintain status	Informational message (swsrp2401) is output and the command ends normally
Replication established status (i.e., halt status)	Start of incremental copy

- Before starting the replication, run the replication pre-processing script for the replication destination volume. The contents of this script can be customized. For more information about this, refer to "[Appendix C Pre-processing and Post-processing of Replication](#)". If you do not want to implement the pre-processing script, use the -t option.

If replication is started with the -Xgroup option specified, the replication pre-processing script is not executed for the copy destination volume. Therefore, pre-processing must be executed for all copy destination volumes in the group before this command is executed.

- Before starting replication that targets a Logical Unit (disk), execute pre-processing and post-processing of the Logical Unit (disk), refer to "[Appendix C Pre-processing and Post-processing of Replication](#)" for details.
- To protect processing that accesses a replication destination volume, set the replication destination volume before starting replication so that other processes cannot access it. To do so, execute a dismount command. If you cannot execute dismount resulting from any of the conditions described below, this command fails and terminates.
 - There is a directory in which other volumes are mounted under the mount point.
 - A file in the volume is being used. In addition, if the replication destination volume is being used on a server other than the server on which this command is entered, the replication-operation administrator must disable access to the replication destination volume by cancellation of the share setting, unmounting, or by some other method.
- In the case of a SynfinityFILE file system whose replica volume consists of more than one device (multipartition), replication cannot be performed by executing the unmount/mount instruction inside the pre-processing and post-processing scripts. The unmounting of the file system must be performed manually before replication is started and re-mounted manually after that replication has completed.

- When the SynfinityFILE and AdvancedCopy Manager configurations are as follows:

```
# sfxinfo /dev/dsk/clt3d10s1
ID      special                                size  mount  Type
0       /dev/dsk/clt3d10s1(800019)                25986 /mnt   META
0       /dev/dsk/clt3d10s1(800019)                105751 /mnt   DATA
1       /dev/dsk/clt3d15s3(800028)                 5120 /mnt   LOG
2       /dev/dsk/clt3d16s3(800030)                131736 /mnt   DATA
```

Setting AdvancedCopy Manager

```
# /opt/FJSVswsrp/bin/swsrpvoinfo
Server Original-Volume      Size      Replica-Volume      Size      Copy      Op-Server
SV1    /dev/dsk/clt3d10s1@SV1    128.0 Mbyte /dev/dsk/clt3d20s1@SV1 128.0 Mbyte bi-direction
original
SV1    /dev/dsk/clt3d15s3@SV1    5.0 Mbyte /dev/dsk/clt3d25s3@SV1 5.0 Mbyte bi-direction
original
SV1    /dev/dsk/clt3d16s3@SV1    128.0 Mbyte /dev/dsk/clt3d26s3@SV1 128.0 Mbyte bi-direction
original
```

- Backup procedure (original volume: SynfinityFILE)

```
# /opt/FJSVswsrp/bin/swsrpstartsync /dev/dsk/clt3d10s1 /dev/dsk/clt3d20s1
FROM=/dev/dsk/clt3d10s1@SV1,TO=/dev/dsk/clt3d20s1 swsrpstartsync completed
# /opt/FJSVswsrp/bin/swsrpstartsync /dev/dsk/clt3d15s3 /dev/dsk/clt3d25s3
FROM=/dev/dsk/clt3d15s3@SV1,TO=/dev/dsk/clt3d25s3 swsrpstartsync completed
# /opt/FJSVswsrp/bin/swsrpstartsync /dev/dsk/clt3d16s3 /dev/dsk/clt3d26s3
FROM=/dev/dsk/clt3d16s3@SV1,TO=/dev/dsk/clt3d26s3 swsrpstartsync completed
:
:(Checking equivalence-held state)
:
# umount /mnt
# /opt/FJSVswsrp/bin/swsrpmake /dev/dsk/clt3d10s1 /dev/dsk/clt3d20s1
FROM=/dev/dsk/clt3d10s1@SV1,TO=/dev/dsk/clt3d20s1 swsrpmake completed
# /opt/FJSVswsrp/bin/swsrpmake /dev/dsk/clt3d15s3 /dev/dsk/clt3d25s3
FROM=/dev/dsk/clt3d15s3@SV1,TO=/dev/dsk/clt3d25s3 swsrpmake completed
# /opt/FJSVswsrp/bin/swsrpmake /dev/dsk/clt3d16s3 /dev/dsk/clt3d26s3
FROM=/dev/dsk/clt3d16s3@SV1,TO=/dev/dsk/clt3d26s3 swsrpmake completed
```

```
# mount -F sfxfs /dev/dsk/ctl3d10s1 /mnt
#
```

- Restore procedure (replica volume: SynfinityFILE)

```
# umount /mnt
# /opt/FJSVswsrp/bin/swsrpstartsync /dev/dsk/ctl3d20s1 /dev/dsk/ctl3d10s1
FROM=/dev/dsk/ctl3d20s1@SV1,TO=/dev/dsk/ctl3d10s1 swsrpstartsync completed
# /opt/FJSVswsrp/bin/swsrpstartsync /dev/dsk/ctl3d25s3 /dev/dsk/ctl3d15s3
FROM=/dev/dsk/ctl3d25s3@SV1,TO=/dev/dsk/ctl3d15s3 swsrpstartsync completed
# /opt/FJSVswsrp/bin/swsrpstartsync /dev/dsk/ctl3d26s3 /dev/dsk/ctl3d16s3
FROM=/dev/dsk/ctl3d26s3@SV1,TO=/dev/dsk/ctl3d16s3 swsrpstartsync completed:
:(Checking equivalence-held state)
:
# /opt/FJSVswsrp/bin/swsrpmake /dev/dsk/ctl3d20s1 /dev/dsk/ctl3d10s1
FROM=/dev/dsk/ctl3d20s1@SV1,TO=/dev/dsk/ctl3d10s1 swsrpmake completed
# /opt/FJSVswsrp/bin/swsrpmake /dev/dsk/ctl3d25s3 /dev/dsk/ctl3d15s3
FROM=/dev/dsk/ctl3d25s3@SV1,TO=/dev/dsk/ctl3d15s3 swsrpmake completed
# /opt/FJSVswsrp/bin/swsrpmake /dev/dsk/ctl3d26s3 /dev/dsk/ctl3d16s3
FROM=/dev/dsk/ctl3d26s3@SV1,TO=/dev/dsk/ctl3d16s3 swsrpmake completed
# mount -F sfxfs /dev/dsk/ctl3d10s1 /mnt
#
```

- When SDX objects of PRIMECLUSTER GD are used: Refer to "[Configuration and Conditions of Unsupported SDX Objects](#)" and "[10.1.5.2 Notes on SDX Object Operations \(Replication Management\)](#)".
- Refer to "[10.1.1 General Notes](#)" for notes about starting replication.
- An error occurs if the synchronous process is restarted with -Xda option. When changing the access permission for replication destination volume, first cancel the current replication processing with the [swsrpcancel](#) command and then restart this command.
- This command terminates abnormally with the following error message if the -Xda option is set for an ETERNUS Disk storage system that does not support the Destination Access Permission function.

```
swsrp0719 EC cannot be run. Error=17(not support.)
```

- This command cannot be executed while any of the following commands are running:
 - [swsrpsetvol](#)
 - [swsrpdelvol](#)
 - [swsrprecoverres](#)
- When the Storage Cluster Continuous Copy function is used:
 - When the status of the TFO group is "Normal", and synchronous processing cannot start on either the Primary Storage or Secondary Storage, the command may terminate abnormally. Check the cause of the error on both the Primary Storage and Secondary Storage to address the error message in the storage device where the error is occurring.
 - When the status of the TFO group is "Normal" and its phase is "Maintenance", or when the status of the TFO group is "Halt", the command starts synchronous processing on the storage that has an "Active" TFO status and terminates normally. However, synchronous processing could not be started on the storage that has a "Standby" TFO status. In this case, swsrp2882 message may be output.
If a swsrp2882 message is output, after the status of TFO group has become "Normal" and the phase is not "Maintenance", refer to "[7.10 When Recovering Storage Cluster Continuous Copy Function](#)" to take action.

9.3.2.2 swsrpmake (Replication Creation Command)

NAME

swsrpmake - creates a replication volume

SYNOPSIS

To Execute on Managed Server

Synchronized Replication

```
/opt/FJSSVswsrp/bin/swsrpmake [-m] [-f] [-t] ([-j | -Xconcur]) [-Xgds-slice-skipchk]
fromVolumeName toVolumeName
```

Snapshot Replication

```
/opt/FJSSVswsrp/bin/swsrpmake [-m] [-f] [-t] ([-T | -C | -P]) [-Xgds-slice-skipchk]
fromVolumeName toVolumeName
```

To Execute on Management Server

Synchronized Replication

```
/opt/FJSSVswsrp/bin/swsrpmake [-h serverName] [-m] [-f] [-t] ([-j | -Xconcur]) [-Xgds-slice-
skipchk] fromVolumeName toVolumeName
```

Snapshot Replication

```
/opt/FJSSVswsrp/bin/swsrpmake [-h serverName] [-m] [-f] [-t] ([-T | -C | -P]) [-Xgds-slice-
skipchk] fromVolumeName toVolumeName
```

To Execute on Managed Server: Group Specification

Synchronized Replication

```
/opt/FJSSVswsrp/bin/swsrpmake [-m] [-f] [-t] ([-j | -Xconcur]) [-Xreverse] -Xgroup groupName
```

Snapshot Replication

```
/opt/FJSSVswsrp/bin/swsrpmake [-m] [-f] [-t] ([-T | -C | -P]) [-Xconcur-opc] [-Xreverse] -Xgroup
groupName
```

To Execute on Management Server: Group Specification

Synchronized Replication

```
/opt/FJSSVswsrp/bin/swsrpmake [-h serverName] [-m] [-f] [-t] ([-j | -Xconcur]) [-Xreverse] -
Xgroup groupName
```

Snapshot Replication

```
/opt/FJSSVswsrp/bin/swsrpmake [-h serverName] [-m] [-f] [-t] ([-T | -C | -P]) [-Xconcur-opc] [-
Xreverse] -Xgroup groupName
```

DESCRIPTION

When synchronization processing (EC) is not performed, this command starts snapshot processing (OPC/QuickOPC/SnapOPC/SnapOPC+) to create a replication volume.

When synchronization processing is performed, check whether the status of the processing is in the equivalency maintenance state. If so, suspend the synchronization processing and create a replication volume. The status in which synchronization processing is suspended is referred to as the replication established status. After this command is executed, the replication destination volume can be accessed.

The behavior of this command could differ depending on the status of the Advanced Copy during progress and the following should be noted:

- When synchronous processing is not performed, the command starts snapshot processing and creates replication volumes. This form of replication is called snapshot replication.
- When synchronous processing is performed, the command checks the state of synchronous processing and, if the synchronous processing is in the equivalency maintenance state, the command suspends the processing in order to create a replication volume. This form of replication is called synchronous replication. The state in which synchronous processing is suspended is called the copy established state.

In either case, the copy destination volume is made accessible after this command is executed.

This command executes the replication pre-processing and post-processing before a replication is created. For details of these processes, refer to "[Appendix C Pre-processing and Post-processing of Replication](#)".

If the Storage Cluster Continuous Copy function is used, replication volumes are created on both the Primary Storage and Secondary Storage.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

When performing the server-to-server replication, specify the name of the operation server configured with the [swsrpsetvol](#) command.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdisprv](#) command.

-m

For server-to-server replication, this specifies that communication processing is not performed with the non-operation server. When this option is specified, volume status check processing and pre-processing and post-processing are not performed for the non-operation server volume.

This option is valid only during server-to-server replication.

-f

Specifies that the pre-processing and post-processing are not performed for a replication source volume.

Use this option only when AdvancedCopy Manager does not execute pre-processing and post-processing of the copy source volume and unique pre-processing and post-processing scripts are inserted before and after (respectively) the AdvancedCopy Manager command, or when you judge that copy pre-processing and post-processing are unnecessary (e.g., for a database constructed on a raw device).

-t

Specifies that the pre-processing and post-processing are not performed for a replication destination volume.

Use this option when applied to either of the following conditions:

- Pre-processing and post-processing for copy destination volume is not performed by AdvancedCopy Manager but independent pre-processing and post-processing is added in before and after AdvancedCopy Manager commands.
- Pre-processing and post-processing for copy destination volume is determined to be unnecessary (e.g. Database created on RAW device)

If you specify the -t option to execute the [swsrpstartsync](#) command, specify the -t option also when executing this command.

-j

Specifies that copy processing is forcibly suspended even if the execution status of copying is "sync: Total copy or incremental copy is being performed" or "halt: Hard suspend status" in synchronization processing of intra-box replication. This option can only be used for intra-box replication; if attempted for intra-box replication, an error occurs.

When this option is specified, the data in the destination volume is not guaranteed.

Additionally, when this option is specified, pre-processing and post-processing for the source and destination volumes are not performed. When restarting synchronous processing (REC) between boxes, execute the [swsrpstartsync](#) command with the -t option specified.

-T

Specifies that differential snapshot processing is executed.

This option cannot be specified while synchronous processing is performed.

This option is valid only for a replication within the ETERNUS Disk storage system and when the ETERNUS Disk storage system supports the QuickOPC function.

For a snapshot replication without specifying this option, ordinary snapshot processing (OPC without using the QuickOPC function) is started. This option must be specified to perform the replication operation using differential snapshots.

-Xgroup *groupName*

Specifies that a replication is created for each group.

-Xreverse

Specifies that a replication is created for each group in the reverse direction (i.e., from replica volumes to original volumes). This option is valid only when the **-Xgroup** option is specified.

-Xconcur

Specifies that a replication is created using the Concurrent Suspend function.

This option is not supported when PRIMECLUSTER GD Snapshot linkage is being used.

-Xconcur-opc

Specifies that a replication is created using the Concurrent OPC function.

This option is not supported when linking with PRIMECLUSTER GD Snapshot.

-C

Specifies to start SnapOPC.

This option is only valid for intra-box replications and for ETERNUS Disk storage system that supports the SnapOPC function.

-P

Specifies to start SnapOPC+.

This option is only valid with intra-box replications and for ETERNUS Disk storage system that supports the SnapOPC+ function.

An error occurs in the following cases:

- The same SnapOPC+ session exists at the copy source volume and the copy destination volume.

-Xgds-slice-skipchk

Skips the SDX object status check processing that is usually performed as part of the copy source and copy destination pre-processing if SDX objects are operated in slice units.

If this option is specified for a stopped volume, the **-f** and **-t** options must also be specified.

OPERANDS

fromVolumeName

Specifies a replication source volume.

Specifies the replication source volume or the replication destination volume that was set with the [swsrpsetvol](#) command.

Specifies the [AdvancedCopy Manager device name](#) for the volume name.

toVolumeName

Specifies a replication destination volume.

Specifies the replication source volume or the replication destination volume that was set with the [swsrpsetvol](#) command.

Specifies the [AdvancedCopy Manager device name](#) for the volume name.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- The Management Server instructs Managed Server: SRC-SV to create a copy of /dev/dsk/c1t2d0s3 in /dev/dsk/c1t2d0s3@TARG-SV:

```
# /opt/FJSVswsrp/bin/swsrpmake -h SRC-SV /dev/dsk/c1t0d2s3 /dev/dsk/c1t2d0s3@TARG-SV
FROM=/dev/dsk/c1t2d0s3@SRC-SV,TO=/dev/dsk/c1t2d0s3@TARG-SV swsrpmake completed
#
```


- Creates a copy for the group "GRP1":

```
# /opt/FJSVswsrp/bin/swsrpmake -Xgroup GRP1
GROUP=GRP1 swsrpmake completed
#
```

- Creates a copy for the group "GRP1" by using the Concurrent OPC:

```
# /opt/FJSVswsrp/bin/swsrpmake -Xgroup GRP1 -Xconcur-opc
GROUP=GRP1 swsrpmake completed
#
```

NOTES

- Replication cannot be created if:
 - The specified replication source and replication destination volumes are not set as replication volumes.
 - The synchronization processing from the specified replication source volume to the replication destination volume is not in the equivalency maintenance status.
 - After operation starts, physical volume information (box identifier, OLU, EXTENT start position, and EXTENT size) is changed. In this case, operation for the corresponding volume cannot continue. Stop the replication operation for the volume, and then delete and reset the replication volume information.
 - If a replication source volume is an SDX object (SynfinityDISK), the status of the SDX object does not match the following:
 - The mirror volume is ACTIVE or STOP.
 - The mirror slice status is ACTIVE or TEMP
When the mirror slice is TEMP, a reproduction former volume is not pre-processed.
 - The SDX disk status is ENABLE.
 - For server-to-server replication, the Managed Server on which the command is executed is not the operation server of the replication volume.
 - For server-to-server replication, communication with a destination server fails.
 - When the copy source volume or destination volume constitutes a SynfinityFILE of multi-volume composition and the file system is mounted. In this case, unmount all the volumes.
 - When the original volume or replica volume is a logical volume of VxVM and the structure of the logical volume is changed to a structure not supported by AdvancedCopy Manager or deleted from VxVM.
 - If the -j, -T, or -C option is specified when the function is used with PRIMECLUSTER GD Snapshot.
 - The backup management function is executing a process for a specified source or destination volume
 - From the copy destination server, the command was executed with both -Xconcur and -m options specified simultaneously for any of the following ETERNUS Disk storage systems:
 - ETERNUS DX80 S2 (for earlier than V10L30 firmware)
 - ETERNUS DX90 S2 (for earlier than V10L30 firmware)
 - ETERNUS DX400/DX400 S2 series (for earlier than V10L30 firmware)
 - ETERNUS DX8000/DX8000 S2 series (for earlier than V10L30 firmware)
 - If the -Xgroup option is specified, processing is 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 is terminated immediately if an error occurs.
 - Note the following when you specify the -Xconcur-opc option:
 - Specify this option together with the -Xgroup option.

- If an error is detected by the preliminary check for the replication volumes to be processed, processing is terminated without creating a copy.
Refer to the following table for details of the preliminary checks.

Type	Details of Preliminary Checks
Common	Check if a pair of synchronized replications is not exist in the group.
When the - T option is specified	Check if a pair of volumes which is in the tracking state and a pair of volumes on which the tracking is not executed are not mixed in the group.

- If this command is executed during the snapshot processing, start or restart the snapshot processing. Refer to the following table for details of copy operation.

OPC/SnapOPC/SnapOPC+

Operation Status (Display by swsrpstat -L)	Copy Operation
<p>When the operation status of the replication volume in the group is in one of the following:</p> <ul style="list-style-type: none"> - The Concurrent OPC is defined. (Status is "readying".) - Copy is not executed. (Status is "----" and Trk is "----".) - Snapshot processing is in progress and tracking is not executed. (Status is "snap" and Trk is "off".) - Snapshot processing is in progress and tracking is in progress. (Status is "snap" and Trk is "on".) - Tracking is in progress. (Status is "----" and Trk is "on".) - SnapOPC is in progress. (Status is "copy-on-write".) - SnapOPC+ is in progress. (Status is "copy-on-write(active)" or "copy-on-write(inactive)".) - Error suspend status (Status is "failed".) 	<ol style="list-style-type: none"> 1. Copying that is in progress is terminated. 2. Define the Concurrent OPC. Snapshot processing is newly started with the Concurrent OPC specified.

QuickOPC

Operation Status (Display by swsrpstat -L)	Copy Operation
<p>When the operation status of the replication volume in the group is in one of the following:</p> <ul style="list-style-type: none"> - The Concurrent OPC is defined. (Status is "readying".) - Copy is not executed. (Status is "----" and Trk is "----".) - Snapshot processing is in progress and tracking is not executed. (Status is "snap" and Trk is "off".) 	<ol style="list-style-type: none"> 1. Copying that is in progress is terminated. 2. Define the Concurrent OPC. 3. Differential snapshot processing is newly started with the Concurrent OPC specified.

Operation Status (Display by swsrpstat -L)	Copy Operation
<ul style="list-style-type: none"> - SnapOPC is processing (Status is "copy-on-write") - SnapOPC+ is processing (Status is "copy-on-write(active)" or "copy-on-write(inactive)") - Error suspend status (Status is "failed".) 	
<p>When the operation status of the replication volume in the group is in one of the following:</p> <ul style="list-style-type: none"> - Snapshot processing is in progress and tracking is in progress. (Status is "snap" and Trk is "on".) - Tracking is in progress. (Status is "----" and Trk is "on".) 	<p>Differential snapshot processing is restarted with the Concurrent OPC specified.</p>

- When the replication to be performed satisfies either of the following conditions, perform dismount of the copy target volume before executing a command.
 - When a replication is performed between different OSs
 - When the sizes of replication source volume and replication destination volume differ
 - When the file systems of replication source volume and replication destination volume differ

If dismount is performed in the abovementioned case, post-processing of the copy target volume at the time of duplicate creation produces an error.

- Before a replica is created, the pre-processing for the source volume and destination volume are executed (for synchronized replication, only pre-processing is executed). After the replica has been created, the post-processing for the source volume and destination volume is executed. The contents of the script can be customized. For more information about this refer to "[Appendix C Pre-processing and Post-processing of Replication](#)". If the -f or -t option is specified, the pre-processing and post-processing are not performed.

If a replica is created with the -Xgroup option specified, the replication, pre-processing and post-processing scripts are not executed for the copy source volume/copy destination volume. Therefore, execute pre-processing for all copy source volumes/copy destination volumes in the group before this command is executed.

(However, for synchronous replication, pre-processing for the copy destination volume is executed when replication starts, so it need not be executed.)

In addition, post-processing must be executed after this command is executed.

- If this command is executed during execution of snapshot processing, stop the copy processing in progress and restart the snapshot processing.
- To protect the data, this command sets the replication source volume before creating a replica so that other processes cannot access the volume. The replication source volume is unmounted in order to disable to access it. If you cannot unmount the source volume because of the situations described below, this command fails and terminates.
 - There is a directory in which other volumes are mounted under the mount point.
 - A file on the volume is being used. In addition, if the replication source volume is being used from a server other than the server on which this command is entered, the backup-operation administrator must disable access to the replication source volume by cancellation of the share setting, unmounting, or by using some other method. If for whatever reason you do not want to remove the share setting of the replication source volume, or execute unmount or use some other method, refer to "[Appendix C Pre-processing and Post-processing of Replication](#)" and note the following:
 - For replication pre-processing, execute the UNIX sync command to synchronize volumes and the file system.
 - In replication post-processing, execute the <fsck> command to check the file system at the copy destination.

- When the replication source volume is being used by backup management, dismount may not be possible. In this case, "swsrp2613 An error occurred in the pre-processing script for creating a replica. Error Code = 2" is output.
- When you are using the original copy source volume in the SynfinityFILE file system, use the following procedures (UNIX commands) to mount the copy source volume.

```
# sfxadm <The list of RAW device names which constitute SynfinityFILE>
# fsck -F sfxfs [-y] <The representation RAW device name of SynfinityFILE>
# mount -F sfxfs <The representation device name of SynfinityFILE Mount point>
```

- Where a SynfinityFILE file system whose original or replica volume consists of more than one device (multipartition), replication cannot be performed by executing the unmount/mount instruction inside the pre-processing and post-processing scripts. The unmounting of the file system must be performed manually before replication is started and re-mounted manually after that replication has completed.

- When the SynfinityFILE and AdvancedCopy Manager configurations are as follows:

```
# sfxinfo /dev/dsk/clt3d10s1
ID      special                                size    mount    Type
0       /dev/dsk/clt3d10s1(800019)              25986   /mnt     META
0       /dev/dsk/clt3d10s1(800019)              105751  /mnt     DATA
1       /dev/dsk/clt3d15s3(800028)              5120    /mnt     LOG
2       /dev/dsk/clt3d16s3(800030)              131736  /mnt     DATA

Setting AdvancedCopy Manager
# /opt/FJSVswsrp/bin/swsrpvoinfo
Server Original-Volume      Size      Replica-Volume      Size      Copy      Op-Server
SV1 /dev/dsk/clt3d10s1@SV1    128.0 Mbyte /dev/dsk/clt3d20s1@SV1 128.0 Mbyte bi-direction
original
SV1 /dev/dsk/clt3d15s3@SV1     5.0 Mbyte /dev/dsk/clt3d25s3@SV1  5.0 Mbyte bi-direction
original
SV1 /dev/dsk/clt3d16s3@SV1    128.0 Mbyte /dev/dsk/clt3d26s3@SV1 128.0 Mbyte bi-direction
original
```

- Backup procedure (original volume: SynfinityFILE (snapshot))

```
# umount /mnt
# /opt/FJSVswsrp/bin/swsrpmake /dev/dsk/clt3d10s1 /dev/dsk/clt3d20s1
FROM=/dev/dsk/clt3d10s1@SV1,TO=/dev/dsk/clt3d20s1 swsrpmake completed
# /opt/FJSVswsrp/bin/swsrpmake /dev/dsk/clt3d15s3 /dev/dsk/clt3d25s3
FROM=/dev/dsk/clt3d15s3@SV1,TO=/dev/dsk/clt3d25s3 swsrpmake completed
# /opt/FJSVswsrp/bin/swsrpmake /dev/dsk/clt3d16s3 /dev/dsk/clt3d26s3
FROM=/dev/dsk/clt3d16s3@SV1,TO=/dev/dsk/clt3d26s3 swsrpmake completed
# mount -F sfxfs /dev/dsk/clt3d10s1 /mnt
```

- Backup procedure (original volume: SynfinityFILE (synchronous))

```
# /opt/FJSVswsrp/bin/swsrpstartsync /dev/dsk/clt3d10s1 /dev/dsk/clt3d20s1
FROM=/dev/dsk/clt3d10s1@SV1,TO=/dev/dsk/clt3d20s1 swsrpstartsync completed
# /opt/FJSVswsrp/bin/swsrpstartsync /dev/dsk/clt3d15s3 /dev/dsk/clt3d25s3
FROM=/dev/dsk/clt3d15s3@SV1,TO=/dev/dsk/clt3d25s3 swsrpstartsync completed
# /opt/FJSVswsrp/bin/swsrpstartsync /dev/dsk/clt3d16s3 /dev/dsk/clt3d26s3
FROM=/dev/dsk/clt3d16s3@SV1,TO=/dev/dsk/clt3d26s3 swsrpstartsync completed
:
:(Checking equivalence-held state)
:
# umount /mnt
# /opt/FJSVswsrp/bin/swsrpmake /dev/dsk/clt3d10s1 /dev/dsk/clt3d20s1
FROM=/dev/dsk/clt3d10s1@SV1,TO=/dev/dsk/clt3d20s1 swsrpmake completed
# /opt/FJSVswsrp/bin/swsrpmake /dev/dsk/clt3d15s3 /dev/dsk/clt3d25s3
FROM=/dev/dsk/clt3d15s3@SV1,TO=/dev/dsk/clt3d25s3 swsrpmake completed
# /opt/FJSVswsrp/bin/swsrpmake /dev/dsk/clt3d16s3 /dev/dsk/clt3d26s3
FROM=/dev/dsk/clt3d16s3@SV1,TO=/dev/dsk/clt3d26s3 swsrpmake completed
```

```
# mount -F sfxfs /dev/dsk/clt3d10s1 /mnt
#
```

- Restore procedure (original volume: SynfinityFILE (snapshot))

```
# umount /mnt
# /opt/FJSVswsrp/bin/swsrpmake /dev/dsk/clt3d20s1 /dev/dsk/clt3d10s1
FROM=/dev/dsk/clt3d20s1@SV1,TO=/dev/dsk/clt3d10s1 swsrpmake completed
# /opt/FJSVswsrp/bin/swsrpmake /dev/dsk/clt3d25s3 /dev/dsk/clt3d15s3
FROM=/dev/dsk/clt3d25s3@SV1,TO=/dev/dsk/clt3d15s3 swsrpmake completed
# /opt/FJSVswsrp/bin/swsrpmake /dev/dsk/clt3d26s3 /dev/dsk/clt3d16s3
FROM=/dev/dsk/clt3d26s3@SV1,TO=/dev/dsk/clt3d16s3 swsrpmake completed
# mount -F sfxfs /dev/dsk/clt3d10s1 /mnt
#
```

- Restore procedure (original volume: SynfinityFILE (synchronous))

```
# umount /mnt
# /opt/FJSVswsrp/bin/swsrpstartsync /dev/dsk/clt3d20s1 /dev/dsk/clt3d10s1
FROM=/dev/dsk/clt3d20s1@SV1,TO=/dev/dsk/clt3d10s1 swsrpstartsync completed
# /opt/FJSVswsrp/bin/swsrpstartsync /dev/dsk/clt3d25s3 /dev/dsk/clt3d15s3
FROM=/dev/dsk/clt3d25s3@SV1,TO=/dev/dsk/clt3d15s3 swsrpstartsync completed
# /opt/FJSVswsrp/bin/swsrpstartsync /dev/dsk/clt3d26s3 /dev/dsk/clt3d16s3
FROM=/dev/dsk/clt3d26s3@SV1,TO=/dev/dsk/clt3d16s3 swsrpstartsync completed
:
:(Checking equivalence-held state)
:
# /opt/FJSVswsrp/bin/swsrpmake /dev/dsk/clt3d20s1 /dev/dsk/clt3d10s1
FROM=/dev/dsk/clt3d20s1@SV1,TO=/dev/dsk/clt3d10s1 swsrpmake completed
# /opt/FJSVswsrp/bin/swsrpmake /dev/dsk/clt3d25s3 /dev/dsk/clt3d15s3
FROM=/dev/dsk/clt3d25s3@SV1,TO=/dev/dsk/clt3d15s3 swsrpmake completed
# /opt/FJSVswsrp/bin/swsrpmake /dev/dsk/clt3d26s3
/dev/dsk/clt3d16s3FROM=/dev/dsk/clt3d26s3@SV1,TO=/dev/dsk/clt3d16s3 swsrpmake completed
# mount -F sfxfs /dev/dsk/clt3d10s1 /mnt
#
```

- For Snap Data Volume, only SnapOPC/SnapOPC+ can be executed.
- When restoring data from the copy destination volume of SnapOPC/SnapOPC+, note that there must be a SnapOPC session between the copy source volume and the copy destination volume.
- If this command is executed during snapshot processing, snapshot processing is restarted. Refer to the following table for details of copy operation:

Table 9.9 Copy Operation When Command Is Executed During Snapshot Processing

Operation Status (Display by swsrpstat -L)	-T Option Specification	Copy Operation
Copy is not executed. (Status is "----" and Trk also is "----".)	No	Ordinary snapshot processing is started.
Copy is not executed. (Status is "----" and Trk also is "----".)	Yes	Differential snapshot processing is started.
Snapshot processing is in progress and tracking is not executed. (Status is "snap" and Trk is "off".)	No	Copying that is in progress is terminated and ordinary snapshot processing is restarted.
Snapshot processing is in progress and tracking is not executed. (Status is "snap" and Trk is "off".)	Yes	Copying in progress is terminated and differential snapshot processing is restarted.
Snapshot processing is in progress and tracking is in progress. (Status is "snap" and Trk is "on".)	No	Copying in progress is terminated and ordinary snapshot processing is restarted.

Operation Status (Display by swsrpstat -L)	-T Option Specification	Copy Operation
Snapshot processing is in progress and tracking is in progress. (Status is "snap" and Trk is "on".)	Yes	Differential snapshot processing is restarted.
Tracking is in progress. (Status is "----" and Trk is "on".)	No	Tracking processing is terminated and ordinary snapshot processing is restarted.
Tracking is in progress. (Status is "----" and Trk is "on".)	Yes	Differential snapshot processing is restarted.

- When SDX objects of PRIMECLUSTER GD are used: Refer to "[Configuration and Conditions of Unsupported SDX Objects](#)", and "[10.1.5.2 Notes on SDX Object Operations \(Replication Management\)](#)".
- Refer to "[10.1.1 General Notes](#)" for notes on creating replication.
- This command cannot be executed while any of the following commands are running:
 - [swsrpsetvol](#)
 - [swsrpdelvol](#)
 - [swsrprecoverres](#)
- When the Storage Cluster Continuous Copy function is used:
 - If the status of the TFO group is "Normal", and replication volume cannot create on either the Primary Storage or Secondary Storage, the command may terminate abnormally. Check the cause of the error on both the Primary Storage and Secondary Storage to address the error message in the storage device where the error is occurring.
 - When the status of the TFO group is "Normal" and its phase is "Maintenance", or when the status of the TFO group is "Halt", the command starts synchronous processing on the storage that has an "Active" TFO status and terminates normally. However, synchronous processing could not be started on the storage that has a "Standby" TFO status. In this case, swsrp2882 message may be output.
If a swsrp2882 message is output, after the status of TFO group has become "Normal" and the phase is not "Maintenance", refer to "[7.10 When Recovering Storage Cluster Continuous Copy Function](#)" to take action.

9.3.2.3 swsrpstat (Operation Status Display Command)

NAME

swsrpstat - displays the operation status

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsrp/bin/swsrpstat [ [-L] [-Xdate] [-Xda] [-Xstandby-s] | [-Xall] ] [-E] [-H] [-O]
[originalVolumeName]
```

To Execute on Management Server

```
/opt/FJSVswsrp/bin/swsrpstat [-h serverName] [ [-L] [-Xdate] [-Xda] [-Xstandby-s] | [-Xall] ] [-E]
[-H] [-O] [originalVolumeName]
```

To Execute on Managed Server: Group Specification

```
/opt/FJSVswsrp/bin/swsrpstat [ [-L] [-Xdate] [-Xda] [-Xstandby-s] | [-Xall] ] [-E] [-H] [-O] -
Xgroup groupName
```

o Execute on Management Server: Group Specification

```
/opt/FJSVswsrp/bin/swsrpstat [-h serverName] [ [-L] [-Xdate] [-Xda] [-Xstandby-s] | [-Xall] ] [-E]
[-H] [-O] -Xgroup groupName
```

DESCRIPTION

This command displays the operation status of the specified volume. If a volume is not specified, the operation status of each replication volume is displayed.

The following information is displayed.

Title	Description
Server	Displays a Managed Server name.
Original-Volume	Displays a replication source volume name.
Replica-Volume	Displays a replication destination volume name.
Direction	<p>Displays the copy direction if a copy session exists.</p> <p>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.</p>
Status	<p>Displays the execution status.</p> <p>If the Storage Cluster Continuous Copy function is used, the execution status of the ETERNUS Disk storage system where the TFO status is "Active" is displayed.</p> <p>---- : Copying is not being performed. readying : The Concurrent OPC has been defined. sync : Total copy or incremental copy is being performed. equivalent : Equivalency maintained status. replicated : Replication established status. copy-on-write : SnapOPC processing is being performed. copy-on-write(active) : SnapOPC+ processing is being performed (most recent snap generation). copy-on-write(inactive) : SnapOPC+ processing is being performed (past snap generation). failed : Error suspend status. failed(badsector) : There is a bad sector in the copy source. failed(overflow) : There is insufficient Snap Data Volume or Snap Data Pool capacity. failed(other) : The error suspend status is neither failed(badsector) or failed(overflow). halt : Hard suspend status. halt(sync) : Hard suspend status (hard suspend status has been entered during a REC total copy or differential copy). halt(equivalent) : Hard suspend status (hard suspend status has been entered during an REC equivalency maintenance state). halt(use-buffer) : A line fault occurred and the REC buffer is in use. (*1) halt(use-disk-buffer) : A line fault occurred and the REC Disk buffer is in use. (*1) snap : OPC or QuickOPC processing is being performed. ????? : Failed to get copy status. gds-error : Copy using PRIMECLUSTER GD terminated abnormally.</p> <p>*1: "halt(use-buffer)" and "halt(use-disk-buffer)" are output when the following two conditions are fulfilled:</p> <ul style="list-style-type: none"> - The -H option was specified, and - In inter-server replication, this command was executed on the replication source server of REC.
Execute	<p>Displays the copy progress rate as a percentage. "----" is displayed during SnapOPC or when no copying is being performed.</p> <p>If the Storage Cluster Continuous Copy function is used, the replication progress rate of the ETERNUS Disk storage system where the TFO status is "Active" is displayed.</p>
Trk	<p>Displays whether tracking processing is in progress.</p> <p>If the Storage Cluster Continuous Copy function is used, the presence of the tracking process of the ETERNUS Disk storage system where the TFO status is "Active" is displayed.</p>

Title	Description
	<p>on : Tracking processing is in progress. off : Tracking processing is not in progress. ---- : Snapshot processing is not being performed.</p> <p>Always displays "----" when the PRIMECLUSTER GD Snapshot linkage function is enabled.</p>
Update	<p>Displays the percentage of data that has been updated since the last copy was created, in the cases shown below.</p> <p>If the Storage Cluster Continuous Copy function is used, the size of the data that has been updated of the ETERNUS Disk storage system where the TFO status is "Active" is displayed.</p> <p>For SnapOPC+, this displays the percentage of data that has been updated since the previous snap generation.</p> <ul style="list-style-type: none"> - When the Status column is "----" and the Trk column is "on" (that is, tracking is in progress and the physical copy is not being executed). - When the Status column is "copy-on-write" (that is, when SnapOPC processing is in progress) - When the Status column is "copy-on-write(active)" (that is, when SnapOPC+ (most recent snap generation) is being performed) - When the Status column is "copy-on-write(inactive)" (that is, when SnapOPC+ (past snap generation) is being performed) <p>Displays "----" during physical copying or when tracking is not in progress.</p> <p>Always displays "----" when the PRIMECLUSTER GD Snapshot linkage function is enabled.</p>
Rcv	<p>Indicates the Recovery mode for Inter-ETERNUS synchronization.</p> <p>auto : Automatic Recovery mode manual : Manual Recovery mode ---- : Inter-ETERNUS synchronous replication (REC) is not performed.</p>
Split	<p>Indicates the Split mode for Inter-ETERNUS synchronization.</p> <p>auto : Automatic Split mode manual : Manual Split mode ---- : Inter-ETERNUS synchronous replication (REC) is not performed using the Synchronous mode.</p>
Xfer	<p>Indicates the Transfer for Inter-ETERNUS synchronization.</p> <ul style="list-style-type: none"> - sync : Synchronous mode async : Through mode stack : Stack mode consist : Consistency mode <p>---- : Inter-ETERNUS synchronous replication (REC) is not performed.</p>
Snap-Gen	<p>When SnapOPC+ is performed, displays the snap generation number. When SnapOPC+ is not performed, "----" is displayed. 1 is the oldest snap generation, and the snap generation number increases until the most recent generation is reached.</p>
Copy-DA (Copy Destination Access)	<p>Displays the access permission for the replication destination volume.</p> <p>If the Storage Cluster Continuous Copy function is used, the access permission of the ETERNUS Disk storage system where the TFO status is "Active" is displayed.</p> <p>on : Access is not being permitted. off : Access is being permitted.</p> <p>When an ETERNUS Disk storage system does not support the Destination Access Permission function or the SDX objects is specified with the units of logical volumes, 'off' is displayed.</p>

Title	Description
Date	<p>If a copy session exists, a timestamp in the [yyyy/mm/dd hh:mm:ss] format is displayed for the last status transition of the copy session as described below.</p> <p>If the Storage Cluster Continuous Copy function is used, the timestamp of the ETERNUS Disk storage system where the TFO status is "Active" is displayed.</p> <p>The timestamp is relative to the time zone settings of the specified storage system.</p> <p>[For the OPC and QuickOPC]</p> <ul style="list-style-type: none"> - Starting the logical copy (Execution the "swsrpmake" or "swsrpmake -T") - Changing to error suspend status (The Status is "failed") <p>[For the SnapOPC and SnapOPC+]</p> <ul style="list-style-type: none"> - Starting the logical copy (Execution the "swsrpmake -C" or "swsrpmake -P") - Changing to error suspend status (The Status is "failed") <p>[For the EC]</p> <ul style="list-style-type: none"> - Suspend (Execution the "swsrpmake") - Concurrent Suspend (Execution the "swsrpmake -Xconcur") - Changing to error suspend status (The Status is "failed") <p>[For the REC]</p> <ul style="list-style-type: none"> - Starting the REC by initial copy skip function (Execution the "swsrpstartsync -Xskip") - Suspend (Execution the "swsrpmake") (*1) - Concurrent Suspend (Execution the "swsrpmake -Xconcur") - Changing to error suspend status (The Status is "failed") - Changing to hard suspend status (The Status is "halt") <p>*1: When a forced suspend is performed on a REC session in hardware suspend condition, the timestamp of the last execution is initialized as "----/--/-- --:--:--". When a forced suspend is performed while copy is executing, its timestamp is not updated.</p> <p>"----/--/-- --:--:--" is displayed in the following situations:</p> <ul style="list-style-type: none"> - There is no copy session (Status column is "----") - Copy status could not be obtained from the storage system (Status column is "?????") - Transition between no copy session and start of a synchronous EC/REC initial copy (swsrpstartsync) is started under the condition of no copy session. - The REC session is recovered automatically from the hardware suspend status (The Status is "halt") - The storage system where last execution timestamp function is not supported - The SDX objects being set per logical volume
Standby-Session	<p>If the Storage Cluster Continuous Copy function is used, the Standby-Side Continuous Copy Session status is displayed.</p> <ul style="list-style-type: none"> - exist: A Standby-Side Continuous Copy Session exists. - not-exist: A Standby-Side Continuous Copy Session does not exist. (An Active-Side Continuous Copy Session exists)

Title	Description
	<ul style="list-style-type: none"> - ----: A Standby-Side Continuous Copy Session does not exist. (An Active-Side Continuous Copy Session also does not exist) - ?????: The Standby-Side Continuous Copy Session status is unknown.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdisprv](#) command.

-L

Specifies to display in extension format.

For the items displayed by this option, refer to the table in the description above.

-E

An option which when set displays 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 is always displayed in the normal direction (from the copy source to the copy destination) when a bi-directional OPC physical copy is taking place.

In case this option is set, the operational status in the normal direction is not displayed. Only the operational status in the opposite direction is displayed.

This option can be used to check the operational status of an OPC session performed in the opposite direction if a QuickOPC/SnapOPC session in the normal direction exists.

-Xgroup *groupName*

Specifies that the operation statuses are displayed by the group.

If this option is set, all the operational statuses relating to the specified group are displayed.

-H

Changes the Status column display method used when a REC enters a hard-suspend (halt) status.

If this option is specified, then "halt (sync)" is displayed when hard-suspend status occurs during a total copy or a differential copy.

"halt (equivalent)" is displayed when hard-suspend status occurs during an equivalency maintenance state.

If the transfer mode is Consistency mode and a REC Disk buffer is set, "halt(use-disk-buffer)" is displayed when issue of a hard suspend causes transfer data to be saved to the REC Disk buffer.

If the transfer mode is Consistency mode and a REC buffer is set, "halt(use-buffer)" is displayed when issue of a hard suspend causes transfer data to be saved to the REC buffer.

Note that hard-suspend status is still displayed as "halt" even if this option is specified for PRIMECLUSTER GD snapshot linkage.

-O

Changes the error suspend(failed) status Status column display method.

If this option is specified, the display is as follows:

There is a bad sector: "failed(badsector)"

When there is insufficient Snap Data Volume or Snap Data Pool capacity: "failed(overflow)"

All other cases: "failed(other)"

However, when linked with PRIMECLUSTER GD Snapshot, only "failed" is displayed when an error suspend status occurs even if -O option is specified.

-Xda

Displays the access permission settings for copy destination volume.

-Xdate

Displays the timestamp of the last copy session status transition.

-Xstandby-s

If the Storage Cluster Continuous Copy function is used, the Standby-Side Continuous Copy Session status is displayed.

-Xall

Displays the all information.

For the items displayed by this option, refer to the example following.

Note that the items displayed by this option may be subject to change in a future release.

OPERANDS

originalVolumeName

Specifies a replication source volume name.

For the volume names of other servers, use the format "*Volume-Name@Managed-Server-Name*".

Specifies the [AdvancedCopy Manager device name](#) for the *Volume-Name*.

Displays all operation statuses for the specified replication source volume.

If this operand is omitted, all the operation statuses of the Managed Server on which this command is executed are displayed.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Direct that all the replication operation statuses for Managed Server (SRC-SV) be displayed on the Management Server:

```
# /opt/FJSVswsrp/bin/swsrpstat -h SRC-SV
Server Original-Volume          Replica-Volume          Direction Status          Execute
SRC-SV /dev/dsk/clt2d0s3@SRC-SV /dev/dsk/clt2d0s3@TARG-SV regular replicated      ----
SRC-SV /dev/dsk/clt2d1s3@SRC-SV /dev/dsk/clt2d1s3@TARG-SV regular      ----      ----
SRC-SV /dev/dsk/clt2d2s3@SRC-SV /dev/dsk/clt2d2s3@TARG-SV reverse snap          45%
SRC-SV /dev/dsk/clt1d0s3@SRC-SV /dev/dsk/clt1d1s3@TARG-SV regular copy-on-write(inactive) ----
SRC-SV /dev/dsk/clt1d0s3@SRC-SV /dev/dsk/clt1d2s3@TARG-SV regular copy-on-write(inactive) ----
SRC-SV /dev/dsk/clt1d0s3@SRC-SV /dev/dsk/clt1d3s3@TARG-SV regular copy-on-write(active)  ----
#
```

- Direct that all the replication operation statuses for Managed Server (SRC-SV) be displayed on the Management Server (with the -L option specified):

```
# /opt/FJSVswsrp/bin/swsrpstat -h SRC-SV -L
Server Original-Volume          Replica-Volume          Direction Status          Execute
Trk Update Rcv Split Xfer Snap-Gen
SRC-SV /dev/dsk/clt2d0s3@SRC-SV /dev/dsk/clt2d0s3@TARG-SV regular replicated      ----
---- auto ---- async ----
SRC-SV /dev/dsk/clt2d1s3@SRC-SV /dev/dsk/clt2d1s3@TARG-SV regular      ----      ----
on 6% ---- ---- ---- ----
SRC-SV /dev/dsk/clt2d2s3@SRC-SV /dev/dsk/clt2d2s3@TARG-SV reverse snap          45%
off ---- ---- ---- ----
SRC-SV /dev/dsk/clt1d0s3@SRC-SV /dev/dsk/clt1d1s3@TARG-SV regular copy-on-write(inactive)
---- off 0% ---- ---- ---- 1
SRC-SV /dev/dsk/clt1d0s3@SRC-SV /dev/dsk/clt1d2s3@TARG-SV regular copy-on-write(inactive)
---- off 5% ---- ---- ---- 2
SRC-SV /dev/dsk/clt1d0s3@SRC-SV /dev/dsk/clt1d3s3@TARG-SV regular copy-on-write(active) ----
off 6% ---- ---- ---- 3
#
```

- Direct that all the replication operation statuses for Managed Server (SRC-SV) be displayed on the Management Server (with the -Xall option specified):

```
# /opt/FJSVswsrp/bin/swsrpstat -h SRC-SV -Xall
Server Original-Volume          Replica-Volume          Direction Status          Execute Trk Update
Rcv Split Xfer Snap-Gen Copy-DA Date          Standby-Session
```

```

SRC-SV /dev/dsk/clt2d0s3@SRC-SV /dev/dsk/clt2d0s3@TARG-SV regular replicated ---- ---- ----
auto ---- async ---- on 2010/10/20 22:00:45 exist
SRC-SV /dev/dsk/clt2d1s3@SRC-SV /dev/dsk/clt2d1s3@TARG-SV ---- ---- ---- on 6%
---- ---- ---- off 2010/10/20 22:10:39 not-exist
SRC-SV /dev/dsk/clt2d2s3@SRC-SV /dev/dsk/clt2d2s3@TARG-SV reverse snap 45% off ----
---- ---- ---- on 2010/10/21 22:00:23 ?????
#

```

- Direct that all the replication operation statuses for Managed Server (SRC-SV) be displayed on the Management Server (with the -Xdate option specified):

```

# /opt/FJVSwsrpb/bin/swsrpstat -h SRC-SV -Xdate
Server Original-Volume Replica-Volume Direction Status Execute Date
SRC-SV /dev/dsk/clt2d0s3@SRC-SV /dev/dsk/clt2d0s3@TARG-SV regular replicated ---- 2010/10/20
22:00:45
SRC-SV /dev/dsk/clt2d1s3@SRC-SV /dev/dsk/clt2d1s3@TARG-SV ---- ---- 2010/10/20
22:10:39
SRC-SV /dev/dsk/clt2d2s3@SRC-SV /dev/dsk/clt2d2s3@TARG-SV reverse snap 45% 2010/10/21
22:00:23
#

```

- Direct that all the replication operation statuses for Managed Server (SRC-SV) be displayed on the Management Server (with the -Xstandby-s option specified):

```

# /opt/FJVSwsrpb/bin/swsrpstat -h SRC-SV -Xstandby-s
Server Original-Volume Replica-Volume Direction Status Execute Standby-Session
SRC-SV /dev/dsk/clt2d0s3@SRC-SV /dev/dsk/clt2d0s3@TARG-SV regular replicated 45% exist
SRC-SV /dev/dsk/clt2d1s3@SRC-SV /dev/dsk/clt2d1s3@TARG-SV regular replicated 10% not-exist
SRC-SV /dev/dsk/clt2d2s3@SRC-SV /dev/dsk/clt2d2s3@TARG-SV reverse equivalent 100% ?????
#

```

NOTES

- When the execution status (the status displayed in the "Status" column) is "failed", "halt", or "?????", refer to ["8.4 Troubleshooting: Hardware or Other Error During Replication"](#) and take action.
- If the execution status (displayed in the Status column) is "gds-error":
A PRIMECLUSTER GD error probably occurred. Remove the cause of the PRIMECLUSTER GD error, then use the `swsrpcancel` command to stop the replication operation.
- The operation status cannot be displayed if:
 - The volumes have not been set as replication volumes.
 - When the original volume or replica volume is a logical volume of VxVM and the structure of the logical volume is changed to a structure not supported by AdvancedCopy Manager or the logical volume is deleted from VxVM.
- This command cannot be executed while the following command is running:
 - `swsrprecoverres`
- If the Storage Cluster Continuous Copy function is used, some corrective action may be required according to what is displayed in the **Standby-Session** column.
 - For "not-exist":
Refer to ["7.10 When Recovering Storage Cluster Continuous Copy Function"](#).
 - Other than those above:
No action is required.

9.3.2.4 swsrpcancel (Replication Cancellation Command)

NAME

`swsrpcancel` - stops snapshot processing and synchronization processing

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsrp/bin/swsrpcancel [-c] {[[-f] [-t] [-m] [-b]]|[-T]|[-Xforce]} [-Xgds-slice-skipchk]
fromVolumeName toVolumeName
```

To Execute on Management Server

```
/opt/FJSVswsrp/bin/swsrpcancel [-h serverName] [-c] {[[-f] [-t] [-m] [-b]]|[-T]|[-Xforce]} [-Xgds-
slice-skipchk] fromVolumeName toVolumeName
```

To Execute on Managed Server: Group Specification

```
/opt/FJSVswsrp/bin/swsrpcancel [-c] {[[-f] [-t] [-m] [-b]]|[-T]|[-Xforce]} [-Xreverse] -Xgroup
groupName
```

To Execute on Management Server: Group Specification

```
/opt/FJSVswsrp/bin/swsrpcancel [-h serverName] [-c] {[[-f] [-t] [-m] [-b]]|[-T]|[-Xforce]} [-
Xreverse] -Xgroup groupName
```

DESCRIPTION

This command stops snapshot processing (OPC/QuickOPC/SnapOPC/SnapOPC+) and synchronization processing (EC) of the specified replication volume. Use this command when an error is detected during the replication operation and when you want to stop replication processing.

If the Storage Cluster Continuous Copy function is used, the replication process is stopped on both the Primary Storage and Secondary Storage.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

When performing the server-to-server replication, specify the name of the operation server configured with the [swsrpsetvol](#) command.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdisprv](#) command.

-c

Specifies that the copy processing is forcibly stopped from a non-operation server when the operation server cannot be used because of a system failure. If this option is used, the copy processing is stopped without implementing pre-processing and post-processing.

This option cannot be specified on the operation server.

-m

For server-to-server replication, this option specifies that communication processing is not performed with non-operation servers. When this option is specified, volume status check processing and pre-processing and post-processing are not performed for the replication destination volume of a non-operation server.

This option is valid only during server-to-server replication.

-f

When stopping synchronous processing under the equivalency maintained status, specify that pre-processing and post-processing for the copy source volume is not being executed.

Use this option only when AdvancedCopy Manager does not execute pre-processing and post-processing of the copy source volume and unique pre-processing and post-processing scripts are inserted before and after (respectively) the AdvancedCopy Manager command, or when you judge that copy pre-processing and post-processing are unnecessary (e.g., for a database constructed on a raw device).

-t

When stopping the synchronous processing under the equivalency maintained status, this specifies that post-processing for the copy destination volume is not executed.

Use this option only when AdvancedCopy Manager does not execute pre-processing and post-processing of the copy destination volume and unique pre-processing and post-processing scripts are inserted before and after (respectively) the AdvancedCopy Manager command, or when you judge that copy pre-processing and post-processing are unnecessary (e.g., for a database constructed on a raw device).

-b

Specifies that the inter-box copy operation in the halt status be stopped.

For replication on single server, the copying is cancelled for both boxes.

For server-to-server replication, the copying is cancelled only for the box connected to the server on which this command was executed.

Use the option only for a copy operation in the halt status and whose recovery is not expected to be done successfully.

The option is valid only for a copy operation in the HALT status.

The option can be used on the server that is not specified as the operation server.

This option cannot be used for the logical volume of an SDX object.

If this option is specified together with the -Xgroup option, copy processing is only stopped for sessions in the group whose status is halt.

-T

Specifies that differential snapshot processing is terminated.

This option only terminates tracking processing and does not terminate snapshot processing (OPC physical copying) when it is in progress. This option releases the OPC session if snapshot processing (i.e., OPC physical copying) is complete.

This option can only be used while physical copying or tracking is in progress.

This option cannot be used for the logical volume of an SDX object.

If this option is specified together with the -Xgroup option, only tracking processing within the group is stopped.

-Xgroup *groupName*

Specifies that copy processing is stopped for each group.

-Xreverse

Specifies that copy processing that is being executed in the reverse direction (from replica volumes to original volumes) is stopped

This option is valid only when the -Xgroup option is specified.

-Xgds-slice-skipchk

Skips the SDX object status check processing that is usually performed as part of the copy source and copy destination pre-processing. This takes place if SDX objects are operated in slice units and the synchronization processing for equivalency maintenance state is stopped.

If this option is specified for a stopped volume, the -f and -t options must also be specified.

-Xforce

Specification of this option forcibly stops a SnapOPC+ session.

Forcibly stopping SnapOPC+ stops the specified snap generation and all earlier snap generations.

This option can be specified even if there is only one snap generation.

This option is a SnapOPC+ fixed option.

If this option is specified at the same time as the -Xgroup option, the SnapOPC+ sessions in that group are stopped, and all the snap generations and snap generations before it in the group are also stopped.

OPERANDS

fromVolumeName

Specifies a copy source volume name.

Specifies the replication source volume or the replication destination volume set with the [swsrpsetvol](#) command.

Specifies the [AdvancedCopy Manager device name](#) for the volume name.

toVolumeName

Specifies a copy destination volume name.

Specifies the replication source volume or the replication destination volume set with the [swsrpsetvol](#) command.

Specifies the [AdvancedCopy Manager device name](#) for the volume name.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- On the Management Server, direct that copying from /dev/dsk/c1t2d0s3 to /dev/dsk/c1t2d0s3@TARG-SV be stopped for Managed Server SRC-SV:

```
#/opt/FJSVswsrp/bin/swsrpcancel -h SRC-SV /dev/dsk/c1t2d0s3 /dev/dsk/c1t2d0s3@TARG-SV
FROM=/dev/dsk/c1t2d0s3@SRC-SV,TO=/dev/dsk/c1t2d0s3@TARG-SV swsrpcancel completed
#
```

- Stops copy processing for group "GRP1":

```
# /opt/FJSVswsrp/bin/swsrpcancel -Xgroup GRP1
GROUP=GRP1 swsrpcancel completed
#
```

NOTES

- Following execution with the -b option specified in replication on single server, the command stops normally when cancellation of the copying succeeds in either of the boxes.
- Copy processing cannot be performed if:
 - The specified replication source volume and replication destination volume have not been set as replication volumes.
 - The copy processing is not performed for the specified replication source and replication destination volumes.
 - For server-to-server replication, the Managed Server on which the command is executed is not the operation server of the replication volume.
 - For server-to-server replication, communication with a destination server fails.
 - When the original volume or replica volume is a logical volume of VxVM and the structure of the logical volume is changed to a structure not supported by AdvancedCopy Manager or the logical volume is deleted from VxVM.
 - The backup management function is executing a process for a specified source or destination volume
- If the -Xgroup option is specified, processing is performed for each active replication volume, but processing is terminated without cancelling replication in the following case:
 - If no replication volumes have been copied
 - When you execute -b option specifying a nonexistent copy in the "halt" state in the group.
 - When you execute -T option specifying a nonexistent copy in the tracking process of QuickOPC in the group.
 - When you execute with the -Xforce option specified and there are no SnapOPC+ sessions in the group

Additionally, if the -Xgroup option is specified, replication volumes are processed in order, but processing is terminated immediately if an error occurs.

- When stopping SnapOPC+ sessions, the oldest snap generation must be stopped first.
Specify the -Xforce option to forcibly stop the second and subsequent snap generations. However, note that this option stops the specified snap generation and all earlier snap generations.
If the -Xforce and -Xgroup options are specified at the same time, the SnapOPC+ sessions in the group are stopped, and all the snap generations and snap generations before it in the group are also stopped.

- The volume information for the specified volume is not checked (check for matching of the volume information in the management list and the actual physical information). The operation is cancelled unconditionally using the volume information in the management list.
- The following processing is performed depending on the operation status of the specified replication source volume and replication destination volume.

Table 9.10 Processing Differences in Accordance with Operating Status

Operation Status	Processing
Reverse copy direction	Suspends the processing.
Not operating yet	Suspends the processing.
Total or incremental copy is performed	Suspends synchronization processing. The replication destination volume cannot be used.
Equivalency maintain status	Suspends synchronization processing. The replication destination volume can be used as a copy
Replication established status	Suspends synchronization processing. The replication destination volume can be used as a copy
Snapshot processing is being performed	Stops the snapshot processing. The replication destination volume cannot be used.

- Pre-processing and post-processing for a replication source and replication destination volumes are performed only when the status of synchronization processing is equivalency maintenance status.
- Synchronization processing cannot be stopped in equivalency maintenance status in the following situations:
 - When a replication source volume is an SDX object (SynfinityDISK) and the status of the SDX object is not one of the following:
 - The mirror volume is ACTIVE or STOP.
 - The mirror slice status is ACTIVE or TEMP
When the mirror slice is TEMP, a reproduction former volume is not pre-processed.
 - The SDX disk status is ENABLE.
 - When copy source volume or copy destination volume is a SynfinityFILE of multi-volume composition and the file system is mounted. In this case, perform unmount for all the volumes with SynfinityFILE.
- When the -T option is not specified command processing varies depending on the operation status of the source and destination volumes.

Table 9.11 Processing When -T Option Is Not Specified

Operation Status (Display by swsrpstat -L)	Processing
Reverse copy	Processing is interrupted.
Not in operation (Status is "----" and Trk also is "----".)	Processing is interrupted.
Total or difference copying is in progress. (Status is "sync".)	Synchronous processing is terminated. The destination volume cannot be used.
Equivalency maintained status (Status is "equivalent".)	Synchronous processing is terminated. The destination volume can be used as a replica.
Replica created status (Status is "replicated".)	Synchronous processing is terminated. The destination volume can be used as a replica.
Snapshot processing is in progress. (Status is "snap")	Synchronous processing is terminated. The destination volume cannot be used. Tracking processing, if in progress, is also terminated.

Operation Status (Display by swsrpstat -L)	Processing
Snapshot processing complete status and tracking processing in progress. (Status is "----" and Trk also is "on".)	Tracking processing is terminated. The destination volume can be used as a replica.

- When the -T option is specified command processing varies depending on the operation status of the source and destination volumes.

Table 9.12 Processing When -T Option Is Specified

Operation status (Display by swsrpstat -L)	Processing
Reverse copy	Processing is interrupted.
Not in operation (Status is "----" and Trk also is "----".)	Processing is interrupted.
Total or difference copying is in progress. (Status is "sync".)	Processing is interrupted.
Equivalency maintained status (Status is "equivalent".)	Processing is interrupted.
Replica created status (Status is "replicated".)	Processing is interrupted.
Snapshot processing is in progress and tracking processing is not in progress. (Status is "snap" and Trk is "off")	Processing is interrupted.
Snapshot processing is in progress and tracking processing is progress. (Status is "snap" and Trk is "on".)	Only tracking processing is terminated. Snapshot processing (ie, physical copying) is continued. The destination volume can be used as a replica.
Snapshot processing complete status and tracking processing in progress. (Status is "----" and Trk also is "on".)	Tracking processing is terminated. The destination volume can be used as a replica.

- When SDX objects of PRIMECLUSTER GD are used: Refer to "[Configuration and Conditions of Unsupported SDX Objects](#)" and "[10.1.5.2 Notes on SDX Object Operations \(Replication Management\)](#)".
- This command cannot be executed while any of the following commands are running:
 - [swsrpsetvol](#)
 - [swsrpdelvol](#)
 - [swsrprecoverres](#)
- When the Storage Cluster Continuous Copy function is used:
 - When the status of the TFO group is "Normal", and replication processing cannot stop on either the Primary Storage or Secondary Storage, the command may terminate abnormally. Check the cause of the error on both the Primary Storage and Secondary Storage to address the error message in the storage device where the error is occurring.
 - When the status of TFO group is "Normal" and its phase is "Maintenance", or when the status of the TFO group is "Halt", the command starts synchronous processing on the storage that has an "Active" TFO status and terminates normally. However, synchronous processing could not be started on the storage that has a "Standby" TFO status. In this case, swsrp2882 message may be output.
If a swsrp2882 message is output, after the status of TFO group has become "Normal" and its phase is not "Maintenance", refer to "[7.10 When Recovering Storage Cluster Continuous Copy Function](#)" to take action.

9.3.2.5 swsrpchsnc (Synchronization Mode Change Command)

NAME

swsrpchsnc - changes the operation mode

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsrp/bin/swsrpchsnc [-F {Sync|Async|Stack|Consist}] [-R {Auto|Manual}] [-S {Auto|Manual}] fromVolumeName toVolumeName
```

To Execute on Management Server

```
/opt/FJSVswsrp/bin/swsrpchsnc [-h serverName] [-F {Sync|Async|Stack|Consist}] [-R {Auto|Manual}] [-S {Auto|Manual}] fromVolumeName toVolumeName
```

To Execute on Managed Server: Group Specification

```
/opt/FJSVswsrp/bin/swsrpchsnc [-F {Sync|Async|Stack|Consist}] [-R {Auto|Manual}] [-S {Auto|Manual}] [-Xreverse] -Xgroup groupName
```

To Execute on Management Server: Group Specification

```
/opt/FJSVswsrp/bin/swsrpchsnc [-h serverName] [-F {Sync|Async|Stack|Consist}] [-R {Auto|Manual}] [-S {Auto|Manual}] [-Xreverse] -Xgroup groupName
```

DESCRIPTION

This command changes the operation mode of inter-box synchronization (transfer mode, recovery mode, or split mode).

OPTIONS

-h *serve-Name*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

When performing the server-to-server replication, specify the name of the operation server configured with the [swsrpsetvol](#) command.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdispsrv](#) command.

-F {Sync|Async|Stack|Consist}

Changes the transfer mode.

You can specify one of the following parameters in the operand:

Sync : Changes to Synchronous mode.

Async : Changes to Through mode.

Stack : Changes to Stack mode.

Consist : Changes to Consistency mode.

Changing the transfer mode requires this option.

You can only specify the -S option together with the -F Sync option.

Specifying the -F Sync option to change the transfer mode from a non-synchronous mode to the synchronous mode without specifying the -S option sets the automatic split mode.

-R {Auto|Manual}

Changes the recovery mode.

You can specify either of the following parameters in the operand:

Auto : Changes to Automatic Recovery mode.

Manual : Changes to Manual Recovery mode.

Changing the recovery mode requires this option.

-S {Auto|Manual}

Changes the split mode.

You can specify either of the following parameters in the operand:

Auto : Changes to Automatic Split mode.

Manual : Changes to Manual Split mode.

Changing the split mode requires this option.

To specify this option, you must also specify -F Sync, or the transfer mode of the current synchronization must be synchronous mode.

-Xgroup *groupName*

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) is changed.

This option is valid only when the -Xgroup option is specified.

OPERANDS

fromVolumeName

Specifies a copy source volume name.

Specifies the replication source volume or the replication destination volume set with the [swsrpsetvol](#) command.

Specifies the [AdvancedCopy Manager device name](#) for the volume name.

toVolumeName

Specifies a copy destination volume name.

Specifies the replication source volume or the replication destination volume set with the [swsrpsetvol](#) command.

Specifies the [AdvancedCopy Manager device name](#) for the volume name.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- From Management Server to Managed Server: Instruct SRC-SV to change the transfer mode for inter-box synchronization between /dev/dsk/c1t2d0s3@SRC-SV and /dev/dsk/c1t2d10s3@TARG-SV from Consistency mode to Through mode:

```
#!/opt/FJVSwsrp/bin/swsrpstat -h SRC-SV -L /dev/dsk/c1t2d0s3@SRC-SV
Server Original-Volume      Replica-Volume      Direction Status Execute Trk  Rcv  Split
Xfer   Snap-Gen
SRC-SV /dev/dsk/c1t2d0s3@SRC-SV /dev/dsk/c1t2d10s3@TARG-SV regular  sync  99%   ---- auto ----
consist ----
:
#!/opt/FJVSwsrp/bin/swsrpchsnc -h SRC-SV -F Async /dev/dsk/c1t2d0s3@SRC-SV /dev/dsk/
c1t2d10s3@TARG-SV
FROM=/dev/dsk/c1t2d0s3@SRC-SV, TO=/dev/dsk/c1t2d10s3@TARG-SV swsrpchsnc completed
:
#!/opt/FJVSwsrp/bin/swsrpstat -h SRC-SV -L /dev/dsk/c1t2d0s3@SRC-SV
Server Original-Volume      Replica-Volume      Direction Status      Execute Trk  Rcv
Split Xfer  Snap-Gen
SRC-SV /dev/dsk/c1t2d0s3@SRC-SV /dev/dsk/c1t2d10s3@TARG-SV regular  equivalent 100%   ---- auto
----  async ----
#
```

- Specify that the transfer mode of inter-box synchronous processing for group "GRP1" be changed from Consistency mode to Through mode:

```
# /opt/FJSVswsrp/bin/swsrpchsnc -F Async -Xgroup GRP1
GROUP=GRP1 swsrpchsnc completed
#
```

NOTES

- In any of the following cases, you cannot change the operation mode of inter-box synchronization:
 - The specified original/replica volume has not been configured as a replica volume.
 - None of the -F, -R and -S options are specified.
 - In inter-server replication, the Managed Server on which you execute the command is not the operation server for the replica volume.
 - Synchronization from the specified source volume to the destination volume is not secured.
 - Synchronization from the specified source volume to the destination volume is in error-suspended status (with "failed" indicated as the Status by the [swsrpstat](#) command or hard-suspended status (with "halt" indicated as the Status by the [swsrpstat](#) command.
 - The ETERNUS Disk storage system in which the specified original/replica volume is placed do not support the operation mode change function.
 - When the original volume or replica volume is a logical volume of VxVM and the structure of the logical volume is changed to a structure not supported by AdvancedCopy Manager or the logical volume is deleted from VxVM.
 - The specified volume is being used with PRIMECLUSTER GD Snapshot.
- If the -Xgroup option is specified, the synchronous processing for all replication volumes being processed is changed to the same operation mode. However, the operation mode cannot be changed if there is at least one replication volume that meets the conditions above. Additionally, if the -Xgroup option is specified, replication volumes are processed in order, but processing is terminated immediately if an error occurs.
- This command cannot be executed while any of the following commands are running:
 - [swsrpsetvol](#)
 - [swsrpdelvol](#)
 - [swsrprecoverres](#)

9.3.2.6 swsrprevsync (Synchronization Reverse Command)

NAME

swsrprevsync - reverses the copying direction of synchronization

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsrp/bin/swsrprevsync fromVolumeName toVolumeName
```

To Execute on Management Server

```
/opt/FJSVswsrp/bin/swsrprevsync [-h serverName] fromVolumeName toVolumeName
```

To Execute on Managed Server: Group Specification

```
/opt/FJSVswsrp/bin/swsrprevsync [-Xreverse] -Xgroup groupName
```

To Execute on Management Server: Group Specification

```
/opt/FJJSVswsrp/bin/swsrprevsync [-h serverName] [-Xreverse] -Xgroup groupName
```

DESCRIPTION

This command reverses the copying direction of synchronization when in suspended status.

Note that this command cannot be used to handle a logical volume of an SDX object.

If the Storage Cluster Continuous Copy function is used, the replication process is reversed on both the Primary Storage and Secondary Storage.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

When performing the server-to-server replication, specify the name of the operation server configured with the [swsrpsetvol](#) command.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdisprv](#) command.

-Xgroup *groupName*

Changes the direction of synchronous processing for each group.

-Xreverse

Specifies that the direction of synchronous processing that is being executed in reverse direction (i.e., from replica volumes to original volumes) is inverted.

This option is valid only when the -Xgroup option is specified.

OPERANDS

fromVolumeName

Specifies a copy source volume name. (as the destination in the reverse direction)

Specifies the replication source volume or the replication destination volume set with the [swsrpsetvol](#) command.

Specifies the [AdvancedCopy Manager device name](#) for the volume name.

toVolumeName

Specifies a copy destination volume name. (as the source in the reverse direction)

Specifies the replication source volume or the replication destination volume set with the [swsrpsetvol](#) command.

Specifies the [AdvancedCopy Manager device name](#) for the volume name.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- From Management Server to Managed Server: Instruct SRC-SV to reverse the copying direction of synchronization between /dev/dsk/c1t2d0s3@SRC-SV and /dev/dsk/c1t2d10s3@TARG-SV:

```
#!/opt/FJJSVswsrp/bin/swsrpstat -h SRC-SV /dev/dsk/c1t2d0s3@SRC-SV
Server Original-Volume      Replica-Volume      Direction  Status    Execute
SRC-SV /dev/dsk/c1t2d0s3@SRC-SV /dev/dsk/c1t2d10s3@TARG-SV reverse    replicated ----
#!/opt/FJJSVswsrp/bin/swsrprevsync /dev/dsk/c1t2d10s3@TARG-SV /dev/dsk/c1t2d0s3@SRC-SV
FROM=/dev/dsk/c1t2d0s3@SRC-SV,TO=/dev/dsk/c1t2d10s3@TARG-SV swsrprevsync completed
#!/opt/FJJSVswsrp/bin/swsrpstat /dev/dsk/c1t2d0s3@SRC-SV
Server Original-Volume      Replica-Volume      Direction  Status    Execute
SRC-SV /dev/dsk/c1t2d0s3@SRC-SV /dev/dsk/c1t2d10s3@TARG-SV regular    replicated ----
#
```

- Change the direction of synchronous processing for group "GRP1":

```
# /opt/FJSVswsrp/bin/swsrprevsync -Xgroup GRP1
GROUP=GRP1 swsrprevsync completed
#
```

NOTES

- In any of the following cases, you cannot reverse the direction of synchronization:
 - The specified original/replica volume has not been configured as a replica volume.
 - Unidirectional copying has been specified in the copying attribute for the replica volume.
 - In inter-server replication, the Managed Server on which you execute the command is not the operation server for the replica volume.
 - Replication has not been established in synchronization from the specified source volume to the destination volume.
 - The ETERNUS Disk storage system in which the specified original/replica volume is placed do not support the reverse function.
 - When the original volume or replica volume is a logical volume of VxVM and the structure of the logical volume is changed to a structure not supported by AdvancedCopy Manager or the logical volume is deleted from VxVM.
 - The specified volume is being used with PRIMECLUSTER GD Snapshot.
- If the -Xgroup option is specified, the direction of synchronous processing for all replication volumes being processed is changed. However, the direction of synchronous processing cannot be changed if there is at least one replication volume that meets the conditions above. Additionally, if the -Xgroup option is specified, replication volumes are processed in order, but processing is terminated immediately if an error occurs.
- If the -Xgroup option is specified, processing is not performed for those sessions in the group where the direction has already been changed.
- This command cannot be executed while any of the following commands are running:
 - [swsrpsetvol](#)
 - [swsrpdelvol](#)
 - [swsrprecoverres](#)
- When the Storage Cluster Continuous Copy function is used:
 - When the status of the TFO group is "Normal", and synchronous processing cannot reverse on either the Primary Storage or Secondary Storage, the command may terminate abnormally. Check the cause of the error on both the Primary Storage and Secondary Storage to address the error message in the storage device where the error is occurring.
 - When the status of the TFO group is "Normal" and its phase is "Maintenance", or when the status of TFO group is "Halt", the command reverses the synchronous processing on the storage that has an "Active" TFO status and terminates normally.

9.3.3 Maintenance Commands

This section describes the maintenance commands for replication management.

9.3.3.1 swsrprecoverres (Resource Adjustment Command)

NAME

swsrprecoverres - executes recovery to make replication management information

SYNOPSIS

To Execute on Managed Server

```
/opt/FJSVswsrp/bin/swsrprecoverres [-r] [-p]
```

To Execute on Management Server

```
/opt/FJVSwsrp/bin/swsrprecoverres [-h serverName] [-r] [-p]
```

DESCRIPTION

This command executes recovery to make replication management information consistent if a system failure or other problem has made the replication management information inconsistent.

OPTIONS

-h *serverName*

When executing on the Management Server, specify the name of the target Managed Server.

When executing on the target Managed server, it is not necessary to specify this option.

For the server name, accurately specify the same letters, including upper and lower case, as the server name displayed with the [stgxfwcmdisprv](#) command.

-r

Specify this option when an incorrect status exists in the replication management files. E.g., when the [swsrpsetvol](#) or [swsrpdelvol](#) command that update the replication management files terminate suddenly because of a forced stop, a system failure or the replication management files have been deleted by mistake.

When this option is specified, the following processing is performed to recreate the replication management files.

Replication management information in the Management Server repository is downloaded to the Managed Server, and the replication volume information is restored.

The ETERNUS Disk storage system is queried, and the execution information for the synchronization processing is restored.

-p

Specifies to not perform any post-processing that has not already been performed. This option can be used when the [swsrpmake](#) command processing terminates suddenly because of a forced stop or system failure.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Perform resource adjustment:

```
#!/opt/FJVSwsrp/bin/swsrprecoverres
swsrprecoverres completed
#
```

NOTES

- The resource adjustment processing cannot be performed if:
 - Communication with a Management Server fails and the -r option is specified.
 - The -r option is not specified when this command is executed after processing by the [swsrpsetvol](#) and [swsrpdelvol](#) commands has been forcibly terminated.
 - When the original volume or replica volume is a logical volume of VxVM and the structure of the logical volume is changed to a structure not supported by AdvancedCopy Manager or the logical volume is deleted from VxVM.
- The consistency of resource information is restored only for the Managed Server on which this command is executed. When a resource inconsistency occurs during processing related to server-to-server replication, this command must be executed both on the Managed Server on which the processing was performed and on the Managed Server with which communication is being made.
- If this command is executed for the copy source volume and the copy destination volume in the status where the [swsrpmake](#) command is interrupted without completing its process and the remaining processing must be completed, the replication post-processing is executed.

- This command cannot be executed while any of the following commands are running:
 - `swsrpsetvol`
 - `swsrpvolinfo`
 - `swsrpdelvol`
 - `swsrpstartsync`
 - `swsrpmake`
 - `swsrpstat`
 - `swsrpcancel`
 - `swsrpchsync`
 - `swsrprevsync`
 - `swsrprecoverres`

9.4 Start/Stop Daemon Commands

This section describes how to start and stop AdvancedCopy Manager daemons.



Notes when using commands:

- Confirm that the environment variable `SWSTGNODE` is not defined.
- Execute this command using the same environment variable `LANG` that of the code system used by the system.
- Only a root user can execute this command.
- For cluster operation, refer to "[Starting and Stopping Daemons](#)".

9.4.1 Starting and Stopping Communication Daemon

Starting Communication Daemon

Enter the following command to stop the communication daemon:

```
# /opt/FJSVswstf/bin/stgfwcom start
```

Stopping Communication Daemon

Enter the following command to stop the communication daemon:

```
# /opt/FJSVswstf/bin/stgfwcom stop
```



Termination of the communication daemon may take up to ten seconds to take effect.

In order to verify if the communication daemon has been launched normally, check if the message `stgxfws0060` has been generated in the `syslog`.

9.4.2 Starting and Stopping AdvancedCopy Manager Daemons

The daemons listed below are started and stopped together:

Table 9.13 Daemons Started and Stopped as Package

Server Type	Daemon Name
Managed Server	Communication daemon
Management Server	Communication daemon

Starting AdvancedCopy Manager Daemon

Enter the following command to start a AdvancedCopy Manager daemon:

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

Stopping AdvancedCopy Manager Daemon

Enter the following command to stop a AdvancedCopy Manager daemon:

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

9.5 SnapOPC/SnapOPC+ Support Commands

This section explains the commands can be used to execute SnapOPC/SnapOPC+.



Note

Operating Environment for Command Execution

The table below lists the requirements for the operating environment in which to execute the respective commands.

Table 9.14 Operating Environment When Executing SnapOPC/SnapOPC+ Support Commands

Command	Operating Environment of AdvancedCopy Manager
swstestupdate	-
swstsdv	-



Note

In Cluster Operation

In cluster operation, execute a command by performing the following steps:

1. If necessary, specify the logical node name to the transaction in the environment variable SWSTGNODE. Specify the environment variable as follows:

```
# SWSTGNODE=logicalNodeName
# export SWSTGNODE
```

2. Execute the command.

The table below lists the commands available in cluster operation, whether environment variables must be configured, and the nodes on which commands are executable.

Table 9.15 Commands When Cluster Operation Is Used

Command	Environment Variable	Command Execution Node	Remarks
swstestupdate	Not required	Transaction operation node	-
swstsdv	Not required	Transaction operating node	-

9.5.1 swstestupdate (Update Volume Measurement Command)

NAME

swstestupdate - measures the physical capacity of a Snap Data Volume

SYNOPSIS

To Start Update Volume Measurements (Setting Up Monitor Session)

```
/opt/FJSSVswsts/bin/swstestupdate start volumeName
```

To Look Up Update Volume (Displaying Status of Monitor Session)

```
/opt/FJSSVswsts/bin/swstestupdate status volumeName
```

To Finish Update Volume Measurements (Cancelling Monitor Session)

```
/opt/FJSSVswsts/bin/swstestupdate stop volumeName
```

DESCRIPTION

This command measures the update amount in the volume.

To Start Update Volume Measurements (Setting Up Monitor Session)

Specify "start" as the subcommand to start measure of the physical capacity of a Snap Data Volume.

To Look Up Update Volume (Displaying Status of Monitor Session)

Specify "status" as the subcommand to display the physical capacity of a Snap Data Volume.

The following information is displayed.

Title	Description
Volume-Name	The volume name.
Update	The amount of data that has been updated since measurements started in the number of blocks (1 block = 512 bytes). Displays "----" if there is no monitor session. Displays "failed" if the monitor session is error-suspended.

To Finish Update Volume Measurements (Cancelling Monitor Session)

Specify "stop" as the subcommand to stop measure of the physical capacity of a Snap Data Volume.

OPERANDS

volumeName

- Specifies the name of the measurement target volume.

PRIMECLUSTER GD logical volumes cannot be specified.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Measure the update volume:

```
(Start taking measurements by starting a monitor session)
#/opt/FJSSVswsts/bin/swstestupdate start /dev/dsk/c1t0d0s1
/dev/dsk/c1t0d0s1 swstestupdate completed
```

```

(First measurement)
#/opt/FJJSVswsts/bin/swstestupdate status /dev/dsk/c1t0d0s1
Volume-Name      Update
/dev/dsk/c1t0d0s1 644333

(Second measurement)
#/opt/FJJSVswsts/bin/swstestupdate status /dev/dsk/c1t0d0s1
Volume-Name      Update
/dev/dsk/c1t0d0s1 1045000

(Stop taking measurements by cancelling the monitor session)
# /opt/FJJSVswsts/bin/swstestupdate stop /dev/dsk/c1t0d0s1
/dev/dsk/c1t0d0s1 swstestupdate completed
#

```

NOTES

- The amount of data updated cannot be measured in the following cases:
 - When the specified volume does not exist.
 - When the specified volume is TFOV.
 - When the monitor session has already been set up (i.e., when the "start" option is specified).
 - When the monitor session does not exist (i.e., when the "stop" option is specified).

9.5.2 swstsdv (Snap Data Volume Operation/Reference Command)

NAME

swstsdv - Initialize or display the Snap Data Volume

SYNOPSIS

To Initialize Snap Data Volume

```

/opt/FJJSVswsts/bin/swstsdv init volumeName
/opt/FJJSVswsts/bin/swstsdv init -G accessVolumeName luNumber

```

To Display Snap Data Volume Status

```

/opt/FJJSVswsts/bin/swstsdv stat volumeName
/opt/FJJSVswsts/bin/swstsdv stat -G accessVolumeName luNumber

```

To Display Snap Data Pool Status

```

/opt/FJJSVswsts/bin/swstsdv poolstat -G volumeName

```

DESCRIPTION

Snap Data Volume initialization/status display and Snap Data Pool status display are all possible. Specify one of the following subcommands according to the contents you want to operate on, and then execute the subcommand.

Snap Data Volume Initialization

Specify "init" as the subcommand to initialize the Snap Data Volume.

Initialize the Snap Data Volume when SnapOPC/SnapOPC+ is not running.

All the data in the Snap Data Volume is destroyed in the initialization of the Snap Data Volume. For this reason, the following procedure is recommended.

(Creation of the Access volume and collection of the LU-Number must be completed before logical volume manager operations.)

1. If there is necessary data in the Snap Data Volume, back it up in advance using the relevant OS command.

2. Initialize the disk management information.
Set the target Snap Data Volume (physical disk) outside logical volume manager management when there is a logical volume manager operation.
3. Initialize the Snap Data Volume using this command.
Initialize the information using the Access volume when there is a logical volume manager operation.

Snap Data Volume Status Display

Specify "stat" as the subcommand to display Snap Data Volume logical capacity, physical capacity, used space capacity, and unused space capacity.

The contents that are displayed are as follows:

Title	Description
BoxID	Displays the Box ID.
LUN	Displays the logical unit number on the ETERNUS Disk storage system. The format used for the display is "<decimal><hexadecimal>".
Rate	Displays the Snap Data Volume usage as a percentage (%).
Logical	Displays the Snap Data Volume logical capacity. The unit used for display is 'sector'.
Physical	Displays the Snap Data Volume physical capacity. The unit used for display is 'sector'.
Used	Displays the amount of Snap Data Volume used space. The unit used for display is 'sector'. The capacity for both Copy usage (*1) and Host usage (*1) is displayed.
Copy	For the space used in the Snap Data Volume, Copy usage (*1) is displayed. The unit used for display is 'sector'.
Host	For the space used in the Snap Data Volume, Host usage (*1) is displayed. The unit used for display is 'sector'.
Pool	Displays the Snap Data Pool usage allocated to the Snap Data Volume. The unit used for display is 'sector'.

*1:

Copy usage: Usage in which the update occurs while SnapOPC/SnapOPC+ is running.

Host usage: Update of Snap Data Volume while SnapOPC/SnapOPC+ is not running.

Snap Data Pool Status Display

Specify "poolstat" as the subcommand to display the Snap Data Pool total amount, used space capacity, and unused space capacity.

The contents that are displayed are as follows:

Title	Description
BoxID	Displays the Box ID.
Pool-Type	Displays the Snap Data Pool type. Normal Encrypted
Rate	Displays the Snap Data Pool usage rate (%).
Total	Displays the Snap Data Pool total amount. The unit used for display is 'sector'.
Used	Displays the Snap Data Pool used space capacity. The unit used for display is 'sector'. The capacity for both Copy usage (*1) and Host usage (*1) is displayed.
Copy	For the space used in the Snap Data Pool, Copy usage (*1) is displayed. The unit used for display is 'sector'.
Host	For the space used in the Snap Data Pool, Host usage (*1) is displayed. The unit used for display is 'sector'.

Title	Description
Free	Displays the Snap Data Pool unused space capacity. The unit used for display is 'sector'.

*1:

Copy usage: Usage in which the update occurs while SnapOPC/SnapOPC+ is running.

Host usage: Update of Snap Data Volume while SnapOPC/SnapOPC+ is not running.

OPTIONS

Specific Options of Snap Data Volume Initialization(init)

-G accessVolumeName

Specifies the Access volume.

For details on Access volume, refer to "[Creating Access Volume](#)".

Specific Options of Snap Data Volume Status Display(stat)

-G accessVolumeName

Specifies the Access volume.

For details on Access volume, refer to "[Creating Access Volume](#)".

Specific Options of Snap Data Pool Status Display(poolstat)

-G volumeName

Specifies the ETERNUS Disk storage system volume.

If a Snap Data Volume is specified with insufficient capacity however, an error occurs in the command. In this case, specify an Access volume for which there is sufficient capacity.

For details on Access volume, refer to "[Creating Access Volume](#)".

OPERANDS

Specific Operands of Snap Data Volume Initialization(init)

volumeName

Specifies the volume name (AdvancedCopy Manager device name) of the Snap Data Volume to be initialized.

If a partition (slice) has been specified, the Logical Unit (disk) on which the partition (slice) exists is initialized.

luNumber

Specifies the logical unit number in the ETERNUS Disk storage system.

This value is specified as a decimal.

Specific Operands of Snap Data Volume Status Display(stat)

volumeName

Specifies the Snap Data Volume volume name (AdvancedCopy Manager device name).

If a partition (slice) has been specified, the information about the Logical Unit (disk) on which the partition (slice) exists is displayed.

luNumber

Specifies the logical unit number in the ETERNUS Disk storage system.

This value is specified as a decimal.

EXIT STATUS

=0: Completed successfully

>0: Terminated abnormally

EXAMPLES

- Snap Data Volume is initialized:

```
# /opt/FJSVswsts/bin/swstsdv init /dev/dsk/c2t18d4s2
swstsdv init completed
#
```

- The Snap Data Volume is initialized using the Access volume following insufficient capacity:

```
# /opt/FJSVswsts/bin/swstsdv init -G /dev/dsk/c2t18d1s2 110
swstsdv init completed
#
```

- The Snap Data Volume status is displayed:

```
# /opt/FJSVswsts/bin/swstsdv stat /dev/dsk/c2t18d4s2
BoxID = 00E4000M3#####E450S20A####KD4030639004##
LUN = 77 (0x4D)
Rate Logical(sector) Physical(sector) Used(sector) Copy(sector) Host(sector) Pool(sector)
0% 8388608 1048576 0 0 0 0
#
```

- The LU-Number (110) status is displayed using the Access volume:

```
# /opt/FJSVswsts/bin/swstsdv stat -G /dev/dsk/c2t18d1s2 110
BoxID = 00E4000M3#####E450S20A####KD4030639004##
LUN = 110 (0x6E)
Rate Logical(sector) Physical(sector) Used(sector) Copy(sector) Host(sector) Pool(sector)
0% 8388608 1048576 0 0 0 0
#
```

- The Snap Data Pool status is displayed by specifying the Access volume:

```
# /opt/FJSVswsts/bin/swstsdv poolstat -G /dev/dsk/c2t18d1s2
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
#
```

NOTES

- When there is insufficient Snap Data Volume capacity, operations are possible using the Access volume and logical unit number (LU-Number) in the ETERNUS Disk storage system.
When there is insufficient Snap Data Volume capacity, however, and "stat" is specified in the subcommand and an attempt is made to check the LU-Number of the Snap Data Volume with insufficient Snap Data Volume capacity, an access error occurs when the attempt to access the Snap Data Volume with insufficient Snap Data Volume capacity is made. In this case, an error occurs in the command and it is not possible to check the LU-Number.
For this reason, you must check the LU-Number in advance and record it.
- By specifying TFOV in the volume and access volume, Snap Data Volume cannot be initialized and Snap Data Volume's status or Snap Data Pool's status cannot be displayed.

9.6 Troubleshooting Data Collection Command

This section explains the command used to collect troubleshooting data when an AdvancedCopy Manager fault occurs.

9.6.1 esfsnap (AdvancedCopy Manager's Manager Troubleshooting Data Collection Command)

NAME

esfsnap - Collection of AdvancedCopy Manager's manager troubleshooting data

SYNOPSIS

```
/opt/FJSVesfcm/bin/esfsnap -dir dirname [-all]
```

DESCRIPTION

This command collects AdvancedCopy Manager troubleshooting data. Data is collected for AdvancedCopy Manager's manager installed on the Management Server where this command is executed.

Only the root user can execute this command.

Send the collected data to Fujitsu Technical Support.

OPERANDS

-dir *dirname*

Specify a directory where the troubleshooting data is stored. It is necessary to specify the full path name for *dirname*.

The collected troubleshooting data is stored in the directory which its name begins with "esfsnap_". Before executing this command, check that *dirname* directory has the following free space.

Collected Information	Required Free Space
Initial investigation information	more than 40 MB
All troubleshooting information	more than "80 + (2 * number of registered device)" MB

-all

Collects all troubleshooting data on servers. Data includes lots of information, which increases the size of the data. This option is not necessary for initial investigation.

EXIT STATUS

=0: Completed successfully
>0: Terminated abnormally

EXAMPLES

- Collect data to the /var/tmp directory:

```
# /opt/FJSVesfcm/bin/esfsnap -dir /var/tmp
```

NOTES

- Refer to [NOTES] of "esfsnap (CCM Server Troubleshooting Data Collection Command)" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module*.

9.6.2 acmcapture (AdvancedCopy Manager's Agent Troubleshooting Data Collection Command)

NAME

acmcapture - Collection of AdvancedCopy Manager's agent troubleshooting data

SYNOPSIS

```
/opt/FJSVswstf/bin/acmcapture dirname
```

DESCRIPTION

This command collects AdvancedCopy Manager's agent troubleshooting data. Data is collected for AdvancedCopy Manager's agent installed on the Managed Server where this command is executed.

Only the root user can execute this command.

Troubleshooting Data File Name

The collected data is stored in a file with the "acmcapture_HOSTNAME_YYYYMMDD_ttmss.tar.Z" format under the directory specified for "dirname".

The name of the host where the command is executed is used for HOSTNAME. If characters other than alphanumerics and hyphens are used in the host name, then "NOHOSTNAME" is used as the host name. Values are truncated to 15 characters.

- For YYYYMMDD, YYYY is the year, MM is the month, and DD is the day.
- For ttmss, tt is the hour, mm is minutes, and ss is seconds.

OPERANDS

dirname

Directory where the troubleshooting data is stored.

Specify a directory with plenty free space (500 MB or more) is recommended. Regard 500 MB as a rough guide, since the size of the troubleshooting data varies depending on the system environment and conditions. If there is not enough free space, then this command terminates abnormally without collecting troubleshooting data.

Specify an existing directory, otherwise this command terminates abnormally without collecting troubleshooting data.

EXIT STATUS

- =0: Completed successfully
- >0: Terminated abnormally

EXAMPLES

- Collect data to the /var/tmp directory:

```
# /opt/FJSVswstf/bin/acmcapture /var/tmp
Executing.....
Completed.
OUTPUT=/var/tmp/acmcapture_srv01_20090323_094614.tar.Z
```

NOTES

- In a clustered system, execute this command at all nodes. To gather information on the shared disk, mount the shared disk on a node and execute this command on that node.
- Execute this command when other AdvancedCopy Manager commands are in the stopped state. Data collection may not be performed correctly if other AdvancedCopy Manager commands are running.

Chapter 10 Notes on Operations

This chapter provides notes on the operations of AdvancedCopy Manager.

10.1 Notes on Backup and Replication Operation

10.1.1 General Notes

10.1.1.1 Backup and Replication Exclusions

The following devices must not be backed up or replicated:

- The device on which AdvancedCopy Manager has been installed
- The device on which the management list of AdvancedCopy Manager resides

10.1.1.2 Consistency of Data on Volumes

Table 10.1 Consistency of Data on Volumes

Data in Transaction Volume	Maintenance of Data Integrity	Operation
Oracle database space that is currently in use	Linkage with Oracle is required.	Use a sample script designed for linkage with Oracle. For details, refer to " Appendix B Sample Script for Backup Operations ".
File system	An AdvancedCopy Manager command unmounts the file system to maintain the integrity.	Refer to " Appendix A Pre-processing and Post-processing of Backup and Restoration " and " Appendix C Pre-processing and Post-processing of Replication ".
Other than the above	The data integrity must be maintained in operations.	Take appropriate action, such as stopping transactions during execution of backup or replication.

10.1.1.3 Fetching Device Information

Before starting either the backup or replication operation, The "[3.4.4 Fetching Device Information on Managed Server](#)" that is managed by all Managed Servers using Web Console. The time taken to complete this operation is proportional to the total number of devices defined on the selected Managed Servers. If many devices are involved, perform this operation while the CPU or I/O load on the system is low.

As a guide, about 0.5 seconds are required per device (i.e., partition) when there is no load on the system.

10.1.1.4 Running Software

All the devices are always in use while any of the following software is running. Therefore, before performing backup, restoration, or replication, perform the following steps:

Software Which Locks Disk (e.g., Antivirus Software, System Surveillance Software, etc.)

When performing backup/restore operation and replication operation, to prevent software from accessing a device targeted for service, suspend the software or specify the device for services other than check or monitoring by the software. For the setting method, refer to the manuals for respective software.

10.1.1.5 Maximum Number of Simultaneously Executing Operations

There is an upper limit (maximum value) to the number of Advanced copy sessions that can be simultaneously executed. For this reason, make an operational design so that Advanced copies exceeding the maximum number of copy sessions cannot simultaneously

be executed.

Also, when creating partitions in one Logical Unit (LU) and running Advanced copy at the same time, be sure not to exceed the maximum number of copy sessions simultaneously executable per LU.

For the upper limit to the number of copy sessions simultaneously executable, refer to the ETERNUS Disk storage system manuals.

10.1.1.6 Slice Containing VTOC

The header slice in VTOC-labelled disks may contain disk management information (VTOC).

Slices that contain this type of VTOC cannot be used in backup operations.

These slices can be used in replication operations if the `-Xvtoc` option is specified when the replication volume information is set.



Note

- The above restriction does not apply if there are separate operations for each physical disk that makes up the VxVM disk group.
- When a slice that contains VTOC is copied to the header slice (containing VTOC) of another disk, the VTOC of the copy destination disk is rewritten.
If a slice that contains VTOC is to be copied to the header slice (containing VTOC) of another disk, prepare the copy destination disk as a VTOC-labelled disk and set the same disk configuration, and then execute the copy process.
- Do not copy a slice that does not contain VTOC to the header slice (containing VTOC) of another disk.
If a slice that does not contain VTOC is copied to the header slice (containing VTOC) of another disk, the VTOC of the copy destination disk is rewritten, and access to the slices in the copy destination disk may no longer be possible.

10.1.1.7 SynfinityFILE and PRIMECLUSTER GFS Local File Systems

SynfinityFILE file systems comprised of multiple slices and the global file systems of PRIMECLUSTER GFS cannot be processed.

To backup or restore SynfinityFILE file systems comprised of multiple slices or the global file systems of PRIMECLUSTER GFS, process the multiple slices simultaneously in the unmounted state.

10.1.1.8 SynfinityFILE/Global File System and PRIMECLUSTER GFS Shared File System

Note the following points about backup, restoration, and replication of the SynfinityFILE/Global file system and PRIMECLUSTER GFS shared file system.

- The target file system must be unmounted in advance.
- Execute the operations on all partitions of the target file system.

The file system is unmounted as follows.

- SynfinityFILE/Global file system: Stop all daemons/services
- PRIMECLUSTER GFS shared file system: Unmount all nodes

10.1.1.9 Mounting of Disks from Other Servers

Before backup/restoration or replication of a volume that can be mounted from multiple servers, unmount and cancel any pending mount on the other servers.

For disks that need not be mounted from other servers, prevent multiple servers from detecting or accessing the same partition by setting up hardware, such as devices in the ETERNUS Disk storage system and Fibre Channel switches, accordingly.

10.1.1.10 Notes on Copy Operations Targeted at File Systems

For a copy operation targeted at a file system, unmount the relevant volume to prevent data access and maintain the data integrity. If backup/restoration and replication are performed without unmounting the volume, an error occurs since a volume in use cannot be unmounted.

Note the following about processing where unmounting is enabled:

- Another application is not using the volume. If another application is using it, suspend the application temporarily.

- Users are not using the volume. If users are using it, prohibit them from using it temporarily.
- Another volume is not mounted on the volume. If it is mounted, unmount it temporarily.
- The volume is not shared with NFS sharing. If it is shared, unshare it temporarily.

Unmounting is required only during command execution. Normal operations can be resumed after command execution.

10.1.1.11 Notes on Executing Logical Unit (Disk) Unit Copy

Logical Unit (disk unit) copy can be used for replication operations.

Note the following points if Logical Unit (disk unit) copy is executed:

Copying Between Logical Units (Disks) of Different Sizes

When copying between Logical Units (disks) of different sizes, the physical copy matches the size of the smaller Logical Unit (disk). If the copy direction is from the smaller Logical Unit (disk) to the larger Logical Unit (disk), part of the area on the larger disk is not physically copied to, and data from before the copy operation remains. If this remaining area contains the Logical Unit (disk) management information, partitioning tools may identify the disk format incorrectly. Therefore, format the copy destination Logical Unit (disk) before copying to it.

About GPT Disks

To copy the GPT disks by each Logical Unit (disk), create the source and destination volumes in the same size for replication in the unit of bytes.

When copying between volumes of different sizes, the GPT label of the destination volume may be changed. If you restore the destination volume whose GPT label is changed, you cannot use the destination volume.

Duplicate Disk Management Information

Copying between Logical Units (disks) of EFI labelled disks also copies the UUID disk management information that is used to uniquely identify disks and slices included in the Logical Unit (disk). Thus, after copy is executed, duplicate disk management information may exist within a system.

Disk management information is duplicated when the following types of copying are performed:

- Copying between Logical Units (disks) within a server
- Copying between Logical Units (disks) on different servers, if there are multiple copy destinations within the same copy destination server for the same copy source
- Copying from a volume that has previously been copied to a Logical Unit (disk) that is not the copy source but is on the same server as the copy source

Refer to "Solaris 10 System Administrator Collection" - "System Administration Guide: Devices and File Systems", "29. Copying UFS Files and File Systems (Tasks)", "Example 29-2" for the action to take if disk management information (UUID) is duplicated.

Notes When Executing Copy

- If the replication source/destination volume areas for replication management, and the transaction/backup volume areas for backup management duplicate each other, replication management and backup management operations cannot be performed simultaneously for volumes that have duplicate areas.
- Execute copy operations when the slices included in the Logical Unit (disk) are in the unmounted state. If copy is executed when the slices are in the mounted state, the slices included in the copy destination disk may not be able to be made consistent as a file system.

SDX Objects

Logical Unit (disk unit) copying cannot be used for SDX objects.

10.1.1.12 Notes on Executing Copy Between Different Operating Systems

Copies can be made between different operating systems by using replication operations.

Use the following combinations to execute copy between different operating systems:

Table 10.2 Combinations for Executing Copy Between Different Operating Systems

Copy Source	Copy Destination	
	Slices	Logical Unit (Disk)
slices	A	N/A
Logical Unit (disk)	A	N/A

(A=Available, N/A=Not Available)

10.1.1.13 Notes on Executing Copy from Logical Unit (Disk) to Slice (Partition) Using Replication Function

When a copy is performed, the Logical Unit (disk) data is written in the replication destination volume (slice). If a copy is performed with a file system existing on the replication destination volume (slice), because a data overwrite can occur to the file system depending on the copy, not only will the data in the file system become inaccessible but the copied data will become inaccessible through the file system as well.

10.1.1.14 Operations in ETERNUS Multipath Driver Environment

The access special files used for the ETERNUS Multipath Driver are as follows:

- Solaris standard special file (/dev/[r]dsk/c*t*d*s*)
- Mplb special file (/dev/FJSMplb/[r]dsk/mplb*s*)

The access special files that can be used vary depending on the version of the ETERNUS Multipath Driver. Refer to the ETERNUS Multipath Driver manual for details of the usable special access files.



Note

- If the Solaris standard special file is used
System volume backup operations are not possible.
This is because the ETERNUS Multipath Driver cannot use the Solaris standard special file as a boot device.
- If the mplb special file is used
Logical unit (disk unit) copy of EFI disk label disks is not possible.

10.1.1.15 Notes When Using -m Option

When replication occurs between servers, the -m option can be specified in the following commands in order to avoid communication processing to non-operating servers:

- [swsrpstartsync](#)
- [swsrpmake](#)
- [swsrpcancel](#)

When the -m option is specified, any one of the following conditions must apply in order to ensure that problems do not arise as a result of not performing volume pre-processing and post-processing at the non-operating server:

- The non-operating server has been shut down.
- Manual pre-processing and post-processing is performed for the volumes on the non-operating server.
- The target volume is a RAW device (Oracle, etc., that does not need pre-processing and post-processing), rather than a file system.

10.1.1.16 Notes When Using -t Option

When you execute the [swsrpstartsync](#) command with the -t option, volume lock fails with the [swsrpmake](#) command and the swsrp3607 message may be output.

If you execute the [swsrpstartsync](#) command with the -t option to perform independent pre-processing, check whether the pre-processing has properly been done.

If you determine that pre-processing and post-processing for replication destination volume is not needed and execute the [swsrpstartsync](#) command with the -t option, execute the [swsrpmake](#) command with the -t option.

10.1.1.17 Copy Processing Performance of Advanced Copy

The Advanced Copy processing is performed by the ETERNUS Disk storage system. Therefore, direct enquiries concerning copy processing performance to the support department for the ETERNUS Disk storage system.

10.1.1.18 Moving of VM Guest of VMware

Refer to the manual of the VMware vCenter Server for information on how to move the VM guest by using the functions of the VMware vCenter Server.

Do not move the VM guest that changes the disk type by utilizing the functions of the VMware vCenter Server. When the disk type has been changed, the following situations may occur.

- A "?????" is displayed in the status of the configuration information of the AdvancedCopy Manager. The status can be confirmed from the following information displayed with operations of the Command Line Interface or Web Console.

Confirmation Means	Command Name/Screen Name	Displayed Field
Command Line Interface	swstbackstat	Status
	swstsyncstat	Status
	swstreststat	Status
	swsthistdisp	Status
	swstexecstat	Execute
	swsttrkstat	Status
	swsrpstat	Status
Web Console	Backup Operation screen	Status
	Replication Operation screen	Status

- The following error is displayed and the operation of the Command Line Interface or Web Console ends.

Executed Operation	Errors that Occurs
Backup management	swst0603, swst0634, swst0719, swst0740
Replication management	swsrp0603, swsrp0634, swsrp0719, swsrp0740, swsrp2606, swsrp2648

When the above-mentioned situation has occurred, confirm the device configuration of the virtual machine and confirm if whether the type of disk has been changed.

When the type of disk has been changed, reconfigure the configuration information of the AdvancedCopy Manager.

When the type of disk has not been changed, there is the possibility that a hardware failure has occurred.

Confirm the status of the hardware, and when a hardware failure has occurred, re-execute the operations after removing the cause of the error. Refer to the following sections for details.

- ["8.3 Troubleshooting: Hardware Error During Backup"](#)
- ["8.4 Troubleshooting: Hardware or Other Error During Replication"](#)

10.1.2 Notes on 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. Accordingly, snapshot backup is disabled if:

- All the specified number of generations have been backed up, and
- There is no new volume that can be used as a backup volume.

Simultaneous Backups

Performing snapshot high-speed backup simultaneously with the same transaction volume specified performs two backups in parallel. However, performing snapshot high-speed backup simultaneously more often than the number of preservation generations cancels the backup operations starting from the oldest one. In other words, you cannot perform more simultaneous backups than the specified number of generations.

Notes on Combining Cascade Copy and Concurrent OPC in Operation

If you specify a volume used in cascade copy as the copy source and then execute a Concurrent OPC by the `srsrpmake` command, it may not be possible to preserve the sequence of data updates across multiple volumes. In this case, an error message is displayed and the command terminates in an error.

If the error cannot be resolved by responding to the message, we recommend you change the operating method of Advanced Copy or reconfigure the system.

If you cannot change the operating method of Advanced Copy or reconfigure the system, create a file (`check.ini`) that sets the sequence of data updates across multiple volumes. Then, disable sequence preservation before you retry Concurrent OPC. The sequence of data updates is not preserved across multiple volumes.

Create `check.ini` in the following path of the server on which to execute the `srsrpmake` command:

- In a non-cluster environment

```
/etc/opt/FJSVswsrp/data/DEFAULT
```

- In a cluster environment

```
/etc/opt/FJSVswsrp/logicalNodeName/data/DEFAULT
```

Table 10.3 File Format of `check.ini`

Section Name	Key Name	Value
[check]	ConcurOPCCascade	Set preservation of the sequence of data updates across multiple volumes. By setting "No", preservation of the sequence is disabled. By setting any value other than "No", preservation of the sequence is enabled.

[Example for disabling sequence preservation]

```
[check]  
ConcurOPCCascade=No
```

Point

For any of the below ETERNUS Disk storage systems, Concurrent OPC can be executed regardless of sequence preservation of data updates. However, the sequence of data updates is not preserved across multiple volumes.

- ETERNUS DX80 S2 (for earlier than V10L30 firmware)
- ETERNUS DX90 S2 (for earlier than V10L30 firmware)
- ETERNUS DX400/DX400 S2 series (for earlier than V10L30 firmware)

- ETERNUS DX8000/DX8000 S2 series (for earlier than V10L30 firmware)

10.1.3 Notes on Synchronous High-speed Backup

Processing of Backup

The `swstbackup` command cannot be executed until the transaction and backup volumes are in the equivalency maintain state.

Backup Policy Settings


A backup policy can still be set up if the number of registered volumes is enough for snapshot type high-speed backup, even if the number is not enough for synchronous high-speed backup. In this case, synchronous high-speed backup cannot be performed.

10.1.4 Notes on Restoration

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 is overwritten.

Figure 10.1 Notes on Restoration of AdvancedCopy Manager



TD: Transaction volume
BD: Backup volume

10.1.5 Notes on SDX Object Operations

10.1.5.1 Notes on SDX Object Operations (Backup Management)

For `swstdevinforeset` Command

Note Following Points about Operation with SDX Object in Units of Logical Volumes:

- An SDX object cannot be set as a transaction volume for operation of an SP-series product.
- A volume in a disk group to which the transaction volume belongs cannot be registered as a backup volume.
- A volume in a single disk to which the transaction volume belongs cannot be registered as a backup volume.

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

For **swststartsync** Command

Note Following Points about Operation with SDX Object in Units of Logical Volumes:

Backup synchronous processing of a PRIMECLUSTER GD transaction volume cannot be started in the following circumstances:
In the case of any of the above, you should take appropriate action according to the [System Administrator Corrective Action] for the error message as displayed.

- The volume structure of the transaction volume is changed.
- The transaction volume is set as a proxy object.
- A backup volume cannot be stopped.
- The slices that make up the transaction volume include a temporarily separated slice or a slice being copied.

Backup synchronous processing with EC cannot be started in the following circumstances:

In these cases, only software-copy with PRIMECLUSTER GD is available. To use the software-copy function of PRIMECLUSTER GD, specify the -Xgds-softcopy option.

- An SDX object that does not support the EC function is specified.
- All available backup volumes are in a mirroring structure.
- The transaction volume belongs to a mirror group connected to any subordinate group.
- An SDX object is in the root class.

If a backup volume satisfies any of the following conditions, it is not selected as the backup destination because it cannot be used as a proxy object of the transaction volume. Even if it is specified as the backup destination, it cannot be used.

- The backup volume is not the same class as the transaction volume.
- The slices that make up the backup volume include a temporarily separated slice or a slice that is being copied.
- The backup volume is set as the master object associated with another proxy of PRIMECLUSTER GD.
- The backup volume is mounted.
- The backup volume is a logical volume whose size is different from that of the transaction volume.
- The backup volume is in a mirroring structure (excluding cases where -Xgds-softcopy or -Xgds-selectcopy option is specified).
- No disk is connected directly to the backup volume (excluding cases where -Xgds-softcopy or -Xgds-selectcopy option is specified).

For **swstcancelsync** Command

Note Following Points about Use with SDX Object in Units of Logical Volumes:

The backup volume enters the invalid data state (INVALID) if synchronous processing is cancelled in the following circumstances during the backup of an SDX object:

- When the synchronous processing is not in the equivalency maintenance state or copy established state

Restart synchronous processing or snapshot processing again, and the volume should enter the stop state (STOP), which automatically clears the data invalid state.

To use an alternative method to restore the state, refer to "Errors Related to Volume States" in the *PRIMECLUSTER Global Disk Services User's Guide (Solaris(TM) Operating System)*, and perform appropriate action.

For **swstbackup** Command

Note Following Points about Use with SDX Object in Units of Logical Volumes:

Backup of a PRIMECLUSTER GD transaction volume cannot be performed in the following circumstances:

In the case of any of the above, you should take appropriate action according to the [System Administrator Corrective Action] for the error message as displayed.

- The volume structure of the transaction volume is changed.
- The transaction volume is set as a proxy object.
- A backup volume cannot be stopped.
- The slices that make up the transaction volume include a temporarily separated slice or a slice being copied.
- An SDX object that does not support the OPC function is specified.
- All available backup volumes are in a mirroring structure.
- The transaction volume belongs to a mirror group connected to any subordinate group.
- An SDX object is in the root class.

If a backup volume satisfies any of the following conditions, it is not selected as the backup destination because it cannot be used as a proxy object of the transaction volume even if it is specified as the backup destination.

- The backup volume is not the same class as the transaction volume.
- The slices that make up the backup volume include a temporarily separated slice or a slice being copied.
- The backup volume is set as the master object associated with another proxy of PRIMECLUSTER GD.
- The backup volume is mounted.
- The backup volume is a logical volume whose size is different from that of the transaction volume.
- The backup volume is in a mirroring structure.
- No disk is connected directly to the backup volume.

Backup with the output destination backup volume specified cannot be performed in the following circumstances:

- An ordinary volume is specified as the transaction volume and a PRIMECLUSTER GD logical volume as the backup volume.
- A PRIMECLUSTER GD logical volume is specified as the transaction volume and an ordinary volume as the backup volume.

Note Following Points about Operation with SDX Object in Units of Slices:

Backup cannot be performed in the following circumstances:

In the case of any of the above, you should perform appropriate action according to the [System Administrator Corrective Action] for the error message as displayed.

- An SDX object does not satisfy the following conditions:
- The state of the mirror volume is ACTIVE.
- The state of the SDX disk is ENABLE.
- The state of the mirror slice is ACTIVE or TEMP.

Pre-processing and post-processing scripts cannot be executed in the following circumstances

- The state of the SDX object slice is TEMP.

For **swstorestore** Command

Note Following Points about Operation with SDX Object in Units of Logical Volumes:

A PRIMECLUSTER GD transaction volume cannot be restored in the following circumstances:

In the case of any of the above, you should take appropriate action according to the [System Administrator Corrective Action] for the error message as displayed.

- In restoration to a transaction volume, the backup volume is not the proxy object of the transaction volume.
- In restoration to a transaction volume, the transaction volume cannot be stopped.
- In restoration to a transaction volume, the target backup volume cannot be stopped.

Restoration with OPC cannot be performed for a PRIMECLUSTER GD transaction volume in the following circumstances: In these cases, only software-copy is available. To use the software-copy function of PRIMECLUSTER GD, specify the `-Xgds-softcopy` option.

- A PRIMECLUSTER GD logical volume that does not support the OPC function is specified.
- The backup volume belongs to a mirror group connected to any subordinate group.
- An SDX object is in the root class.

If the `swstrestore` command using the software-copy function of PRIMECLUSTER GD is suspended, such as for a forced termination, the started software-copy processing continues. In this case, wait for the software-copy processing to finish, and then execute the `swstsrsemitch` command.

Note Following Points about Operation with SDX Object in Units of Slices:

If the transaction volume is a slice of an SDX object and it is not a single volume, restoration using this command cannot be performed. This is because Advanced Copy cannot be performed on a mirror volume. For restoration to the transaction volume of an SDX object, follow the procedure below.

1. Confirm the backup volume by using the backup history list view or `swsthistdisp` command.
2. For restoration, use the "dd" command to copy data to the volume.

10.1.5.2 Notes on SDX Object Operations (Replication Management)

For `swsrpsetvol` Command

Note Following Points about Operation with SDX Object in Units of Logical Volumes:

Replication volume information for the replication operation of an SDX object cannot be set in the following circumstances:

- An SDX object and slice, or an SDX object and a volume of another OS are specified.
- A logical volume and the `n` option are specified.
- The original volume and replica volume belong to different classes.
- The original volume and replica volume belong to the same group; alternatively, in operation of single volumes, they belong to the same single disk.

For `swsrpdelvol` Command

Note Following Points about Operation with SDX Object in Units of Logical Volumes:

If the original volume and replica volume are bound, or if they are separated when a copy session exists, the replication volume information cannot be deleted.

For `swsrpstartsync` Command

Note Following Points about Operation with SDX Object in Units of Logical Volumes:

If any of the following is true, the `swsrpstartsync` command in the replication operation of an SDX object cannot be executed:

- The copy source volume and copy destination volume are in the master-proxy relationship when synchronous processing is started.
- The copy source volume and copy destination volume have not been separated from their master-proxy relationship when synchronous processing is resumed.
- The copy source volume and copy destination volume belong to different classes when synchronous processing is started.
- The copy source volume and copy destination volume belong to the same disk group or single disk.

- The copy source volume satisfies any of the following conditions:
 - The volume contains any temporarily separated slice or slice that is being copied.
 - The copy source volume is a proxy object when synchronous processing is started.
 - Copy processing with another volume is in progress.
- The copy destination volume satisfies any of the following conditions:
 - The volume is in a mirroring structure and the -Xgds-softcopy or -Xgds-selectcopy option is not specified when synchronous processing is started. The volume cannot be stopped.
 - The volume is a master volume when synchronous processing is started.
 - The volume contains any temporarily separated slice or slice that is being copied.
 - The volume is a proxy of another master.
- The hardware copy function is not available (excluding cases when the -Xgds-softcopy or -Xgds-selectcopy option is specified). The typical causes are:
 - EC/REC of the ETERNUS Disk storage system is not properly set.
 - The copy source volume belongs to a mirror group connected to any subordinate group.
 - The copy destination volume is in a mirroring structure.
 - The volume is in the root class.

In the replication operation of an SDX object registered with the replication function, do not use the "sdxproxy" command directly to change the state of the SDX object. Otherwise, the "sdxproxy" command may create an inconsistency between the management information of the replication function and the actual object state. In this case, use the [swsrpcancel](#) command to cancel the replication operation.

Note Following Points about Operation with SDX Object in Units of Slices:

Note the following for the replication operation of an SDX object:

Replication cannot be started in the following cases:

- The copy destination volume is an SDX object and the status of the SDX object is not as follows:
 - Mirror volume is either ACTIVE or STOP
 - Mirror slice status is either ACTIVE or STOP
- If the copy destination volume is an SDX object that is not a single volume, replication using this command cannot be performed. Therefore, if either the original volume or replica volume is an SDX object that is not a single volume, you can perform replication only from an SDX object to a non-SDX object.
- To copy data to a volume that is an SDX object but not a single volume, use the dd command.

For [swsrpmake](#) Command

Note Following Points about Operation with SDX Object in Units of Logical Volumes:

If any of the following is true, replication in the replication operation of an SDX object cannot be created:

- The copy source volume and copy destination volume are in the master-proxy relationship in snapshot replication.
- The copy source volume and copy destination volume belong to different classes in snapshot replication.
- The copy source volume and copy destination volume belong to the same disk group or a single disk in snapshot replication.
- The copy source volume satisfies any of the following conditions:
 - The volume contains any temporarily separated slice or slice that is being copied, in snapshot replication.
 - The copy source volume is a proxy object in snapshot replication.
 - Copy processing with another volume is in progress.

- The copy destination volume satisfies any of the following conditions:
 - The volume is in a mirroring structure in snapshot replication.
 - The volume could not be stopped in snapshot replication.
 - The volume is a master volume in snapshot replication.
 - The volume contains any temporarily separated slice or slice that is being copied, in snapshot replication.
 - The volume is a proxy of another master in snapshot replication.
- The OPC function is not available for snapshot replication. The typical causes are:
 - OPC of the disk array unit is not properly set.
 - EC/REC in the equivalency maintenance state or suspended state is between the specified copy source volume and another proxy. However, Advanced Copy can be set in units of mirror slices, which makes it possible to set Advanced Copy for as many copy destination volumes as the number of mirror slices of the copy source volume at the same time. Therefore, this error occurs as the result of an attempt to set Advanced Copy for a greater number of copy source volumes than the number of mirror slices.
 - The copy source volume belongs to a mirror group connected only to some subordinate groups but not directly to a disk. The copy source volume belongs to a mirror group connected to any subordinate group.
 - The copy destination volume is in a mirroring structure.
 - The volume is in the root class.

In the replication operation of an SDX object registered with the replication function, do not use the "sdxproxy" command directly to change the state of the SDX object. Otherwise, the "sdxproxy" command may create an inconsistency between the management information of the replication function and the actual object state. In this case, use the [swsrpcancel](#) command to cancel the replication operation.

Note Following Points about Operation with SDX Object in Units of Slices:

Replication cannot be performed in the following circumstances:

- The copy source volume is an SDX object for replication of an SDX object, and the SDX object is not in the following states:
 - The state of the mirror volume is either ACTIVE or STOP.
 - The state of the mirror slice is either ACTIVE or TEMP.
 - The state of the SDX disk is ENABLE
- The copy destination volume is an SDX object, and the SDX object is not in the following states:
 - The state of the mirror volume is either ACTIVE or STOP.
 - The state of the mirror slice is either ACTIVE or STOP.
- If the copy destination volume is an SDX object that is not a single volume for the replication operation for an SDX object, replication using this command cannot be performed. Therefore, if either the original volume or replica volume is an SDX object that is not a single volume, you can perform replication only from an SDX object to a non-SDX object.
- To copy data to a volume that is an SDX object but not a single volume, use the dd command.

Preprocessing of the copy source volume is not performed in the following circumstances.

- The state of the slice of the SDX object is TEMP.

For [swsrpcancel](#) Command

Note Following Points about Operation with SDX Object in Units of Logical Volumes:

If any of the following is true, replication in the replication operation of an SDX object cannot be cancelled:

- The pair of the target replication volume used for replication cancellation is in the equivalency maintenance state, and copy processing between the copy source volume and another volume is in progress.

If replication cancellation is executed at either of the following times during the replication operation of an SDX object, the copy destination volume enters the invalid data state (INVALID):

- When snapshot replication is neither in the equivalency maintenance state nor copy established state
- When data copy processing of snapshot processing is in progress

Restart synchronous processing or snapshot processing, and the volume enters the stop state (STOP), which automatically clears the data invalid state.

To restore the state using an alternative method, refer to "Errors Related to Volume Statuses" in the *PRIMECLUSTER Global Disk Services User's Guide (Solaris(TM) Operating Environment)*, and perform appropriate action.

Note Following Points about Operation with SDX Object in Units of Slices:

Synchronous processing in the equivalency maintenance state cannot be stopped in the following circumstances:

- The copy source volume is an SDX object, and the SDX object is in any of the following states:
 - The state of the mirror volume is either ACTIVE or STOP.
 - The state of the mirror slice is either ACTIVE or TEMP.
 - The state of the SDX disk is ENABLE.

Note: If the mirror slice is in the TEMP state, pre-processing of the original volume is not performed.

- The file system is mounted in an environment where the copy source volume or copy destination volume makes up a SynfinityFILE of a multi-volume structure. In such cases, unmount the file system, and process all of the volumes that make up the SynfinityFILE.

10.1.6 Notes on Cluster Operation

Note the following general points on cluster operation:

- Do not set a SWSTGNODE registration that configures an environment variable to the entire system.
- AdvancedCopy Manager does not monitor daemons of an active server. If the daemon of a transaction stops for some reason, AdvancedCopy Manager does not automatically restart it, so you must restart it manually. Moreover, the stopping of an active server daemon of AdvancedCopy Manager does not cause the cluster to fail over. For details on the daemons of AdvancedCopy Manager active servers and how to start these daemons, refer to "[Chapter 2 Starting and Stopping Daemons](#)".
- For volumes for which to perform backup and replication, configure their device names, mount points and logical unit numbers on the ETERNUS Disk storage system so that they are the same on all the nodes making up a cluster.

For a Managed Server daemon (or Management Server daemon) operating together with a daemon in a scalable operation, note the following points that do not apply to regular operation or cluster operation:

- For a Management Server transaction, note that the logical IP address you specify in the cluster setup command must be the logical IP address of the Management Server transaction. Specifying the physical IP address or logical IP address of a server in scalable operation results in incorrect setting.
- Do not register any cluster resource of a transaction in scalable operation as a cluster resource of a Managed Server transaction (Management Server transaction).
- When specifying the Managed Server's IP address for the execution of AdvancedCopy Manager commands, specify the logical IP address for the Managed Server transactions. Specifying the physical IP address or logical IP address of a server in scalable operation results in an incorrect setting.
- When specifying the Managed Server's port number for the execution of AdvancedCopy Manager commands, specify the port number for the Managed Server transaction communication daemon (the name of the stgxfws logical node). Specifying a port number besides those applicable to server transaction communication daemons results in an incorrect setting.
- Add backup-use volume with the transaction volume for scalable operation transactions which operate backup and replication.
- You can make a backup or replication on a node operating a Managed Server transaction (Management Server transaction) if it is the active node. You cannot make a backup or replication on any standby node of a Managed Server transaction (Management Server transaction).

- If a transaction (in scalable operation) running on a node operating a Managed Server transaction (Management Server transaction) stops for some reason, you must manually switch the Managed Server transaction (Management Server transaction) over to a node running in scalable operation. The Managed Server transactions (Management Server transactions) are not linked with scalable operation transactions, thus stopping scalable operation transactions does not cause a Managed Server transaction (or Management Server transaction) to fail over.

10.1.7 Notes on Backup in Clustered System

SynfinityCLUSTER or PRIMECLUSTER

For backup operation in a clustered system, note the following points that do not apply to regular operation:

- If failover occurs while running a backup command, you must secure consistency by using the `swstrsemitch` command as a recovery measure.
- Do not register the backup volume in the cluster application when you use the synchronized high-speed backup. After synchronized processing begins, the cluster application occasionally does the failover.
- You must perform backup operation on the active server of the cluster service (userApplication) because AdvancedCopy Manager is placed in the cluster service (userApplication) and operates as part of said cluster service (userApplication). You cannot make a backup from a standby node or another cluster service (userApplication).

SynfinityCLUSTER

Note the following points if a file system has been created on a transaction volume and the file system has been registered as a cluster resource:

- Change the backup pre-processing so that the transaction volume is not unmounted. Refer to "[A.2 Pre-processing and Post-processing of Backup](#)" for information about changing backup pre-processing.
- It is not possible to restore data using the `swstrestore` command during transaction operations. To restore data on the volume, either copy the data using the `dd` command, or mount the backup volume and copy the necessary data. Alternatively, perform the procedure described in "[10.1.7.1 Backup and Restoration During Stoppage of Cluster](#)".

PRIMECLUSTER

If a file system has been created on a transaction volume, and the file system has been registered as a cluster resource, use the following procedure to back up and restore data.

Follow the procedure below to perform backup/restore.

Snapshot-type Backup/Restore Procedure Using AdvancedCopy Manager in PRIMECLUSTER Environment

1. Switch userApplication to maintenance mode.

```
# hvutil -m on userApplication
```

2. Unmount the transaction volume (this is mandatory to backup or restore a transaction volume in the file system).
3. If the transaction volume was unmounted, then make sure that the cluster resources of the transaction volume are Offline (wait until the return value of the command is 0):

```
# hvassert -s <transaction_volume> <cluster_resource_name> Offline
```

4. Perform backup or restore

```
# /opt/FJVSswsts/bin/swstbackup Device-Name
Device-Name swstbackup completed
```

or

```
# /opt/FJVSswsts/bin/swstrestore Device-Name
Device-Name swstrestore completed
```

5. If the transaction volume was unmounted, then mount it back.
6. Make sure that the cluster resources of the transaction volume are Online (wait until the return value of the command is 0):

```
# hvassert -s <transaction_volume> <cluster_resource_name> Online
```

7. Stop the maintenance mode for the userApplication.

```
# hvutil -m off userApplication
```

Synchronous Backup Procedure Using AdvancedCopy Manager in PRIMECLUSTER Environment

1. Start synchronization processing.

```
# /opt/FJSVswsts/bin/swststartsync Device-Name
Device-Name swststartsync completed
```

2. Check the equivalency maintenance status.
3. Switch userApplication to maintenance mode.

```
# hvutil -m on userApplication
```

4. Unmount the transaction volume (this is mandatory to backup/restore a transaction volume in the file system).
5. If the transaction volume was unmounted , then make sure that the cluster resources of the transaction volume are Offline (wait until the return value of the command is 0):

```
# hvassert -s <transaction_volume> <cluster_resource_name> Offline
```

6. Perform backup.

```
# /opt/FJSVswsts/bin/swstbackup Device-Name
Device-Name swstbackup completed
```

7. If the transaction volume was unmounted, then mount it back.
8. Make sure that the cluster resources of the transaction volume are Online (wait until the return value of the command is 0):

```
# hvassert -s <transaction_volume_cluster_resource_name> Online
```

9. Stop maintenance mode for userApplication.

```
# hvutil -m off userApplication
```



Note

- In either procedure, if the transaction volume is unmounted, then the pre-processing and post-processing scripts are not executed (they are only executed when the transaction or replication source volumes are mounted). Take this into account if the applicable parts of the pre-processing and post-processing scripts are being customized.
- The cluster detector detects the status of cluster resources either periodically or asynchronously. Therefore, it may take a few moments for the cluster resource status to be reflected, depending on the system load, and on the timing of mounting/unmounting the transaction volume. This means that the status may not be able to be checked using the hvassert command. Due to this fact, note the following about the timing of using hvassert to check the cluster resource status:
 1. Once the mounting/unmounting of the file system of a cluster FSystem resource has finished, only use hvassert after the time elapsed is equal to or greater than the cluster resource monitoring interval. You can specify the maximum time to wait until the cluster reaches the desired status by using hvassert with the -w option:

```
# hvassert -w <replication_vol> Offline|Online <max_wait_time(secs)>
```

2. When using hvassert with the -w option, the system load and cluster resource reflection processing time must also be taken into account.

Due to that, make sure to specify a value that suits the system operating environment, purpose, and requirements.

In PRIMECLUSTER 4.2, the default monitoring interval for the PRIMECLUSTER Fsystem resource is 10 seconds - check the default value for your version.

For details on the hvassert command, refer to "hvassert(1M)" in the *PRIMECLUSTER Installation and Administration Guide*.

VERITAS Cluster Server

For backup operation in a clustered system, note the following points that do not apply to regular operation:

- If failover occurs while running a backup command, you must ensure consistency by using the [swtsrsemitch](#) command as a recovery measure.
- To use synchronous high-speed backup, do not register a backup to any cluster transaction.
- You must perform backups on the active server of the service group because AdvancedCopy Manager is placed into the service group and operates as a part of it.
You cannot make a backup from a standby node or another service group.

- To perform backup/restoration using AdvancedCopy Manager, you usually must unmount transaction volumes beforehand and mount the volumes later.

If the mount point of transaction volumes have been registered to a cluster transaction, set the mount point resource online or offline from the cluster transaction instead of mounting or unmounting the volumes. Alternatively, modify the backup/restoration pre-processing and post-processing scripts so that the transaction volumes are set offline or online instead of being unmounted or mounted.

Moreover, the Volume is immediately placed offline/online for a mount point resource whereas there is a time delay when mount/unmount is used. Therefore, add into the pre-processing script, after the statement where the success or failure value of offline/online is returned, the `df` or `sleep` command to function until mount/unmount is actually complete.

For details and examples on the pre-processing and post-processing scripts, refer to "[Appendix A Pre-processing and Post-processing of Backup and Restoration](#)". For details on the commands used to set resources offline and online, refer to the relevant cluster software manual.

10.1.7.1 Backup and Restoration During Stoppage of Cluster

SynfinityCLUSTER or PRIMECLUSTER

When a cluster service (userApplication) to which AdvancedCopy Manager belongs is operating, backup operation is enabled only on the active node. The standby node environment is not ready (to provide necessary resources) for backup operation, thus disabling backup operation. Likewise, if the cluster service (userApplication) has been stopped, the environment is not ready and disables backup operation.

However, only when the cluster service (userApplication) has been stopped is backup operation enabled by temporarily arranging the minimum essentials.

Note

- When transactions in scalable operation are handled together, only stop cluster service (userApplication) in scalable operation; do not stop cluster service (userApplication) for Managed Server (or Management Server) transactions.
Enabling backup operation disks (transaction volume and backup volume to operate) allows you to perform backup operation in the same way as for regular operations.
If a backup operation disk cannot be enabled, you cannot perform backup operation while user applications in scalable operation are stopped.
- If any of the following types of shared disks cannot be enabled (i.e., set online or mounted in a system), backup operation is disabled:
 - Shared disk for AdvancedCopy Manager shared data

- Backup operation disk (transaction volume/backup volume to operate)
- The following operations are not allowed in backup operation:
 - Backup/restoration through Web Console
 - Manipulation for backup operation from a Management Server (host specification with the -h option)
 - Addition or deletion of a transaction volume/backup volume, or change in policy
- To execute a display command for backup management to a Management Server transaction that also functions as a Managed Server transaction, you must specify options for the command. For details on the display command for backup management, refer to "9.1 Backup Management Commands".

Perform backup operation while a cluster service (userApplication) is stopped by taking the following steps:

1. Confirm that the cluster service (userApplication) has been stopped on both nodes.
For details on how to stop a cluster service (userApplication), refer to the PRIMECLUSTER manual.
2. Log in to the node on which you want to make a backup through telnet, for example.
You cannot use the logical IP address. Directly use a node that is operable by using the physical IP address.
3. Enable the shared disk.
Start (set online) the shared disk.
Mount the shared disk for AdvancedCopy Manager shared data.
When the transaction volume is a file system, mount it.



Note

Be sure to enable the shared disk from either node. Do not enable the shared disk from both nodes.

4. In the environment variable SWSTGNODE, set the logical node name for the relevant task.
Set the environment variable as shown below.

<Execution example>

```
# SWSTGNODE=logicalNodeName
# export SWSTGNODE
```

5. Perform backup operation.
You can perform backup operation in the same way as for regular operations for cluster operation.
You can use the [swstrestore](#) command to perform restoration.
6. Release all the shared disks you enabled in step 3.
Unmount the file systems you mounted.
Stop (set offline) the shared disk.
7. Start (set online) the cluster service (userApplication).
Start the cluster service (userApplication), if necessary.
For details on how to start a cluster service (userApplication), refer to the PRIMECLUSTER manual.

VERITAS Cluster Server

When a service group to which AdvancedCopy Manager belongs is operating, backup is enabled only on the active node. The standby node environment is not ready (to provide necessary resources) for backup operation, therefore backup is disabled. Likewise, if the service group has been stopped, the environment is not ready and disables backup operation. However, backup is temporarily enabled when the service group is stopped.

Note

- If any of the following types of shared disks cannot be enabled (i.e., set online or mounted in a system), backup is disabled:
 - Shared disk for AdvancedCopy Manager shared data
 - Backup operation disk (transaction volume/backup volume to operate)
- The following operations are not allowed in backup:
 - Backup/restoration through Web Console
 - Manipulation for backup operation from a Management Server (host specification with the -h option)
 - Addition or deletion of a transaction volume/backup volume, or change in policy
- To execute a display command for backup management to a Management Server transaction that also functions as a Managed Server transaction, you must specify options for the command. For details on the display command for backup management, refer to "9.1 Backup Management Commands".

If the service group has been stopped, perform a backup by taking the following steps:

1. Confirm that the service group has been stopped on both nodes.
For details on how to stop a service group, refer to the VERITAS Cluster Server manual.
2. Log in (for example, using Telnet) to the node on which you want to make a backup.
You cannot use the logical IP address. Use a node that is operable via the physical IP address.
3. Enable the shared disk.
Start (set online) the shared disk.
Mount the shared disk for AdvancedCopy Manager shared data.
When the transaction volume is a file system, mount it.

Note

Be sure to enable the shared disk from either node. Do not enable the shared disk from both nodes.

4. In the environment variable SWSTGNODE, set the logical node name for the relevant task.
Set the environment variable as shown below.

<Execution example>

```
# SWSTGNODE=logicalNodeName
# export SWSTGNODE
```

5. Perform the backup.
You can perform the backup in the same way as for regular operations in a cluster environment.
You can use the [swstrestore](#) command to perform restoration.
6. Release all the shared disks you enabled in step 3.
unmount the file systems you mounted.
Stop (set offline) the shared disk.
7. Start (set online) the service group.
Start the service group, if necessary.
For details on how to start a service group, refer to the VERITAS Cluster Server manual.

10.1.8 Notes on Replication in Clustered System

SynfinityCLUSTER or PRIMECLUSTER

For replication operation in a clustered system, note the following points that do not apply to regular operation:

- If failover occurs while running a replication command, you must secure consistency by using the [swsrprecoverres](#) command as a recovery measure.
- Do not register the copy destination volume in the cluster application when you use a synchronized replication. After synchronized processing begins, the cluster application occasionally does the failover.
- You must perform replication operation on the active system of the cluster service (userApplication) because AdvancedCopy Manager is placed in the cluster service (userApplication) and operates as part of said cluster service (userApplication). You cannot perform replication from a standby node or another cluster service (userApplication).

SynfinityCLUSTER

Note the following points if a file system has been created on a copy source volume and the file system has been registered as a clustered system resource:

- Change the pre-processing and post-processing for replication so that the copy source volume is not unmounted.
- Do not register the copy destination volume in the cluster application.
- Data cannot be restored from the copy destination volume to the copy source volume using the [swsrpmake](#) command during transaction operations. To restore data on the volume, either copy the data using the dd command, or mount the copy destination volume and copy the necessary data. Alternatively, perform the procedure described in "[10.1.8.1 Replication During Stoppage of Cluster](#)".

PRIMECLUSTER

Note the following points if a file system has been created on a copy source volume and the file system has been registered as a clustered system resource:

- Do not register the copy destination volume in the cluster application.

Follow the procedure below to perform replication.

Snapshot-type Replication Procedure Using AdvancedCopy Manager in PRIMECLUSTER Environment

1. Switch userApplication to maintenance mode.

```
# hvutil -m on userApplication
```

2. Unmount the replication source volume (this is mandatory to replicate a replication source volume in the file system).
3. If the replication source volume was unmounted, then make sure that the cluster resources of the volume are Offline (wait until the return value of the command is 0):

```
# hvassert -s <replication_source_volume> <cluster_resource_name> Offline
```

4. Perform replication.

```
# /opt/FJSVswsrp/bin/swsrpmake <from_volume_name> <to_volume_name>
FROM=<from_volume_name> TO=<to_volume_name> swsrpmake completed
```

5. If the replication source volume was unmounted, then mount it back.
6. Make sure that the cluster resources of the replication source volume are Online (wait until the return value of the command is 0):

```
# hvassert -s <replication_source_volume> <cluster_resource_name> Online
```

7. Stop maintenance mode for userApplication.

```
# hvutil -m off userApplication
```

Synchronous Replication Procedure Using AdvancedCopy Manager in PRIMECLUSTER Environment

1. Start synchronization processing.

```
# /opt/FJSVswsrp/bin/swsrpstartsync <from_volume_name> <to_volume_name>  
FROM=<from_volume_name> TO=<to_volume_name> swsrpstartsync completed
```

2. Check the equivalency maintenance status.
3. Switch userApplication to maintenance mode.

```
# hvutil -m on userApplication
```

4. Unmount the replication source volume (this is mandatory to replicate a replication source volume in the file system).
5. If the replication source volume was unmounted, then make sure that its cluster resources are Offline (wait until the return value of the command is 0):

```
# hvassert -s <replication_source_volume> <cluster_resource_name> Offline
```

6. Perform replication.

```
# /opt/FJSVswsrp/bin/swsrpmake <from_volume_name> <to_volume_name>  
FROM=<from_volume_name> TO=<to_volume_name> swsrpmake completed
```

7. If the replication source volume was unmounted, then mount it back.
8. Make sure that the cluster resources of the replication source volume are Online (wait until the return value of the command is 0):

```
# hvassert -s <replication_source_volume> <cluster_resource_name> Online
```

9. Stop maintenance mode for userApplication.

```
# hvutil -m off userApplication
```

Note

- In either procedure, if the replication source volume was unmounted, then the pre- and post-processing scripts are not executed (they are only executed when the transaction or replication source volumes are mounted).
If the applicable parts of the pre- and post-processing scripts are being customized, take this point into account.
- The cluster detector detects the status of cluster resources either periodically or asynchronously. Therefore, it may take a few moments for the cluster resource status to be reflected, depending on the system load, and on the timing of mounting/unmounting the replication source volume. This means that the status may not be able to be checked using the hvassert command.
Due to this fact, note the following about the timing of using hvassert to check the cluster resource status:

1. Once the mounting/unmounting of the file system of a cluster FSystem resource has finished, only use hvassert after the time elapsed is equal to or greater than the cluster resource monitoring interval.

You can specify the maximum time to wait until the cluster reaches the desired status by using hvassert with the -w option:

```
# hvassert -w <replic_src_vol> <cluster_res_name> Offline|Online  
<max_wait_time(secs)>
```

2. When using hvassert with the -w option, the system load and cluster resource reflection processing time must also be taken into account.
Due to that, make sure to specify a value that suits the system operating environment, purpose, and requirements.

In PRIMECLUSTER 4.2, the default monitoring interval for the PRIMECLUSTER Fsystem resource is 10 seconds - check the default value for your version.

For details on the hvassert command, refer to "hvassert(1M)" in the *PRIMECLUSTER Installation and Administration Guide*.

VERITAS Cluster Server

For replication in a clustered system, note the following points that do not apply to regular operation:

- If failover occurs while running a replication command, you must ensure consistency by using the [swsrprecoverres](#) command as a recovery measure.
- When you use synchronous replication, do not register any destination volumes to a cluster transaction.
- You must perform replication on the active system of the service groups because AdvancedCopy Manager is placed in the service groups and operates as a part of it.
You cannot perform replication from a standby node or another service group.
- To perform replication, you must unmount original/source volumes beforehand and mount the volumes later for some commands. If the mount points of original/source volumes have been registered to a cluster transaction, set the mount point resources online or offline from the cluster transaction instead of mounting or unmounting the volumes. Alternatively, modify the unmount and mount processes in the replication pre-processing and post-processing scripts to offline and online processes, respectively. Moreover, the Volume is immediately placed offline/online for a mount point resource whereas there is a time delay when mount/unmount is used. Therefore, add into the pre-processing script, after the statement where the success or failure value of offline/online is returned, the df or sleep command to function until mount/unmount is actually complete.
For details and examples on the types of commands that require the volumes to be unmounted and mounted, and the pre-processing and post-processing scripts, refer to "[Appendix C Pre-processing and Post-processing of Replication](#)". For details on the commands used to set resources online and offline, refer to the relevant clustered system manual.

10.1.8.1 Replication During Stoppage of Cluster

SynfinityCLUSTER or PRIMECLUSTER

When a cluster service (userApplication) to which AdvancedCopy Manager belongs is operating, replication operation is enabled only on the active node. The standby node environment is not ready (to provide necessary resources) for replication operation, thus disabling replication operation. Likewise, if the cluster service (userApplication) has been stopped, the environment is not ready and disables replication operation.

However, only when the cluster service (userApplication) has been stopped is replication operation enabled by temporarily arranging the minimum essentials.



-
- When transactions in scalable operation are handled together, only stop cluster service (userApplications) in scalable operation; do not stop cluster service (userApplications) for Management Server transaction and Managed Server transactions.
Enabling replication operation disks (source volume/destination volume to operate) allows you to perform replication operation in the same way as for regular operations.
If a replication operation disk cannot be enabled, you cannot perform replication operation while a cluster service (userApplications) in scalable operation is stopped.
 - If any of the following types of shared disks cannot be enabled (i.e., set online or mounted in a system), replication operation is disabled.
 - Shared disk for AdvancedCopy Manager shared data
 - Replication operation disks (source volume/destination volume to operate)
 - The following operations are not allowed in replication operation:
 - Manipulation for replication operation from a Management Server (host specification with the -h option)
 - Addition or deletion of an original/replica volume
 - Inter-server replication without the -m option specified

- To execute a display command for replication management to a Management Server transaction that also functions as a Managed Server transaction, you must specify options for the command. For details on the display command for replication management, refer to "9.3 Replication Management Commands".

Perform replication operation during stoppage of a cluster service (userApplication) by taking the following steps:

1. Confirm that the cluster service (userApplication) has been stopped on both nodes.
For details on how to stop a userApplication, refer to the SynfinityCLUSTER or PRIMECLUSTER manual.
2. Log in to the node on which you want to perform replication operation through telnet, for example.
You cannot use the logical IP address. Directly use a node that is operable by using the physical IP address.
3. Enable the shared disk.
Start (set online) the shared disk.
Mount the shared disk for AdvancedCopy Manager shared data.
When the original volume and replica volume have been mounted as file systems for operation, mount the volumes.

Note

Be sure to enable the shared disk from either node. Do not enable the shared disk from both nodes.

4. In the environment variable SWSTGNODE, set the logical node name for the relevant task.
Set the environment variable as shown below.

<Execution example>

```
# SWSTGNODE=logicalNodeName  
# export SWSTGNODE
```

5. Perform replication operation.
You can perform replication operation in the same way as for regular operations for cluster operation.
You can also perform restoration from the replica volume to the original volume.
6. Release all the shared disks you enabled in step 3.
Unmount the file systems you mounted.
Stop (set offline) the shared disk.
7. Start (set online) the cluster service (userApplication).
Start the userApplication, if necessary.
For details on how to start a cluster service (userApplication), refer to the SynfinityCLUSTER or PRIMECLUSTER manual.

VERITAS Cluster Server

When a service group to which AdvancedCopy Manager belongs is operating, replication is enabled only on the active node. The standby node environment is not ready (to provide necessary resources) for replication, and replication is, therefore, disabled. Likewise, if the service group has been stopped, the environment is not ready and disables replication. However, replication is temporarily enabled when the service group has been stopped.

Note

- If any of the following types of shared disks cannot be enabled (i.e., set online or mounted in a system), replication is disabled.
 - Shared disk for AdvancedCopy Manager shared data
 - Replication operation disks (source volume/destination volume to operate)
- The following operations are not allowed in replication:
 - Manipulation for replication from a Management Server (host specification with the -h option)
 - Addition or deletion of an original/replica volume

- Inter-server replication without the -m option specified
- To execute a display command for replication management to a Management Server transaction that also functions as a Managed Server transaction, you must specify options for the command. For details on the display command for replication management, refer to "9.3 Replication Management Commands".

Perform replication during stoppage of the service group by taking the steps below.

1. Confirm that the service group has been stopped on both nodes.
For details on how to stop a service group, refer to the VERITAS Cluster Server manual.
2. Log in (for example, through Telnet) to the node on which you want to perform replication.
You cannot use the logical IP address. Use a node that is operable via the physical IP address.
3. Enable the shared disk.
Start (set online) the shared disk.
Mount the shared disk for AdvancedCopy Manager shared data.
When the original volume and replica volume have been mounted as file systems for operation, mount the volumes.

Note

Be sure to enable the shared disk from either node. Do not enable the shared disk from both nodes.

4. In the environment variable SWSTGNODE, set the logical node name for the relevant task.
Set the environment variable as shown below.

<Execution example>

```
# SWSTGNODE=logicalNodeName
# export SWSTGNODE
```

5. Perform replication.
You can perform replication in the same way as for regular operations in a cluster environment.
You can also perform restoration from the replica volume to the original volume.
6. Release all the shared disks you enabled in step 3.
Unmount the file systems you mounted.
Stop (set offline) the shared disk.
7. Start (set online) the service group.
Start the service group, if necessary.
For details on how to start a service group, refer to the VERITAS Cluster Server manual.


10.1.9 Use in Combination with ODX (Offloaded Data Transfer)

When using an ETERNUS Disk storage system that supports the Offloaded Data Transfer (hereinafter referred to as "ODX") function, ODX session starts within the ETERNUS Disk storage system by Windows Server 2012 or later Offloaded Data Transfer function when files are copied or transferred. However, Advanced Copy cannot be performed by AdvancedCopy Manager on the condition that a volume on which ODX session is running is taken as copy destination. For this reason, if Advanced Copy is performed by AdvancedCopy Manager on the condition that a volume containing the files or areas being copied or transferred by the Offloaded Data Transfer function is taken as copy destination, an error message is displayed and the command may terminate abnormally.

If the command terminates abnormally, check on Web Console that no ODX session exists on the copy destination volume and then re-execute the command.

If the error is not resolved, refer to "If Error Is Not Resolved" in "Use in Combination with ODX (Offloaded Data Transfer)" in the *ETERNUS SF AdvancedCopy Manager Operation Guide (for Windows)*.

Figure 10.2 Use in Combination with ODX



*1: In the following ETERNUS Disk storage systems whose firmware version number is earlier than V10L50, Advanced Copy cannot be performed by AdvancedCopy Manager with respect to the volume on which ODX session is running.


- ETERNUS DX80 S2
- ETERNUS DX90 S2
- ETERNUS DX400/DX400 S2 series
- ETERNUS DX8000/DX8000 S2 series

10.1.10 Use in Combination with XCOPY (Extended Copy)

When copying data by replicating VMware vSphere virtual machines or running Storage vMotion, Extended Copy (hereinafter referred to as "XCOPY") session starts by the Full Copy function of VMware vStorage APIs for Array Integration (VMware VAAI). However, Advanced Copy cannot be performed by AdvancedCopy Manager on the condition that a volume on which XCOPY session is running is taken as copy destination. For this reason, if Advanced Copy is performed by AdvancedCopy Manager on the condition that a volume on which data is being copied by replicating a virtual machine or running Storage vMotion is taken as copy destination, an error message is displayed and the command may terminate abnormally.

If the command terminates abnormally, check on Web Console that no XCOPY session exists on the copy destination volume and then re-execute the command.

Figure 10.3 Use in Combination with XCOPY



*1: In the following ETERNUS Disk storage systems whose firmware version number is earlier than V10L50, Advanced Copy cannot be performed by AdvancedCopy Manager with respect to the volume on which XCOPY session is running.

- ETERNUS DX80 S2
- ETERNUS DX90 S2
- ETERNUS DX400/DX400 S2 series
- ETERNUS DX8000/DX8000 S2 series

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

This appendix provides information on pre-processing and post-processing of backup and restoration.

A.1 Overview

Pre-processing and post-processing of backup or restoration are started before and after backup or restoration when a backup or restore command is executed.

This processing is required by AdvancedCopy Manager to back up or restore a transaction volume.

A.2 Pre-processing and Post-processing of Backup

Backup using AdvancedCopy Manager must be performed basically while access to a transaction volume from other processes is inhibited.

Normally, the pre-processing acquires the mount status of a transaction volume using the transaction volume name and then performs the following processing:

Table A.1 Pre-processing for Transaction Volume

Transaction Volume Status	Pre-processing
Mounted	Unmount a transaction volume.(*1)
Unmounted	Take no action.

*1 If the transaction volume cannot be unmounted, however, customize the shell script for pre-processing is customized accordingly. For information on this customization, refer to the section on "[When You Do Not Want to Unmount Transaction Volume](#)".

Tasks performed in the post-processing are generally determined by what has been done in the pre-processing.

Table A.2 Backup Pre-processing and Post-processing

Pre-processing	Post-processing
A transaction volume was unmounted.	Remount the transaction volume.
Take no action.	Take no action.

If no transaction volume has been mounted since the start of system operation (e.g., a transaction volume used as a database), neither pre- nor post-processing is performed.

If special pre-processing or post-processing is required, you need to add the additional instructions to the shell script.

When customizing a script, strictly observe the following rules regarding error codes:

Table A.3 Error Code Rules

Error Code	Usage
0-99	Unusable (these are reserved for AdvancedCopy Manager)
100-255	Usable

Note

If pre-processing failed, execute `swstrsemth` ([Resource match command](#)) because the consistency of resource information may be incomplete.

A.2.1 Pre-processing of Backup

The name of a script file for pre-processing of a backup is as follows.

For Non-clustered System

```
/etc/opt/FJSVswsts/sh/OpcBackup.pre
```

For Clustered System

```
/etc/opt/FJSVswsts/logicalNodeName/sh/OpcBackup.pre
```

Pre-processing Script for Backup (OpcBackup.pre)

```
1: #!/bin/sh
2:
3: # AdvancedCopy Manager
4: # All Rights Reserved, Copyright FUJITSU LIMITED, 2002-2009
5:
6: #
7: #   Pre-processing of backup processing
8: #
9: #   Argument: $1 Device name of transaction disk
10: #               $2 Mount point of transaction disk
11: #               $3 Device name of backup disk
12: #
13: #   Error number
14: #       2: Argument error
15: #       10: umount error
16:
17:
18: # Argument check
19: case $# in
20: 1)
21:   ;;
22: 2)
23:   ;;
24: 3)
25:   ;;
26: *)
27:   exit 2
28:   ;;
29: esac
30:
31: device=`echo $1`
32: mount_point=`echo $2`
33: bk_device=`echo $3`
34:
35: # Determination postprocessing file name
36:
37: if [ "$SWSTGNODE" != "" ]
38: then
39:   swstg_node="/`echo $SWSTGNODE`"
40: else
41:   swstg_node=""
42: fi
43:
44: err_log_path="/var/opt/FJSVswsts"$swstg_node"/log"
45:
46: if [ "`echo $device | /usr/bin/grep "/dev/dsk/"`" != "" ]
47: then
48:   # /dev/dsk/c?t?d?s? -> c?t?d?s?
49:   dev=`echo $device | /usr/bin/sed "s/\dev\dsk\///"`
50: elif [ "`echo $device | /usr/bin/grep "/dev/FJSV"`" != "" ]
```

```

51: then
52:   # /dev/FJSMphd/dsk/mplb?s? -> /dev/FJSMphd/dsk/mplb?s?
53:   # /dev/FJSMphd/dsk/mphd?s? -> /dev/FJSMphd/dsk/mphd?s?
54:   dev=`echo $device | /usr/bin/cut -d/ -f5`
55: elif [ "`echo $device | /usr/bin/grep "/dev/sfdsk/"`" != "" ]
56: then
57:   if [ "`echo $device | /usr/bin/grep ":"`" != "" ]
58:   then
59:     # /dev/sfdsk/class/dsk/volume:c?t?d? -> class_volume_c?t?d?
60:     dev=`echo $device | /usr/bin/sed "s/\\/dev\\/sfdsk\\/\\/\\/`"
61:     dev=`echo $dev | /usr/bin/sed "s/\\/dsk\\/\\/\\/`"
62:     dev=`echo $dev | /usr/bin/sed "s:/_/"`"
63:     device=`echo $device | /usr/bin/cut -d: -f1`"
64:   else
65:     # /dev/sfdsk/class/dsk/volume -> _gds_class_volume
66:     dev="_gds_"`echo $device | /usr/bin/sed "s/\\/dev\\/sfdsk\\/\\/\\/`"
67:     dev=`echo $dev | /usr/bin/sed "s/\\/dsk\\/\\/\\/`"
68:   fi
69: elif [ "`echo $device | /usr/bin/grep "/dev/vx/dsk/"`" != "" ]
70: then
71:   # /dev/vx/dsk/volume -> _vx_rootdg_volume
72:   # /dev/vx/dsk/disk-group/volume -> _vx_disk-group_volume
73:   dev=_vx_"`echo $device | /usr/bin/awk -F\ / '{ if (NF == 6) { print $5_"$6 } else print
"rootdg_"$5 }'`"
74: elif [ "`echo $device | /usr/bin/grep "/dev/vx/dmp/"`" != "" ]
75: then
76:   # /dev/vx/dmp/device -> _vx_pv_device
77:   dev=_vx_pv_"`echo $device | /usr/bin/cut -d/ -f5`"
78: else
79:   exit 0
80: fi
81: post_file="/etc/opt/FJSSwsts"$swstg_node"/data/DEFAULT/"$dev".pre"
82: fstype_file="/etc/opt/FJSSwsts"$swstg_node"/data/DEFAULT/"$dev".fstype"
83: bd_file="/etc/opt/FJSSwsts"$swstg_node"/data/DEFAULT/"$dev".bd"
84:
85: if [ "$mount_point" != "" ]
86: then
87:
88: # When device cannot be unmounted
89: #
90: #   if [ "$device" = "/dev/dsk/cXtXdXsX" -o "$device" = "/dev/dsk/cYtYdYsY" ]
91: #   then
92: #       /usr/sbin/lockfs -w $mount_point > /dev/null 2>&1
93: #       if [ "$bk_device" != "" ]
94: #       then
95: #           echo $bk_device > $bd_file
96: #       fi
97: #       df -ln $mount_point | cut -f2 -d: | cut -f2 -d' ' > $fstype_file
98: #       echo "fsck" > $post_file
99: #
100: # When device can be unmounted
101: #
102: #   else
103: #       /usr/bin/df -ln $mount_point | /usr/bin/cut -f2 -d: | /usr/bin/cut -f2 -d' ' > $fstype_file
104: #       /usr/sbin/umount $mount_point 2>/dev/null
105: #       if [ $? != 0 ]
106: #       then
107: #           retry_count=3
108: #           sleep_time=1
109: #           result_flag=1
110: #
111: #           while [ $retry_count -gt 0 ]
112: #           do

```

```

113:         /usr/sbin/umount $mount_point > $err_log_path/$dev.umount 2>&1
114:         if [ $? != 0 ]
115:         then
116:             retry_count=`expr $retry_count - 1`
117:             /usr/bin/sleep $sleep_time
118:         else
119:             /usr/bin/rm -f $err_log_path/$dev.umount
120:             result_flag=0
121:             break
122:         fi
123:     done
124:
125:     if [ $result_flag != 0 ]
126:     then
127:         /usr/sbin/fuser -cu $mount_point> $err_log_path/$dev.fuser 2>&1
128:         /usr/bin/ps -ef > $err_log_path/$dev.ps 2>&1
129:
130:         exit 10
131:     fi
132: fi
133: echo "mount" > $post_file
134:
135: # fi
136:
137: # When device was not mounted
138: #
139: else
140:     echo "none" > $post_file
141: fi
142:
143: exit 0

```

A.2.2 Post-processing of Backup

The name of a script file for post-processing of a backup is as follows.

For Non-clustered System

```
/etc/opt/FJSVswsts/sh/OpcBackup.post
```

For Clustered System

```
/etc/opt/FJSVswsts/logicalNodeName/sh/OpcBackup.post
```

Post-processing Script for Backup (OpcBackup.post)

```

1: #!/bin/sh
2:
3: # AdvancedCopy Manager
4: # All Rights Reserved, Copyright FUJITSU LIMITED, 2002-2006
5:
6: #
7: # Post-processing of backup processing
8: #
9: #Argument: $1 Device name of transaction disk
10: #           $2 Mount point of transaction disk
11: #
12: # Error number
13: #     2: Argument error
14: #     11: mount error
15: #     12: fsck error
16:
17: # Argument check

```

```

18: case $# in
19: 1)
20:   ;;
21: 2)
22:   ;;
23: *)
24:   exit 2
25:   ;;
26: esac
27:
28: device="`echo $1`"
29: mount_point="`echo $2`"
30:
31: # Determination of postprocessing file name
32:
33: if [ "$SWSTGNODE" != "" ]
34: then
35:   swstg_node="`echo $SWSTGNODE`"
36: else
37:   swstg_node=""
38: fi
39:
40: err_log_path="/var/opt/FJSVswsts"$swstg_node"/log"
41:
42: if [ "`echo $device | /usr/bin/grep "/dev/dsk/"`" != "" ]
43: then
44:   # /dev/dsk/c?t?d?s? -> c?t?d?s?
45:   dev="`echo $device | /usr/bin/sed "s/\dev\dsk\///"`"
46: elif [ "`echo $device | /usr/bin/grep "/dev/FJSV"`" != "" ]
47: then
48:   # /dev/FJSVmphd/dsk/implb?s? -> /dev/FJSVmphd/dsk/implb?s?
49:   # /dev/FJSVmphd/dsk/mphd?s? -> /dev/FJSVmphd/dsk/mphd?s?
50:   dev="`echo $device | /usr/bin/cut -d/ -f5`"
51: elif [ "`echo $device | /usr/bin/grep "/dev/sfdsk/"`" != "" ]
52: then
53:   if [ "`echo $device | /usr/bin/grep ":"`" != "" ]
54: then
55:     # /dev/sfdsk/class/dsk/volume:c?t?d? -> class_volume_c?t?d?
56:     dev="`echo $device | /usr/bin/sed "s/\dev\sfdsk\///"`"
57:     dev="`echo $dev | /usr/bin/sed "s/\dsk\//_/"`"
58:     dev="`echo $dev | /usr/bin/sed "s/:/_/"`"
59:     device="`echo $device | /usr/bin/cut -d: -f1`"
60: else
61:   # /dev/sfdsk/class/dsk/volume -> _gds_class_volume
62:   dev="`_gds_`echo $device | /usr/bin/sed "s/\dev\sfdsk\///"`"
63:   dev="`echo $dev | /usr/bin/sed "s/\dsk\//_/"`"
64: fi
65: elif [ "`echo $device | /usr/bin/grep "/dev/vx/dsk/"`" != "" ]
66: then
67:   # /dev/vx/dsk/volume -> _vx_rootdg_volume
68:   # /dev/vx/dsk/disk-group/volume -> _vx_disk-group_volume
69:   dev=_vx_`echo $device | /usr/bin/awk -F\| '{ if (NF == 6) { print $5_"_"$6 } else print
"rootdg_"$5 }'`"
70: elif [ "`echo $device | /usr/bin/grep "/dev/vx/dmp/"`" != "" ]
71: then
72:   # /dev/vx/dmp/device -> _vx_pv_device
73:   dev=_vx_pv_`echo $device | /usr/bin/cut -d/ -f5`"
74: else
75:   exit 0
76: fi
77: post_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".pre"
78: fstype_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".fstype"
79: bd_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".bd"

```

```

80:
81: # Confirmation of postprocessing
82: if [ ! -r $post_file ]
83: then
84:     exit 0
85: fi
86: post="/usr/bin/cat $post_file`"
87:
88: # Confirmation of FStype
89: if [ ! -r $fstype_file ]
90: then
91:     fs=""
92: else
93:     fs="/usr/bin/cat $fstype_file`"
94: fi
95:
96: # No processing
97: if [ "$post" = "none" ]
98: then
99:     /usr/bin/rm -rf $post_file 2> /dev/null
100:    /usr/bin/rm -rf $fstype_file 2> /dev/null
101:    exit 0
102: fi
103:
104: # mount processing
105: if [ "$post" = "mount" ]
106: then
107:     /usr/bin/df -l $device > /dev/null 2>&1
108:     if [ $? != 0 ]
109:     then
110:         if [ ! -r $fstype_file ]
111:         then
112:             /usr/sbin/mount $device $mount_point 2> /dev/null
113:         else
114:             if [ "$fs" = "" ]
115:             then
116:                 /usr/sbin/mount $device $mount_point 2> /dev/null
117:             else
118:                 /usr/sbin/mount -F $fs $device $mount_point 2> /dev/null
119:             fi
120:         fi
121:     if [ $? != 0 ]
122:     then
123:         retry_count=3
124:         sleep_time=1
125:         result_flag=1
126:
127:         while [ $retry_count -gt 0 ]
128:         do
129:             if [ ! -r $fstype_file ]
130:             then
131:                 /usr/sbin/mount $device $mount_point > $err_log_path/$dev.mount 2>&1
132:             else
133:                 if [ "$fs" = "" ]
134:                 then
135:                     /usr/sbin/mount $device $mount_point > $err_log_path/$dev.mount 2>&1
136:                 else
137:                     /usr/sbin/mount -F $fs $device $mount_point > $err_log_path/$dev.mount 2>&1
138:                 fi
139:             fi
140:             if [ $? != 0 ]
141:             then
142:                 retry_count=`expr $retry_count - 1`

```

```

143:         /usr/bin/sleep $sleep_time
144:     else
145:         /usr/bin/rm -f $err_log_path/$dev.mount
146:         result_flag=0
147:         break
148:     fi
149: done
150:
151: if [ $result_flag != 0 ]
152: then
153:     exit 11
154: fi
155: fi
156: fi
157: /usr/bin/rm -rf $post_file 2> /dev/null
158: /usr/bin/rm -rf $fstype_file 2> /dev/null
159: exit 0
160: fi
161:
162: # fsck processing
163: if [ "$post" = "fsck" ]
164: then
165:     /usr/sbin/lockfs -u $mount_point > /dev/null 2>&1
166:     if [ -r $bd_file ]
167:     then
168:         bk_device="`/usr/bin/cat $bd_file`"
169:         fsck_dev="`echo $bk_device | /usr/bin/sed "s/\dsk\\/\rdsk\\"`"
170:         if [ ! -r $fstype_file ]
171:         then
172:             /usr/sbin/fsck -m $fsck_dev > /dev/null 2>&1
173:         else
174:             if [ "$fs" = "" ]
175:             then
176:                 /usr/sbin/fsck -m $fsck_dev > /dev/null 2>&1
177:             else
178:                 if [ "$fs" = "sfxfs" ]
179:                 then
180:                     /usr/sbin/sfxadm $fsck_dev > /dev/null 2>&1
181:                     /usr/sbin/fsck -y -F $fs -o p,log $fsck_dev > /dev/null 2>&1
182:                 else
183:                     /usr/sbin/fsck -m -F $fs $fsck_dev > /dev/null 2>&1
184:                 fi
185:             fi
186:         fi
187:         if [ $? != 0 ]
188:         then
189:             if [ "$fs" = "" ]
190:             then
191:                 result="`/usr/sbin/fsck -n $fsck_dev | /usr/bin/grep \?`"
192:             else
193:                 result="`/usr/sbin/fsck -n -F $fs $fsck_dev | /usr/bin/grep \?`"
194:             fi
195:             if [ "$result" = "FILE SYSTEM STATE IN SUPERBLOCK IS WRONG; FIX?no" ]
196:             then
197:                 if [ "$fs" = "" ]
198:                 then
199:                     /usr/sbin/fsck -y $fsck_dev > $err_log_path/$dev.fsck 2>&1
200:                 else
201:                     /usr/sbin/fsck -y -F $fs $fsck_dev > $err_log_path/$dev.fsck 2>&1
202:                 fi
203:             fi
204:             if [ $? != 0 ]
205:             then
206:                 exit 12

```



```

206:         else
207:             /usr/bin/rm -f $err_log_path/$dev.fsck
208:         fi
209:     else
210:         echo "$result" > $err_log_path/$dev.fsck 2>&1
211:         exit 12
212:     fi
213: fi
214: fi
215: /usr/bin/rm -rf $post_file 2> /dev/null
216: /usr/bin/rm -rf $fstype_file 2> /dev/null
217: /usr/bin/rm -rf $bd_file 2> /dev/null
218: exit 0
219: fi
220:
221: exit 0

```

A.2.3 Customizing Shell Script

The method of customizing the shell script for backup pre- and post-processing is explained below.

When You Do Not Want to Unmount Transaction Volume

Edit the script to delete the comment characters ("##") on lines 90-98, 102 and 135 for a transaction volume that is mounted and that you do not want to unmount. Then specify the target device in the if statement on line 90.

This operation is possible only when the file system is one of the following two:

- ufs
- SynfinityFILE (single-partition configuration)

The pre-processing by the customized scripts is as follows:

Table A.4 Pre-processing by Customized Script

Transaction volume Status	Pre-processing
It is mounted.	Updating to operating volume is deterred.

The post-processing by the customized scripts is as follows:

Table A.5 Post-processing by Customized Script

Pre-processing	Post-processing
If Update to operating volume was prevented.	Restore capacity to update operating volume. Matching of the backup volume is checked.



Note

- Where the standard script is used, the caching of a file system and renewal of subsequent updates are inhibited by unmounting a file system before an Advanced Copy process using (EC/OPC), the integrity of the data of the entire file system is guaranteed. Although flushing the cache of a file system and renewal of subsequent updates are inhibited using "lockfs -w", when customizing a script to prevent unmounting a file system, data integrity or file consistency cannot be guaranteed if the file is in Write Open mode (the database space of Oracle which is supporting online backup).
Moreover, when the renewal of metadata is occurring, the data updates to the file system itself are not guaranteed even with "lockfs -w" specified.
It is therefore critical to consider this and to manage the integrity of data updates accordingly when you elect not to unmount a file system.
- If an SDX object is operated in slice units, do not specify slice name (*/dev/sfdsk/<class name>/dsk/<volume name>:<sdxinfo DEVNAM value>*) in Line 90. Instead, specify logical volume name (*/dev/sfdsk/<class name>/dsk/<volume name>*).

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 the unmount/mount processing in the pre-processing and post-processing scripts into offline/online processing of the mount point resource.

The mount point resource relating to the Volume is immediately offline/online whilst there is a time delay before access is actually prevented when mount/unmount is used.

Note that once the online/offline command is issued there is 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 Unmount Processing of Backup Pre-processing Script (OpcBackup.pre)

[lines 104 and 113]

(before change)	<code>/usr/sbin/umount \$mount_point</code>
(after change)	<code>/opt/VRTSvcs/bin/hares -offline resourceName -sys systemName</code>

[Example]

Adding Unmount Wait Processing to Backup Pre-processing Script (OpcBackup.pre)

[line 133]

(addition)	<code>while /usr/sbin/df -l "\$device">/dev/null 2>&1; do ;; done</code>
------------	--

[Example]

Changing Mount Processing of Backup Post-processing Script (OpcBackup.post)

[lines 112, 116, 118, 131, 135 and 137]

(before change)	<code>/usr/sbin/mount \$device \$mount_point or /usr/sbin/mount -F \$fs \$device \$mount_point</code>
(after change)	<code>/opt/VRTSvcs/bin/hares -online resourceName -sys systemName</code>

[Example]

Adding Mount Wait Processing to Backup Post-processing Script (OpcBackup.post)

[line 157]

(addition)	<code>until /usr/sbin/df -l "\$device">/dev/null 2>&1; do ;; done</code>
------------	--

A.3 Pre-processing and Post-processing of Restoration

In AdvancedCopy Manager, restore processing must be performed while transaction volumes are unmounted.

Thus, in pre-processing, the mount state of the transaction volume is determined using its name and the following processing is then performed:

Table A.6 Pre-processing for Transaction Volume

Transaction Volume Status	Pre-processing
Mounted	Unmount the transaction volume.
Unmounted	Take no action.

Post-processing tasks are determined by what has occurred in pre-processing.

Table A.7 Pre-processing and Post-processing for Restoration

Pre-processing	Post-processing
A transaction volume was unmounted.	Remount the transaction volume.
Take no action.	Take no action.

If special pre-processing or post-processing is required, you need to add the instructions as appropriate to the shell script. When customizing a script, strictly observe the following rules regarding error codes:

Table A.8 Error Code Rules

Error Code	Usage
0-99	Unusable (reserved for AdvancedCopy Manager)
100-255	Usable



Note

If the post-processing fails, the resource information may have inconsistencies. Execute the `swstsrsemrch` command.

A.3.1 Pre-processing of Restoration

The name of a script file for pre-processing of a restoration is as follows.

For Non-clustered System

```
/etc/opt/FJSVswsts/sh/OpcRestore.pre
```

For Clustered System

```
/etc/opt/FJSVswsts/logicalNodeName/sh/OpcRestore.pre
```

Pre-processing Script for Restoration (OpcRestore.pre)

```
1: #!/bin/sh
2:
3: # AdvancedCopy Manager
4: # All Rights Reserved, Copyright FUJITSU LIMITED, 2002-2006
5:
6: #
7: #   Pre-processing of restoration processing
8: #
9: #Argument: $1 Device name of transaction disk
10: #           $2 Mount point of transaction disk
11: #
12: # Error number
13: #     2: Argument error
14: #    10: umount error
15:
16:
17: # Argument check
18: case $# in
19: 1)
20:   ;;
21: 2)
22:   ;;
23: *)
24:   exit 2
```

```

25:    ;;
26: esac
27:
28: device="`echo $1`"
29: mount_point="`echo $2`"
30:
31: # Determination of postprocessing file name
32:
33: if [ "$SWSTGNODE" != "" ]
34: then
35:     swstg_node="`echo $SWSTGNODE`"
36: else
37:     swstg_node=""
38: fi
39:
40: err_log_path="/var/opt/FJSVswsts"$swstg_node"/log"
41:
42: if [ "`echo $device | /usr/bin/grep "/dev/dsk/"`" != "" ]
43: then
44:     # /dev/dsk/c?t?d?s? -> c?t?d?s?
45:     dev="`echo $device | /usr/bin/sed "s/\dev\dsk\///"`"
46: elif [ "`echo $device | /usr/bin/grep "/dev/FJSV"`" != "" ]
47: then
48:     # /dev/FJSVmpbd/dsk/mplb?s? -> /dev/FJSVmpbd/dsk/mplb?s?
49:     # /dev/FJSVmpbd/dsk/mpbd?s? -> /dev/FJSVmpbd/dsk/mpbd?s?
50:     dev="`echo $device | /usr/bin/cut -d/ -f5`"
51: elif [ "`echo $device | /usr/bin/grep "/dev/sfdsk/"`" != "" ]
52: then
53:     if [ "`echo $device | /usr/bin/grep ":"`" != "" ]
54:     then
55:         # /dev/sfdsk/class/dsk/volume:c?t?d? -> class_volume_c?t?d?
56:         dev="`echo $device | /usr/bin/sed "s/\dev\sfdsk\///"`"
57:         dev="`echo $dev | /usr/bin/sed "s/\dsk\//_/"`"
58:         dev="`echo $dev | /usr/bin/sed "s:/:_/"`"
59:         device="`echo $device | /usr/bin/cut -d: -f1`"
60:     else
61:         # /dev/sfdsk/class/dsk/volume -> _gds_class_volume
62:         dev="_gds_"`echo $device | /usr/bin/sed "s/\dev\sfdsk\///"`"
63:         dev="`echo $dev | /usr/bin/sed "s/\dsk\//_/"`"
64:     fi
65: elif [ "`echo $device | /usr/bin/grep "/dev/vx/dsk/"`" != "" ]
66: then
67:     # /dev/vx/dsk/volume -> _vx_rootdg_volume
68:     # /dev/vx/dsk/disk-group/volume -> _vx_disk-group_volume
69:     dev=_vx_"`echo $device | /usr/bin/awk -F\ / '{ if (NF == 6) { print $5_"$6 } else print
"rootdg_"$5 }'`"
70: elif [ "`echo $device | /usr/bin/grep "/dev/vx/dmp/"`" != "" ]
71: then
72:     # /dev/vx/dmp/device -> _vx_pv_device
73:     dev=_vx_pv_"`echo $device | /usr/bin/cut -d/ -f5`"
74: else
75:     exit 0
76: fi
77: post_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".pre"
78: fstype_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".fstype"
79: bd_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".bd"
80:
81: if [ "$mount_point" != "" ]
82: then
83:
84:     # When device can be unmounted
85:     #
86:     /usr/bin/df -ln $mount_point | /usr/bin/cut -f2 -d: | /usr/bin/cut -f2 -d' ' > $fstype_file

```

```

87: /usr/sbin/umount $mount_point 2> /dev/null
88: if [ $? != 0 ]
89: then
90:     retry_count=3
91:     sleep_time=1
92:     result_flag=1
93:
94:     while [ $retry_count -gt 0 ]
95:     do
96:         /usr/sbin/umount $mount_point > $err_log_path/$dev.umount 2>&1
97:         if [ $? != 0 ]
98:         then
99:             retry_count=`expr $retry_count - 1`
100:            /usr/bin/sleep $sleep_time
101:        else
102:            /usr/bin/rm -f $err_log_path/$dev.umount
103:            result_flag=0
104:            break
105:        fi
106:    done
107:
108:    if [ $result_flag != 0 ]
109:    then
110:        /usr/sbin/fuser -cu $mount_point> $err_log_path/$dev.fuser 2>&1
111:        /usr/bin/ps -ef > $err_log_path/$dev.ps 2>&1
112:
113:        exit 10
114:    fi
115: fi
116: echo "mount" > $post_file
117:
118: # When device was not mounted
119: #
120: else
121:     echo "none" > $post_file
122: fi
123:
124: exit 0

```

Restore processing is not possible on a mounted transaction volume that cannot be unmounted. Specify a device at the restore destination.

A.3.2 Post-processing of Restoration

The name of a script file for post-processing of a restoration is as follows.

For Non-clustered System

```
/etc/opt/FJVS/swsts/sh/OpcRestore.post
```

For Clustered System

```
/etc/opt/FJVS/swsts/logicalNodeName/sh/OpcRestore.post
```

Post-processing Script for Restoration (OpcRestore.post)

```

1: #!/bin/sh
2:
3: # AdvancedCopy Manager
4: # All Rights Reserved, Copyright FUJITSU LIMITED, 2002-2006
5:
6: #
7: #   Post-processing of restoration processing
8: #

```

```

 9: #Argument: $1 Device name of transaction disk
10: #           $2 Mount point of transaction disk
11: #
12: # Error number
13: #           2: Argument error
14: #           11: mount error
15:
16: # Argument check
17: case $# in
18: 1)
19:   ;;
20: 2)
21:   ;;
22: *)
23:   exit 2
24:   ;;
25: esac
26:
27: device=`echo $1`
28: mount_point=`echo $2`
29:
30: # Determination of postprocessing file name
31:
32: if [ "$SWSTGNODE" != "" ]
33: then
34:   swstg_node=`echo $SWSTGNODE`
35: else
36:   swstg_node=""
37:   fi
38:
39: err_log_path="/var/opt/FJSVswsts"$swstg_node"/log"
40:
41: if [ "`echo $device | /usr/bin/grep "/dev/dsk/"`" != "" ]
42: then
43:   # /dev/dsk/c?t?d?s? -> c?t?d?s?
44:   dev=`echo $device | /usr/bin/sed "s/\dev\dsk\///"`
45: elif [ "`echo $device | /usr/bin/grep "/dev/FJSV/"`" != "" ]
46: then
47:   # /dev/FJSVmphd/dsk/implb?s? -> /dev/FJSVmphd/dsk/implb?s?
48:   # /dev/FJSVmphd/dsk/mphd?s? -> /dev/FJSVmphd/dsk/mphd?s?
49:   dev=`echo $device | /usr/bin/cut -d/ -f5`
50: elif [ "`echo $device | /usr/bin/grep "/dev/sfdsk/"`" != "" ]
51: then
52:   if [ "`echo $device | /usr/bin/grep ":"`" != "" ]
53:   then
54:     # /dev/sfdsk/class/dsk/volume:c?t?d? -> class_volume_c?t?d?
55:     dev=`echo $device | /usr/bin/sed "s/\dev\sfdsk\///"`
56:     dev=`echo $dev | /usr/bin/sed "s/\dsk\//_"`
57:     dev=`echo $dev | /usr/bin/sed "s:/_/"`
58:     device=`echo $device | /usr/bin/cut -d: -f1`
59:   else
60:     # /dev/sfdsk/class/dsk/volume -> _gds_class_volume
61:     dev="_gds_"`echo $device | /usr/bin/sed "s/\dev\sfdsk\///"`
62:     dev=`echo $dev | /usr/bin/sed "s/\dsk\//_"`
63:   fi
64: elif [ "`echo $device | /usr/bin/grep "/dev/vx/dsk/"`" != "" ]
65: then
66:   # /dev/vx/dsk/volume -> _vx_rootdg_volume
67:   # /dev/vx/dsk/disk-group/volume -> _vx_disk-group_volume
68:   dev=_vx_"`echo $device | /usr/bin/awk -F\ / '{ if (NF == 6) { print $5_"_$6 } else print
"rootdg_"$5 }'`"
69: elif [ "`echo $device | /usr/bin/grep "/dev/vx/dmp/"`" != "" ]
70: then

```

```

71: # /dev/vx/dmp/device -> _vx_pv_device
72: dev=_vx_pv_`echo $device | /usr/bin/cut -d/ -f5`"
73: else
74:   exit 0
75: fi
76: post_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".pre"
77: fstype_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".fstype"
78: bd_file="/etc/opt/FJSVswsts"$swstg_node"/data/DEFAULT/"$dev".bd"
79:
80: # Confirmation of postprocessing
81: if [ ! -r $post_file ]
82: then
83:   exit 0
84: fi
85: post="/usr/bin/cat $post_file`"
86:
87: # Confirmation of FStype
88: if [ ! -r $fstype_file ]
89: then
90:   fs=""
91: else
92:   fs="/usr/bin/cat $fstype_file`"
93: fi
94:
95: # No processing
96: if [ "$post" = "none" ]
97: then
98:   /usr/bin/rm -rf $post_file 2> /dev/null
99:   /usr/bin/rm -rf $fstype_file 2> /dev/null
100: exit 0
101: fi
102:
103: # mount processing
104: if [ "$post" = "mount" ]
105: then
106:   /usr/bin/df -l $device > /dev/null 2>&1
107:   if [ $? != 0 ]
108:   then
109:     if [ ! -r $fstype_file ]
110:     then
111:       /usr/sbin/mount $device $mount_point 2> /dev/null
112:     else
113:       if [ "$fs" = "" ]
114:       then
115:         /usr/sbin/mount $device $mount_point 2> /dev/null
116:       else
117:         if [ "$fs" = "sfxfs" ]
118:         then
119:           dev="`echo $device | /usr/bin/sed "s/\dsk\\/\rdsk\\/"`"
120:           /usr/sbin/sfxadm $dev > /dev/null 2>&1
121:         fi
122:         /usr/sbin/mount -F $fs $device $mount_point 2> /dev/null
123:       fi
124:     fi
125:     if [ $? != 0 ]
126:     then
127:       retry_count=3
128:       sleep_time=1
129:       result_flag=1
130:
131:       while [ $retry_count -gt 0 ]
132:       do
133:         if [ ! -r $fstype_file ]

```

```

134:         then
135:             /usr/sbin/mount $device $mount_point > $err_log_path/$dev.mount 2>&1
136:         else
137:             if [ "$fs" = "" ]
138:             then
139:                 /usr/sbin/mount $device $mount_point > $err_log_path/$dev.mount 2>&1
140:             else
141:                 if [ "$fs" = "sfxfs" ]
142:                 then
143:                     dev=`echo $device | /usr/bin/sed "s/\\/dsk\\/\\/\\/rdsk\\/\\/"/`
144:                     /usr/sbin/sfxadm $dev > $err_log_path/$dev.sfxadm 2>&1
145:                 fi
146:                 /usr/sbin/mount -F $fs $device $mount_point > $err_log_path/$dev.mount 2>&1
147:             fi
148:         fi
149:         if [ $? != 0 ]
150:         then
151:             retry_count=`expr $retry_count - 1`
152:             /usr/bin/sleep $sleep_time
153:         else
154:             /usr/bin/rm -f $err_log_path/$dev.mount
155:             result_flag=0
156:             break
157:         fi
158:     done
159:
160:     if [ $result_flag != 0 ]
161:     then
162:         exit 11
163:     fi
164: fi
165: fi
166: /usr/bin/rm -rf $post_file 2> /dev/null
167: /usr/bin/rm -rf $fstype_file 2> /dev/null
168: exit 0
169: fi
170:
171: exit 0

```

A.3.3 Customizing Shell Script

The method to customize the shell script for restore pre/post-processing is explained below.

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 the unmount/mount processing in the pre-processing and post-processing script into offline/online processing of the mount point resource.

The mount point resource relating to the Volume is immediately offline/online whilst there is a time delay before access is actually prevented when mount/unmount is used.

Note that once the online/offline command is issued there is 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 Unmount Processing of Restore Pre-processing Script (OpcRestore.pre)

[lines 87 and 96]

(before change)	<code>/usr/sbin/umount \$mount_point</code>
(after change)	<code>/opt/VRTSvcs/bin/hares -offline <i>resourceName</i> -sys <i>systemName</i></code>

[Example]

Adding Unmount Wait Processing to Restore Pre-processing Script (OpcRestore.pre)

[line 116]

(addition)	<code>while /usr/sbin/df -l "\$device">/dev/null 2>&1; do ;; done</code>
------------	--

[Example]

Changing Mount Processing of Restore Post-processing Script (OpcRestore.post)

[lines 111, 115, 122, 135, 139 and 146]

(before change)	<code>/usr/sbin/mount \$device \$mount_point or /usr/sbin/mount -F \$fs \$device \$mount_point</code>
(after change)	<code>/opt/VRTSvcs/bin/hares -online <i>resourceName</i> -sys <i>systemName</i></code>

[Example]

Adding Mount Wait Processing to Restore Post-processing Script (OpcRestore.post)

[line 166]

(addition)	<code>until /usr/sbin/df -l "\$device">/dev/null 2>&1; do ;; done</code>
------------	--

Appendix B Sample Script for Backup Operations

This appendix provides information on sample scripts for AdvancedCopy Manager.

A sample script allows you to back up data on a transaction volume with AdvancedCopy Manager. The manager of the backup operation should customize the sample scripts according to their specific requirements and apply them to the backup.

Point

This section explains how the sample script is created for the backup management function. By using it as a reference when creating similar scripts for the replication management function, you can copy Oracle databases.


B.1 Backup and Restoration of Transaction Volume

B.1.1 Backup Procedure

The following is an overview of the backup procedure:

1. Check the work operation status
Check whether processing is in operation or not.
2. Declare the start of online backup
If processing was found to be in operation in step 1, declare the start of online backup for the database to stop updating the data files. This makes the slices that store transaction data ready for online backup.
3. Backup (copy) slices
Backup (copy) the slices that store transaction data using the AdvancedCopy Manager backup function.
4. Declare the end of online backup
If processing was found to be in operation in step 1, declare the end of online backup for the database to resume updating the data files.

Figure B.1 Backup




B.1.2 Restoration Procedure

The following provides an overview of how to perform restoration from a backup volume.

1. Checking the processing stop status
Check that processing has stopped. If processing is in operation, stop the processing completely.
2. Performing restoration from a backup volume
Using the restoration function of AdvancedCopy Manager, restore data from a backup volume to the transaction volume.

Figure B.2 Performing Restoration from Backup Volume



B.2 Sample Scripts for Oracle

B.2.1 Overview of Sample Scripts

The sample shell scripts provided are summarized below. These scripts are stored in the /etc/opt/FJSVswsts/samp directory.

Table B.1 Sample Shell Script for Oracle

No.	Type	Script Name	Associated Transaction	Associated Resource Type
1	Collection of DB information	swst_or_iv	Oracle transaction	-
2	Backup	swst_or_fs_bk	Oracle transaction, general file transaction	File system
3	Backup	swst_or_rd_bk	Oracle transaction	Raw device

B.2.1.1 Collecting DB Information

swst_or_iv outputs information about an Oracle database. The output file names and their contents are:

- spdata.lst : List of table space names and storage destinations
- spdata.srt : Results of spdata.lst, with storage destination as the sort key
- spdata.bgn : Command in the following format:

```
Alter tablespace Table-space-name begin backup;
```

- spdata.end : Command in the following format:

```
Alter tablespace Table-space-name end backup;
```

- spcont.lst : Oracle database control file information
- splog.lst : Oracle database log file information

Use the database information above as a reference for creating shell scripts for backup processing.



Note

For Oracle 9i or later, replace the "svrmgrl" command used in "swst_or_iv" with the sqlplus command.

Format

swst_or_iv

Usage Example

```
# swst_or_iv
swst_or_iv completed.
```

Result Example

```
spdata.lst
TABLESPACE-NAME          DEVICE-NAME
-----
ORA10                    /ora1/ora0.dbf
ORA20                    /ora2/ora0.dbf
RBS                      /oracle/ora/oradata/rbs01.dbf
SYSTEM                  /oracle/ora/oradata/system01.dbf
TEMP                    /oracle/ora/oradata/temp01.dbf
TOOLS                   /oracle/ora/oradata/tools01.dbf
USERS                   /oracle/ora/oradata/users01.dbf
spdata.bgn
alter tablespace ORA10 begin backup;
alter tablespace ORA20 begin backup;
alter tablespace RBS begin backup;
alter tablespace SYSTEM begin backup;
alter tablespace TEMP begin backup;
alter tablespace TOOLS begin backup;
alter tablespace USERS begin backup;
spdata.end
alter tablespace ORA10 end backup;
alter tablespace ORA20 end backup;
alter tablespace RBS end backup;
alter tablespace SYSTEM end backup;
alter tablespace TEMP end backup;
alter tablespace TOOLS end backup;
alter tablespace USERS end backup;
spcont.lst
NAME                      TYPE      VALUE
-----
control_files             string    /oracle/ora/oradata/control
splog.lst
NAME                      TYPE      VALUE
-----
log_archive_dest         string    /oracle/ora/oradata/archive
```

Processing Outline

```
#!/bin/sh
reterr() --- An error message is output before an error is returned.
sqldata()
{
  sqlplus AAA/BBB <<!
  SELECT --- Acquire Oracle table space information.
}
sqllog()
{
  sqlplus svrmgrl <<!
  show --- Acquire Oracle log information.
}
sqlcont()
{
  sqlplus vrmgrl <<!
  show --- Acquire Oracle control information.
```

```

}

# main()

# 0) Environmental variable setting
PATH=$PATH:/usr/bin:/usr/sbin:/opt/FJSVswsts/bin

# 1) Acquisition of Oracle table space information
sqldata() invocation --- Acquire Oracle table space information

# 2) Acquisition of Oracle log information
sqllog() invocation --- Acquire Oracle log information

# 3) Acquisition of Oracle control information
sqlcont() invocation --- Acquire Oracle control information

exit 0

```

B.2.1.2 File System Backup

swst_or_fs_bk, run using AdvancedCopy Manager, backs up processing data stored in an Oracle database on a file system. Processing data is copied to the backup volumes of AdvancedCopy Manager. These scripts correspond to steps 2 to 5 in "[B.2.2.1 DB Information Collection](#)".

Format

```
swst_or_fs_bk {-a | -i} mountPointName
```

Options

-a

Specify when an Oracle database is online. The start of backup processing is declared to the Oracle server.

-i

Specify when an Oracle database is stopped or the file is a general file. The start of backup processing is not declared to the Oracle server.

Operands

mountPointName

Specify the name of the mount point targeted for backup.

Example Using swst_or_fs_bk

While processing is active, back up /gyoumu_db in which an Oracle database is located.

```

# swst_or_fs_bk -a /gyoumu_db
swst_or_fs_bk completed. ( /gyoumu_db -> /dev/dsk/ctl3d0s2/gyoumu_db_SW at 23:00 on
1999.11.01 )

```

Processing Outline of swst_or_fs_bk

```

#!/bin/sh
usage() --- Command syntax is output before an error is returned.
reterr() --- An error message is output before an error is returned.
chkmp() --- Check for a mount point
get_mp_blk() --- Conversion from a mount point to a device
getfs() --- Conversion from a mount point to a file system type
sqlbegin()
{
sqlplus AAA/BBB <<!
alter system switch logfile;
alter system flush shared_pool;

```

```

alter tablespace CCC begin backup; --- Notify Oracle of the start of backup.
}
sqlend()
{
sqlplus AAA/BBB <<!
alter tablespace CCC end backup; --- Notify Oracle of the stopping of backup.
}
getbd() --- Fetch the name of the backup (copy) destination device of AdvancedCopy Manager.
getdate() --- Fetch the backup (copy) date of AdvancedCopy Manager.
gettime() --- Fetch the backup (copy) time of AdvancedCopy Manager.

# main()
Parameter analysis and mount point check

# 0) Environmental variable setting
PATH=$PATH:/usr/bin:/usr/sbin:/opt/FJSVswsts/bin

# 1) Notification issued to Oracle of the start of backup
if [ "-a" operand specification ]
then
sqlbegin() invocation --- Notify Oracle of the start of backup.
Error handling
fi

# 2) Disk backup (copy) by AdvancedCopy Manager command
swstbackup $UDBLK --- Copy the disk.
Error handling

# 3) Notification issued to Oracle of the stopping of backup
if [ "-a" operand specification ]
then
sqlend() invocation --- Notify Oracle of the stopping of backup.
Error handling
fi

# 4) Acquisition of the backup (copy) destination of AdvancedCopy Manager
BDBLK="`getbd $UDBLK`" --- Fetch a backup volume as a character device.
BDDATE="`getdate $UDBLK`"--- Fetch the backup date.
BDTIME="`gettime $UDBLK`"--- Fetch the backup time.

exit 0

```

B.2.1.3 Raw Device Backup

swst_or_rd_bk is run using AdvancedCopy Manager to back up live data using an Oracle database on a raw device. Live data is copied to work disks prepared by users. These scripts apply to steps 2 to 5 in ["B.2.2.1 DB Information Collection"](#).

Format

```
swst_or_rd_bk {-a | -i} deviceName
```

Options

-a

Specify when an Oracle database is online. The start of backup processing is declared to the Oracle server.

-i

Specify when an Oracle database is stopped or the file is a general file. The start of backup processing is not declared to the Oracle server.

Operands

deviceName

Specify the name of the raw device targeted for backup.

Example Using swst_or_rd_bk

While live processing is active, back up /dev/rdskc0t0d0s0 in which an Oracle database is located.

```
# swst_or_rd_bk -a /dev/rdskc0t0d0s6
swst_or_rd_bk completed. ( /dev/rdskc0t0d0s6 -> /dev/rdskc1t3d0s2 at 23:01 on 1999.11.01 )
```

Processing Outline of swst_or_rd_bk

```
#!/bin/sh

usage() --- Command syntax is output before an error is returned.
reterr() --- An error message is output before an error is returned.
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 AAA/BBB <<!
alter system switch logfile;
alter system flush shared_pool;
alter tablespace CCC begin backup; --- Notify Oracle of the start of backup.
}
sqlend()
{
sqlplus AAA/BBB <<!
alter tablespace CCC end backup; --- Notify Oracle of the stopping of backup.
}
getbd() --- Fetch the name of the backup (copy) destination device of AdvancedCopy Manager.
getdate() --- Fetch the backup (copy) date of AdvancedCopy Manager.
gettime() --- Fetch the backup (copy) time of AdvancedCopy Manager.

# main()
Parameter analysis and raw device check

# 0) Environmental variable setting
PATH=$PATH:/usr/bin:/usr/sbin:/opt/FJVSvwsts/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
```

B.2.2 Applying Sample Shell Scripts

B.2.2.1 DB Information Collection

Sample Shell Script Update

Update the procedures below of `swst_or_iv` according to actual operations. The following numbers correspond to those in the scripts:

- 0) Environmental variable setting
- 1) Acquisition of Oracle table space information
- 2) Acquisition of Oracle log information
- 3) Acquisition of Oracle control information

Procedures 1) to 3) change Oracle account information.

These scripts can operate only on business servers of AdvancedCopy Manager.

B.2.2.2 File System Backup

AdvancedCopy Manager Setup

Set up backup volumes so that the mount point targeted for backup is backed up correctly in AdvancedCopy Manager. The generations of the backup volumes to be managed are arbitrary.

Sample Shell Script Update

Update the procedures below of `swst_or_fs_bk` to adjust them for actual operations. The following numbers correspond to those in the scripts.

- 0) Environmental variable setting
- 1) Notification issued to Oracle of the start of backup
- 3) Notification issued to Oracle of the stopping of backup

Procedures 1) and 3) refer to the `spdata.bgn` file and `spdata.end` file created by `swst_or_iv`.

These scripts can operate only on business servers of AdvancedCopy Manager.

B.2.2.3 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 Backup Volume Name and Processing Time

After executing a shell script, a backup volume name and the processing time are output in a message. This information is needed for restoration, so it must be recorded and managed.

B.2.3 Restore Method

B.2.3.1 File system

Use the restore function of AdvancedCopy Manager for restore processing.

B.2.3.2 Raw Device

Use the restore function of AdvancedCopy Manager for restore processing.

B.2.4 Notes on Using 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:

- Initialization parameter file
 - Every time a change occurs, back up this file using the cp command of the OS is executed for backup.
- Control files
 - You are recommended to use mirroring using Oracle.
 - If a change is made in the DB environment, use the following command to create a trace file. If a control file is destroyed, execute the SQL statement in the trace file to recover the control file.

```
alter database backup control file to trace;
```

- Data files
 - The recommended location of data files is on a raw device.
 - To store a data file on the file system, disable the automatic extended function of the data file (AUTOEXTEND OFF). Customization of the sample shell script is required to prevent unmounting during context processing.
 - Create backup copies in AdvancedCopy Manager.
 - The range of table spaces targeted for backup must be within the unit of synchronous backup processing (begin backup/end backup) in Oracle and the unit of backup in AdvancedCopy Manager. For example, if backup copies must be created in units of LUNs in AdvancedCopy Manager, place one table space in one LUN.
 - Archive all of the latest logs by executing the following command after backing up all table spaces in the system:

```
alter system archive log current;
```

- REDO files
 - You are recommended to mirror the online REDO log files and archive log files using Oracle or the OS.
 - No online REDO log file needs to be backed up.



Note

For online backup operations, the target Oracle database must use an archive log.

Recovery of Oracle Databases

After an Oracle database has stopped, restore the database, and restart it. To recover the database to its latest state, use Oracle's recovery function.

Appendix C Pre-processing and Post-processing of Replication

The replication explains the replication pre-processing and post-processing that are performed in AdvancedCopy Manager.

C.1 Overview

The replication pre-processing and post-processing scripts are started before and after the replication (i.e., the copy) processing. AdvancedCopy Manager performs the processing required for a replication source volume and replication destination volume by executing these scripts.



If replication is performed in group units, the replication pre-processing and post-processing scripts are not started. Therefore, replication pre-processing and post-processing must be performed manually when replication is performed in group units.

C.2 Pre-processing and Post-processing of Replication

The pre-processing and post-processing need to be performed for both the replication source volume and the replication destination volume. The purpose of each is as follows.

- Pre-processing and post-processing for a replication source volume
This is intended to freeze the replication source data in a state in which consistency is temporarily ensured. In pre-processing, source data is temporarily frozen, and after the replication is complete, the frozen data is released in post-processing.
- Pre-processing and post-processing for a replication destination volume
This is intended to prevent applications from accessing the copy destination volume during replication. In pre-processing, replication data for the destination is given an access-disabled status. This status is released in post-processing, after the replication is complete.

The specific tasks executed during the pre-processing and post-processing differ according to the data contents of the volume. The next sub section explains the pre-processing and post-processing of the file system.

C.2.1 Pre-processing and Post-processing for File System

The pre-processing and post-processing for a file system is typically the unmounting and remounting of the file system. (For Windows, the processing consists of allocating and cancelling a drive letter.) Pre-processing and post-processing of the replication source volume is performed in order to guarantee the replication source data by disabling access from other processes during execution of synchronous processing and snapshot copy.

If no volume is mounted, the pre-processing and post-processing are not performed.

If the replication source/destination volume is a Logical Unit (disk), pre-processing and post-processing are not performed for the slices included in the Logical Unit (disk). Therefore, the slices included in the Logical Unit (disk) must be unmounted before creating a replica, and the unmounted slices must be mounted after the replica is made.

The implementation status of the pre-processing and post-processing for a file system is shown below. This table shows whether pre-processing and post-processing are performed depending on the command type. It also shows the copy status at the time of the command execution.

Table C.1 Pre-processing and Post-processing for File System

Command	Copy Status	Source/Destination	Pre-processing	Post-processing
swsrpstartsync	Copy not started, or replication establishment	Source	X	X
		Destination	0--*1	X

Command	Copy Status	Source/Destination	Pre-processing	Post-processing
swsrpmake	Equivalency maintenance	Source	0	0
		Destination	X	0--*2
swsrpcancel	Equivalency maintenance	Source	0	0
		Destination	X	0--*2
	Replication establishment	Source	X	X
		Destination	X	X
	Copying	Source	X	X
		Destination	X	X
swsrpmake	Copy not started	Source	0	0
		Destination	0	0


0 : Executed,
X : Not executed

The post-processing for the pre-processing executed in step (*1) is executed in step (*2). In the post-processing in (2), re-mount is performed for the mount point where unmount was performed in the pre-processing in (1).

The pre-processing and post-processing for a file system can be customized. The processing is started in accordance with the table above when a replication management command is executed. For server-to-server replication, these scripts are executed by remote execution using TCP/IP in the pre-processing and post-processing of a volume that is not connected.

The operation image before and after processing is shown in the following figure.

Figure C.1 Synchronized Replication Processing



If processing other than the standard pre-processing and post-processing is required in the pre-processing and post-processing, customize the scripts.

When customizing a script, strictly observe the following rules regarding error codes:

Table C.2 Error Code Rules

Error Code	Usage
0-99	Unusable (reserved for AdvancedCopy Manager)
100-255	Usable



Note

AdvancedCopy Manager does not support pre-processing and post-processing of databases. Database pre-processing and post-processing must be executed uniquely before and after (respectively) the execution of the AdvancedCopy Manager command.

C.2.2 Pre-processing When Replication Is Executed

The replication source volume script (RepSrc.pre) and the replication destination volume script (RepDst.pre) are prepared for a pre-processing script, and it is stored in the following directory. Customize these scripts according to the processing requirements.

For Non-clustered System

```
/etc/opt/FJSVswsrp/sh
```

For Clustered System

```
/etc/opt/FJSVswsrp/logicalNodeName/sh
```

Replication Source Volume Pre-processing Sample Script (RepSrc.pre)

```
1: #!/bin/sh
2:
3: # AdvancedCopy Manager
4: # All Rights Reserved, Copyright FUJITSU LIMITED, 2002-2009
5:
6: #
7: #   Pre-processing 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:
16: # Argument check
17: case $# in
18: 2)
19:   ;;
20: *)
21:   exit 2
22:   ;;
23: esac
24:
25: device=$1
26: mount_point=$2
27:
28: # Determination postprocessing file name
29: if [ "`echo $device | /usr/bin/grep "/dev/dsk/"`" != "" ]
30: then
31:   # /dev/dsk/c?t?d?s? -> c?t?d?s?
32:   dev="`echo $device | /usr/bin/sed "s/\dev\dsk\///"`"
33: elif [ "`echo $device | /usr/bin/grep "/dev/FJSV/"`" != "" ]
34: then
35:   # /dev/FJSVmphd/dsk/mplb?s? -> mplb?s?
36:   # /dev/FJSVmphd/dsk/mphd?s? -> mphe?s?
37:   dev="`echo $device | /usr/bin/cut -d/ -f5`"
38: elif [ "`echo $device | /usr/bin/grep "/dev/sfdsk/"`" != "" ]
39: then
40:   if [ "`echo $device | /usr/bin/grep ":"`" != "" ]
41:   then
42:     # /dev/sfdsk/class/dsk/volume:c?t?d? -> class_volume_c?t?d?
43:     dev="`echo $device | /usr/bin/sed "s/\dev\sfdsk\///"`"
44:     dev="`echo $dev | /usr/bin/sed "s/\dsk\//_"`"
45:     dev="`echo $dev | /usr/bin/sed "s:/:_/"`"
```

```

46:     device=`echo $device | /usr/bin/cut -d: -f1`
47:     else
48:         # /dev/sfdsk/class/dsk/volume -> _gds_class_volume
49:         dev=_gds_`echo $device | /usr/bin/sed "s\/dev\/sfdsk\/\/\/"`
50:         dev=`echo $dev | /usr/bin/sed "s\/dsk\/\/\/"`
51:     fi
52: elif [ "`echo $device | /usr/bin/grep "/dev/vx/dsk/"`" != "" ]
53: then
54:     # /dev/vx/dsk/volume -> _vx_rootdg_volume
55:     # /dev/vx/dsk/disk-group/volume -> _vx_disk-group_volume
56:     dev=_vx_`echo $device | /usr/bin/awk -F\/ '{ if (NF == 6) { print $5_"$6 } else print
"rootdg_"$5 }'`
57: elif [ "`echo $device | /usr/bin/grep "/dev/vx/dmp/"`" != "" ]
58: then
59:     # /dev/vx/dmp/device -> _vx_pv_device
60:     dev=_vx_pv_`echo $device | /usr/bin/cut -d/ -f5`
61: else
62:     exit 0
63: fi
64: post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".spre"
65: fstype_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".sfstype"
66:
67: err_log_path="/var/opt/FJSVswsrp/"$SWSTGNODE"/log"
68:
69: if [ "$mount_point" != "" ]
70: # When device was mounted
71: #
72: then
73:
74: # When device cannot be unmounted
75: #
76: # if [ "$device" = "/dev/dsk/cXtXdXsX" -o "$device" = "/dev/dsk/cYtYdYsY" ]
77: # then
78: #     /usr/sbin/lockfs -w $mount_point > /dev/null 2>&1
79: #     echo "lockfs,$mount_point" > $post_file
80: #
81: # When device can be unmounted
82: #
83: # else
84: #
85:     /usr/bin/df -ln $mount_point | /usr/bin/cut -f2 -d: | /usr/bin/cut -f2 -d' ' > $fstype_file
86:
87:     /usr/sbin/umount $mount_point 2>/dev/null
88:
89:     if [ $? != 0 ]
90:     then
91:         retry_count=3
92:         sleep_time=1
93:         result_flag=1
94:
95:         while [ $retry_count -gt 0 ]
96:         do
97:             /usr/sbin/umount $mount_point > $err_log_path/$dev.umount 2>&1
98:             if [ $? != 0 ]
99:             then
100:                 retry_count=`expr $retry_count - 1`
101:                 /usr/bin/sleep $sleep_time
102:             else
103:                 /usr/bin/rm -f $err_log_path/$dev.umount
104:                 result_flag=0
105:                 break
106:             fi
107:         done

```

```

108:
109:     if [ $result_flag != 0 ]
110:     then
111:         /usr/sbin/fuser -cu $mount_point> $err_log_path/$dev.fuser 2>&1
112:         /usr/bin/ps -ef > $err_log_path/$dev.ps 2>&1
113:         exit 10
114:     fi
115: fi
116:
117:     echo "mount,$mount_point" > $post_file
118:
119: # fi
120:
121: # When device was not mounted
122: #
123: else
124:     echo "none" > $post_file
125: fi
126:
127: exit 0

```

Replication Destination Volume Pre-processing Sample Script (RepDst.pre)

```

1: #!/bin/sh
2:
3: # AdvancedCopy Manager
4: # All Rights Reserved, Copyright FUJITSU LIMITED, 2002-2006
5:
6: #
7: #   Pre-processing 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(system error)
14: #       10: umount error
15:
16: # Argument check
17: case $# in
18: 2)
19:     ;;
20: *)
21:     exit 2
22:     ;;
23: esac
24:
25: device=$1
26: mount_point=$2
27:
28: # Determination postprocessing file name
29: if [ "`echo $device | /usr/bin/grep "/dev/dsk/"`" != "" ]
30: then
31:     # /dev/dsk/c?t?d?s? -> c?t?d?s?
32:     dev="`echo $device | /usr/bin/sed "s/\dev\dsk\///"``"
33: elif [ "`echo $device | /usr/bin/grep "/dev/FJSV/"`" != "" ]
34: then
35:     # /dev/FJSVmphd/dsk/implb?s? -> implb?s?
36:     # /dev/FJSVmphd/dsk/mphd?s? -> mphe?s?
37:     dev="`echo $device | /usr/bin/cut -d/ -f5`"
38: elif [ "`echo $device | /usr/bin/grep "/dev/sfdsk/"`" != "" ]
39: then
40:     if [ "`echo $device | /usr/bin/grep ":"`" != "" ]

```



```

41: then
42:     # /dev/sfdsk/class/dsk/volume:c?t?d? -> class_volume_c?t?d?
43:     dev=`echo $device | /usr/bin/sed "s/\dev\sfdsk\///"`
44:     dev=`echo $dev | /usr/bin/sed "s/\dsk\//_"`
45:     dev=`echo $dev | /usr/bin/sed "s:/_/"`
46:     device=`echo $device | /usr/bin/cut -d: -f1`
47: else
48:     # /dev/sfdsk/class/dsk/volume -> _gds_class_volume
49:     dev=_gds_`echo $device | /usr/bin/sed "s/\dev\sfdsk\///"`
50:     dev=`echo $dev | /usr/bin/sed "s/\dsk\//_"`
51:     post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".dpre"
52: fi
53: elif [ "`echo $device | /usr/bin/grep "/dev/vx/dsk/"`" != "" ]
54: then
55:     # /dev/vx/dsk/volume -> _vx_rootdg_volume
56:     # /dev/vx/dsk/disk-group/volume -> _vx_disk-group_volume
57:     dev=_vx_`echo $device | /usr/bin/awk -F\ | '{ if (NF == 6) { print $5_"_"$6 } else print
"rootdg_"$5 }'`
58: elif [ "`echo $device | /usr/bin/grep "/dev/vx/dmp/"`" != "" ]
59: then
60:     # /dev/vx/dmp/device -> _vx_pv_device
61:     dev=_vx_pv_`echo $device | /usr/bin/cut -d/ -f5`
62: else
63:     exit 0
64: fi
65: post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".dpre"
66:
67: err_log_path="/var/opt/FJSVswsrp/"$SWSTGNODE"/log"
68:
69: if [ "$mount_point" != "" ]
70: # When device was mounted
71: #
72: then
73:
74:     /usr/sbin/umount $mount_point 2>/dev/null
75:
76:     if [ $? != 0 ]
77:     then
78:         retry_count=3
79:         sleep_time=1
80:         result_flag=1
81:
82:         while [ $retry_count -gt 0 ]
83:         do
84:             /usr/sbin/umount $mount_point > $err_log_path/$dev.umount 2>&1
85:             if [ $? != 0 ]
86:             then
87:                 retry_count=`expr $retry_count - 1`
88:                 /usr/bin/sleep $sleep_time
89:             else
90:                 /usr/bin/rm -f $err_log_path/$dev.umount
91:                 result_flag=0
92:                 break
93:             fi
94:         done
95:
96:         if [ $result_flag != 0 ]
97:         then
98:             /usr/sbin/fuser -cu $mount_point> $err_log_path/$dev.fuser 2>&1
99:             /usr/bin/ps -ef > $err_log_path/$dev.ps 2>&1
100:             exit 10
101:         fi
102:     fi

```

```

103:
104:     echo "mount,$mount_point" > $post_file
105:
106: # When device was not mounted
107: #
108: else
109:     echo "none" > $post_file
110: fi
111:
112: exit 0

```

C.2.3 Post-processing When Replication Is Executed

The replication source volume script (RepSrc.post) and the replication destination volume script (RepDst.post) are prepared for a post-processing script, which is stored in the following directory. This script must be customized according to the processing requirements.

For Non-clustered System

```
/etc/opt/FJSVswsrp/sh
```

For Clustered System

```
/etc/opt/FJSVswsrp/logicalNodeName/sh
```

Replication Source Volume Post-processing Sample Script (RepSrc.post)

```

1: #!/bin/sh
2:
3: # AdvancedCopy Manager
4: # All Rights Reserved, Copyright FUJITSU LIMITED, 2002-2006
5:
6: #
7: #   Post-processing of Replication(Source) processing
8: #
9: #   Argument: $1 Device name of Source disk
10: #
11: # Error number
12: #       2: Argument error
13: #       11: mount error
14:
15: # Argument check
16: case $# in
17: 1)
18:     ;;
19: *)
20:     exit 2
21:     ;;
22: esac
23:
24: device=$1
25:
26: # Determination of postprocessing file name
27: if [ "`echo $device | /usr/bin/grep "/dev/dsk/"`" != "" ]
28: then
29:     # /dev/dsk/c?t?d?s? -> c?t?d?s?
30:     dev="`echo $device | /usr/bin/sed "s/\dev\dsk\///"`"
31: elif [ "`echo $device | /usr/bin/grep "/dev/FJSV/"`" != "" ]
32: then
33:     # /dev/FJSVmphd/dsk/mplb?s? -> mplb?s?
34:     # /dev/FJSVmphd/dsk/mphd?s? -> mphe?s?
35:     dev="`echo $device | /usr/bin/cut -d/ -f5`"
36: elif [ "`echo $device | /usr/bin/grep "/dev/sfdsk/"`" != "" ]
37: then

```

```

38:   if [ "`echo $device | /usr/bin/grep ":\`" != "" ]
39:   then
40:       # /dev/sfdsk/class/dsk/volume:c?t?d? -> class_volume_c?t?d?
41:       dev=`echo $device | /usr/bin/sed "s/\dev\sfdsk\///"`
42:       dev=`echo $dev | /usr/bin/sed "s/\dsk\//_"`
43:       dev=`echo $dev | /usr/bin/sed "s:/_/"`
44:       device=`echo $device | /usr/bin/cut -d: -f1`
45:   else
46:       # /dev/sfdsk/class/dsk/volume -> _gds_class_volume
47:       dev=_gds_`echo $device | /usr/bin/sed "s/\dev\sfdsk\///"`
48:       dev=`echo $dev | /usr/bin/sed "s/\dsk\//_"`
49:   fi
50: elif [ "`echo $device | /usr/bin/grep "/dev/vx/dsk/" != "" ]
51: then
52:     # /dev/vx/dsk/volume -> _vx_rootdg_volume
53:     # /dev/vx/dsk/disk-group/volume -> _vx_disk-group_volume
54:     dev=_vx_`echo $device | /usr/bin/awk -F\ / '{ if (NF == 6) { print $5_"_"$6 } else print
"rootdg_"$5 }'`
55: elif [ "`echo $device | /usr/bin/grep "/dev/vx/dmp/" != "" ]
56: then
57:     # /dev/vx/dmp/device -> _vx_pv_device
58:     dev=_vx_pv_`echo $device | /usr/bin/cut -d/ -f5`
59: else
60:     exit 0
61: fi
62: post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".spre"
63: fstype_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".sfstype"
64:
65: err_log_path="/var/opt/FJSVswsrp/"$SWSTGNODE"/log"
66:
67: # Confirmation of postprocessing
68: if [ ! -r $post_file ]
69: then
70:     exit 0
71: fi
72: post=`/usr/bin/cat $post_file | /usr/bin/cut -d',' -f1`
73: mount_point=`/usr/bin/cat $post_file | /usr/bin/cut -d',' -f2`
74:
75: # Confirmation of FStype
76: if [ ! -r $fstype_file ]
77: then
78:     fs=""
79: else
80:     fs=`/usr/bin/cat $fstype_file`
81: fi
82:
83: # No processing
84: if [ "$post" = "none" ]
85: then
86:     /usr/bin/rm -rf $post_file 2> /dev/null
87:     /usr/bin/rm -rf $fstype_file 2> /dev/null
88:     exit 0
89: fi
90:
91: # mount processing
92: if [ "$post" = "mount" ]
93: then
94:     /usr/bin/df -l $device > /dev/null 2>&l
95:     if [ $? != 0 ]
96:     then
97:         if [ ! -r $fstype_file ]
98:         then
99:             /usr/sbin/mount $device $mount_point 2> /dev/null

```

```

100:     else
101:         if [ "$fs" = "" ]
102:         then
103:             /usr/sbin/mount $device $mount_point 2> /dev/null
104:         else
105:             /usr/sbin/mount -F $fs $device $mount_point 2> /dev/null
106:         fi
107:     fi
108:     if [ $? != 0 ]
109:     then
110:         retry_count=3
111:         sleep_time=1
112:         result_flag=1
113:
114:         while [ $retry_count -gt 0 ]
115:         do
116:             if [ ! -r $fstype_file ]
117:             then
118:                 /usr/sbin/mount $device $mount_point > $err_log_path/$dev.mount 2>&1
119:             else
120:                 if [ "$fs" = "" ]
121:                 then
122:                     /usr/sbin/mount $device $mount_point > $err_log_path/$dev.mount 2>&1
123:                 else
124:                     /usr/sbin/mount -F $fs $device $mount_point > $err_log_path/$dev.mount 2>&1
125:                 fi
126:             fi
127:             if [ $? != 0 ]
128:             then
129:                 retry_count=`expr $retry_count - 1`
130:                 /usr/bin/sleep $sleep_time
131:             else
132:                 /usr/bin/rm -f $err_log_path/$dev.mount
133:                 result_flag=0
134:                 break
135:             fi
136:         done
137:
138:         if [ $result_flag != 0 ]
139:         then
140:             exit 11
141:         fi
142:     fi
143: fi
144: /usr/bin/rm -rf $post_file 2> /dev/null
145: /usr/bin/rm -rf $fstype_file 2> /dev/null
146: exit 0
147: fi
148:
149: # lockfs processing
150: if [ "$post" = "lockfs" ]
151: then
152:     /usr/sbin/lockfs -u $mount_point > /dev/null 2>&1
153:     /usr/bin/rm -rf $post_file 2> /dev/null
154:     /usr/bin/rm -rf $fstype_file 2> /dev/null
155:     exit 0
156: fi
157:
158: exit 0

```

Replication Destination Volume Post-processing Sample Script (RepDst.post)

```
1: #!/bin/sh
2:
3: # AdvancedCopy Manager
4: # All Rights Reserved, Copyright FUJITSU LIMITED, 2002-2007
5:
6: #
7: #   Post-processing 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: #           12: fsck error
15:
16: # Argument check
17: case $# in
18: 1)
19:     ;;
20: *)
21:     exit 2
22:     ;;
23: esac
24:
25: device=$1
26:
27: # Determination of postprocessing file name
28: if [ "`echo $device | /usr/bin/grep "/dev/dsk/"`" != "" ]
29: then
30:     # /dev/dsk/c?t?d?s? -> c?t?d?s?
31:     dev="`echo $device | /usr/bin/sed "s/\dev\dsk\///"``"
32: elif [ "`echo $device | /usr/bin/grep "/dev/FJSV"``" != "" ]
33: then
34:     # /dev/FJSVmphd/dsk/mplb?s? -> mplb?s?
35:     # /dev/FJSVmphd/dsk/mphd?s? -> mpdh?s?
36:     dev="`echo $device | /usr/bin/cut -d/ -f5`"
37: elif [ "`echo $device | /usr/bin/grep "/dev/sfdsk/"``" != "" ]
38: then
39:     if [ "`echo $device | /usr/bin/grep ":"``" != "" ]
40:     then
41:         # /dev/sfdsk/class/dsk/volume:c?t?d? -> class_volume_c?t?d?
42:         dev="`echo $device | /usr/bin/sed "s/\dev\sfdsk\///"``"
43:         dev="`echo $dev | /usr/bin/sed "s/\dsk\//_"``"
44:         dev="`echo $dev | /usr/bin/sed "s/:/_/"``"
45:         device="`echo $device | /usr/bin/cut -d: -f1`"
46:     else
47:         # /dev/sfdsk/class/dsk/volume -> _gds_class_volume
48:         dev="_gds_`echo $device | /usr/bin/sed "s/\dev\sfdsk\///"``"
49:         dev="`echo $dev | /usr/bin/sed "s/\dsk\//_"``"
50:     fi
51: elif [ "`echo $device | /usr/bin/grep "/dev/vx/dsk/"``" != "" ]
52: then
53:     # /dev/vx/dsk/volume -> _vx_rootdg_volume
54:     # /dev/vx/dsk/disk-group/volume -> _vx_disk-group_volume
55:     dev=_vx_`echo $device | /usr/bin/awk -F\| '{ if (NF == 6) { print $5_"_$6 } else print
"rootdg_"$5 }'``"
56: elif [ "`echo $device | /usr/bin/grep "/dev/vx/dmp/"``" != "" ]
57: then
58:     # /dev/vx/dmp/device -> _vx_pv_device
59:     dev=_vx_pv_`echo $device | /usr/bin/cut -d/ -f5`"
60: else
```

```

61:     exit 0
62: fi
63: post_file="/etc/opt/FJSVswsrp/"$SWSTGNODE"/data/DEFAULT/"$dev".dpre"
64:
65: err_log_path="/var/opt/FJSVswsrp/"$SWSTGNODE"/log"
66:
67: # Confirmation of postprocessing
68: if [ ! -r $post_file ]
69: then
70:     exit 0
71: fi
72: post="/usr/bin/cat $post_file | /usr/bin/cut -d',' -f1`"
73: mount_point="/usr/bin/cat $post_file | /usr/bin/cut -d',' -f2`"
74:
75: # Confirmation of FStype
76: fs="/usr/sbin/fstyp $device 2> /dev/null`"
77: if [ $? != 0 ]
78: then
79:     fs=""
80: fi
81:
82: # When source device cannot be unmounted, fsck for destination device is needed
83: #
84: #if [ "$device" = "/dev/dsk/cXtXdXsX" -o "$device" = "/dev/dsk/cYtYdYsY" ]
85: #then
86: #   fsck_dev="`echo $device | sed "s/\dsk\\/\rdsk\\/"`"
87: #
88: #   if [ "$fs" = "" ]
89: #   then
90: #       fsck -m $fsck_dev > /dev/null 2>&1
91: #   else
92: #       if [ "$fs" = "sfxfs" ]
93: #       then
94: #           sfxadm $fsck_dev > /dev/null 2>&1
95: #           fsck -y -F $fs -o p,log $fsck_dev > /dev/null 2>&1
96: #       else
97: #           fsck -m -F $fs $fsck_dev > /dev/null 2>&1
98: #       fi
99: #   fi
100: #   if [ $? != 0 ]
101: #   then
102: #       if [ "$fs" = "" ]
103: #       then
104: #           result="`fsck -n $fsck_dev |grep \?`"
105: #       else
106: #           result="`fsck -n -F $fs $fsck_dev |grep \?`"
107: #       fi
108: #       if [ "$result" = "FILE SYSTEM STATE IN SUPERBLOCK IS WRONG; FIX? no" ]
109: #       then
110: #           if [ "$fs" = "" ]
111: #           then
112: #               fsck -y $fsck_dev > /dev/null 2>&1
113: #           else
114: #               fsck -y -F $fs $fsck_dev > /dev/null 2>&1
115: #           fi
116: #           if [ $? != 0 ]
117: #           then
118: #               if [ "$fs" = "" ]
119: #               then
120: #                   fsck -y $fsck_dev > $err_log_path/$dev.fsck 2>&1
121: #               else
122: #                   fsck -y -F $fs $fsck_dev > $err_log_path/$dev.fsck 2>&1
123: #               fi

```

```

124: #         if [ $? != 0 ]
125: #         then
126: #             exit 12
127: #         else
128: #             rm -f $err_log_path/$dev.fsck
129: #         fi
130: #     fi
131: # else
132: #     echo "$result" > $err_log_path/$dev.fsck
133: #     exit 12
134: # fi
135: # fi
136: #fi
137:
138: # No processing
139: if [ "$post" = "none" ]
140: then
141:     /usr/bin/rm -rf $post_file 2> /dev/null
142:     exit 0
143: fi
144:
145: # mount processing
146: if [ "$post" = "mount" ]
147: then
148:     /usr/bin/df -l $device > /dev/null 2>&1
149:     if [ $? != 0 ]
150:     then
151:         if [ "$fs" = "" ]
152:         then
153:             /usr/sbin/mount $device $mount_point 2> /dev/null
154:         else
155:             if [ "$fs" = "sfxfs" ]
156:             then
157:                 dev=`echo $device | /usr/bin/sed "s/\dsk\\/\rdsk\\/"`
158:                 /usr/sbin/sfxadm $dev > /dev/null 2>&1
159:             fi
160:             /usr/sbin/mount -F $fs $device $mount_point 2> /dev/null
161:         fi
162:         if [ $? != 0 ]
163:         then
164:             retry_count=3
165:             sleep_time=1
166:             result_flag=1
167:
168:             while [ $retry_count -gt 0 ]
169:             do
170:                 if [ "$fs" = "" ]
171:                 then
172:                     /usr/sbin/mount $device $mount_point > $err_log_path/$dev.mount 2>&1
173:                 else
174:                     if [ "$fs" = "sfxfs" ]
175:                     then
176:                         dev=`echo $device | /usr/bin/sed "s/\dsk\\/\rdsk\\/"`
177:                         /usr/sbin/sfxadm $dev > $err_log_path/$dev.sfxadm 2>&1
178:                     fi
179:                     /usr/sbin/mount -F $fs $device $mount_point > $err_log_path/$dev.sfxadm 2>&1
180:                 fi
181:                 if [ $? != 0 ]
182:                 then
183:                     retry_count=`expr $retry_count - 1`
184:                     /usr/bin/sleep $sleep_time
185:                 else
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: /usr/bin/rm -rf $post_file 2> /dev/null
198: exit 0
199: fi
200:
201: exit 0

```

C.2.4 Customizing Shell Script

The method to customize the shell script for replication pre/post-processing is explained below.

When You Do Not Want to Unmount Replication Target Volume

To ensure the copy source volume is not unmounted, make the following modifications to the script.

- If you do not want to unmount the copy source volume

Use an editor to remove the comment character ("##") at lines 76 to 79, 83, and 119 in the copy source volume pre-processing script (RepSrc.pre), and specify the target copy source volume name in the if statement on line 76.

- If you do not want to unmount the copy destination volume for the copy source volume

Use an editor to remove the comment character ("##") on lines 84 to 136 in the copy destination volume post-processing script (RepDst.post), and specify the target copy destination volume name in the if statement on line 84. Ensure that the line numbers quoted above are not changed by the insertion of removal of any special processing statements used in the script. Ensure these changes are fully tested before applying the script to live backup processing.

This operation is possible only when the file system is one of the following:

- ufs
- SynfinityFILE (single-partition configuration)

The pre-processing and post-processing of the copy source volume in the customized script is described below.

Table C.3 Copy Source Volume Pre-processing and Post-processing by Customized Script

Copy Source Volume Status	Copy Source Pre-processing	Copy Source Post-processing
It is mounted.	Suppresses updating of the copy source volume	Cancels the suppressed updating of the copy source volume

The pre-processing and post-processing of the copy destination volume in the customized script is described below.

Table C.4 Copy Destination Volume Pre-processing and Post-processing by Customized Script

Copy Destination Post-processing
The consistency of the copy destination volume is checked.

Note

- Where the standard script is used, the caching of a file system and renewal of subsequent updates are inhibited by unmounting a file system before an Advanced Copy process using (EC/OPC), the integrity of the data of the entire file system is guaranteed. Although flushing the cache of a file system and renewal of subsequent updates are inhibited using lockfs-w, when customizing a script to prevent unmounting a file system, data integrity or file consistency cannot be guaranteed if the file is in Write Open mode

(the database space of Oracle which is supporting on-line backup).

Moreover, when the renewal of metadata is occurring, the data updates to the file system itself are not guaranteed, even with "locks" specified.

It is therefore critical to consider this and manage the integrity of data updates accordingly when you elect not to unmount a file system.

- If an SDX object is operated in slice units, do not specify a slice name (*/dev/sfdsk/<class name>/dsk/<volume name>:<DEVNAM value of sdxinfo>*) in Line 76 of the copy source volume pre-processing script (RepSrc.pre) and in Line 84 of the copy destination volume post-processing script (RepDst.post). Instead, specify a logical volume name (*/dev/sfdsk/<class name>/dsk/<volume name>*).

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.

The mount point resource relating to the Volume is immediately offline/online whilst there is a time delay before access is actually prevented when mount/unmount is used.

Note that once the offline/online command is issued there is 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 Unmount Processing of Source Volume Pre-processing Script (RepSrc.pre)

[lines 87 and 97]

(before change)	<code>/usr/sbin/umount \$mount_point</code>
(after change)	<code>/opt/VRTSvcs/bin/hares -offline resourceName -sys systemName</code>

[Example]

Adding Unmount Wait Processing to Source Volume Pre-processing Script (RepSrc.pre)

[line 116]

(addition)	<code>while /usr/sbin/df -l "\$device">/dev/null 2>&1; do ;; done</code>
------------	--

[Example]

Changing Unmount Processing of Destination Volume Pre-processing Script (RepDst.pre)

[lines 74 and 84]

(before change)	<code>/usr/sbin/umount \$mount_point</code>
(after change)	<code>/opt/VRTSvcs/bin/hares -offline resourceName -sys systemName</code>

[Example]

Adding Unmount Wait Processing to Destination Volume Pre-processing Script (RepDst.pre)

[line 103]

(addition)	<code>while /usr/sbin/df -l "\$device">/dev/null 2>&1; do ;; done</code>
------------	--

[Example]

Changing Mount Processing of Source Volume Post-processing Script (RepSrc.post)

[lines 99, 103, 105, 118, 122 and 124]

(before change)	<pre>/usr/sbin/mount \$lvname \$lv_mount_point or /usr/sbin/mount -F \$fs \$device \$mount_point</pre>
(after change)	<pre>/opt/VRTSvcs/bin/hares -online resourceName -sys systemName</pre>

[Example]

Adding Mount Wait Processing to Source Volume Post-processing Script (RepSrc.post)

[line 144]

(addition)	<pre>until /usr/sbin/df -l "\$device">/dev/null 2>&1; do ;; done</pre>
------------	--

[Example]

Changing Mount Processing of Destination Volume Post-processing Script (RepDst.post)

[lines 153, 160, 172 and 179]

(before change)	<pre>/usr/sbin/mount \$lvname \$lv_mount_point or /usr/sbin/mount -F \$fs \$device \$mount_point</pre>
(after change)	<pre>/opt/VRTSvcs/bin/hares -online resourceName -sys systemName</pre>

[Example]

Adding Mount Wait Processing to Destination Volume Post-processing Script (RepDst.post)

[line 197]

(addition)	<pre>until /usr/sbin/df -l "\$device">/dev/null 2>&1; do ;; done</pre>
------------	--


Appendix D System Volume Backup Operation

This appendix provides how AdvancedCopy Manager can be used to back up and restore system volumes.

D.1 Overview

In a SAN boot environment where the system volume is allocated to a machine in the ETERNUS Disk storage system, the system volume can be backed up and restored using the replication function.

Figure D.1 Overview of Backup and Restoration of System Volume



Note


- A hardware device that supports SAN boot is required.
- The backup target server must be in the stopped state to back up or restore the system volume. This is to ensure the consistency between the data in the backup process and the backed up data.

D.2 Preparation

The server must be stopped in order to back up or restore its system volume, so a separate server is required for backup operations. That is, at least two servers are required for this operation.

A configuration example for backup and restoration of a system volume is shown below.

Figure D.2 Configuration Example for Backup and Restoration of System Volume



Point

If the backup source server and the server used for backup operations have different operating systems, refer to "10.1.1.12 Notes on Executing Copy Between Different Operating Systems" in "10.1.1 General Notes".

D.2.1 Registration Managed Server

The server being backed up and the server used for backup operations must be registered as Managed Servers.

Refer to "6.4.3 Registering Managed Server" to register Managed Servers.

D.2.2 Registration of Device Information

The system volume on the server being backed up and the device information of the backup volume on the server used for backup operations must be registered.

Refer to "6.4.4 Fetching Device Information on Managed Server" to register the device information.

Note

To back up a system volume that is a PRIMECLUSTER GD logical volume, the SDX object operation unit on the server must be slices.

D.2.3 Settings of Replication Volume Information

In replication volume information, the system volume to be backed up must be set as the source volume and the backup destination volume as the replica volume.

Refer to "6.4.5 Configuring Source Volume and Replica Volume" to set the replication volume information.

Note

- If the system volume is allocated to a slice that includes VTOC, check the notes in "10.1.1.6 Slice Containing VTOC" under "10.1.1 General Notes".
- If the system volume operates as a Logical Unit (disk unit), refer to "10.1.1.11 Notes on Executing Logical Unit (Disk) Unit Copy" under "10.1.1 General Notes".

Point

Register the backup volume as a volume without a mount point.

Specify the server used for backup operations as the replication destination.

```
# /opt/FJSVswsrp/bin/swsrpsetvol -o REP SYSVOL@TRG-SV SBKVOL@SBK-SV
```

D.3 Backing up System Volume

The replication function is used to back up the system volume.


Stop the server that is to be backed up before performing backup.

Note

If backup is performed while the backup target server is running, the backup may take place while a data update is in progress and without maintaining data consistency. Even though backup is completed, the correct data may not have been backed up.

If the incorrect backed up data is restored, the system may not operate correctly.

Figure D.3 Backing up System Volume



The backup procedure is as follows:

1. Stop the backup target server.
After the server stops, check that the server power is off.
2. Execute the `srsrpmake` command on the server used for backup.

```
# /opt/FJSVswsrp/bin/srsrpmake -m SYSVOL@TRG-SV SBKVOL@SBK-SV
```


(*) The command operand was specified in the format "volume name@Managed Server name", as shown in the examples below.

- name of backup/restore target server: TRG-SV
 - name of system volume: SYSVOL
 - name of server that performs the backup/restore operation: SBK-SV
 - name of backup volume: SBKVOL
3. Activate the backup target server.

D.4 Restoring System Volume

The replication function is used to restore a system volume.

Figure D.4 Restoring System Volume



The restoration procedure is as follows:

1. Stop the restoration target server.
After the server stops, check that the server power is off.
2. If the system volume on the restoration target server is unusable because of a hardware failure, recover the system volume so that the data can be restored.
3. On the server used for restore operations, execute the `srsrpmake` command with replication in the reverse direction.

```
# /opt/FJSSVswsrp/bin/swsrpmake -m SBKVOL@SBK-SV SYSVOL@TRG-SV
```

(*) The command operand was specified in the format "volume name@Managed Server name", as shown in the examples below.

- name of backup/restore target server: TRG-SV
- name of system volume: SYSVOL
- name of server that performs the backup/restore operation: SBK-SV
- name of backup volume: SBKVOL

4. Activate the restoration target server.

Appendix E Backup Operations Using Power-Saving


Enabling Eco Mode on the ETERNUS Disk storage system can control ON/OFF of disk drive power (or spindle rotation) using the Storage Cruiser functions. For example, the Storage Cruiser commands can power ON a disk drive as backup destination at backup and power it OFF on completion of backup.

It is possible to save power consumption under the following operations:

- Backup management
 - Snapshot type high-speed backup (OPC) with no access to the backup volume for a long period of time.
 - Suspend/Resume functionality of synchronous type high-speed backup (EC) with a long suspend time.
- Replication management
 - Snapshot type high-speed replication (OPC) with no access to the replication destination volume for a long period of time.
 - Suspend/Resume functionality of synchronous type high-speed replication (EC, REC) with a long suspend time.

An overview of the system structure is as follows.

Figure E.1 Overview of System Structure



E.1 Environment Settings

This section describes how to configure the environment.

1. [E.1.1 Settings for ETERNUS Disk Storage System](#)
2. [E.1.2 Settings for AdvancedCopy Manager](#)
3. [E.1.3 Settings for Storage Cruiser](#)

E.1.1 Settings for ETERNUS Disk Storage System

Settings in Eco-mode


To use the Eco-mode, it is required to set the device Eco-mode, create the Eco-mode schedule, and then set the Eco-mode scheduling for the RAID group. For more details about the settings, refer to the *ETERNUS SF Storage Cruiser Operation Guide*. For details on how to configure an ETERNUS Disk storage system, refer to the ETERNUS Web GUI manual.

Creating RAID Group and Logical Unit


The Eco-mode is applied and the RAID group is created.

A RAID group that is to store the data to be backed up (hereafter called "backup source RAID group") and a RAID group that is to store the backup data (hereafter called "backup destination RAID group") are created. A Logical Unit is allocated to each RAID group. To effectively reduce power consumption, the backup destination RAID group sets a Logical Unit for each backup that is performed simultaneously. Other Logical Units should not be included in the backup destination RAID group.

The example below illustrates the procedure to simultaneously back up three Logical units (Backup operation A). There are no Logical Units other than the ones for backup operation A in the backup destination RAID group. Therefore, a disk drive can be stopped, except when a backup is running.




The example below illustrates the case where Logical Units for both backup operation B and for business operations are included in the backup destination RAID group, in addition to the Logical Units for backup operation A. As a result, although backup operation A has ended, the Logical Units for backup operation B and for business applications are required to be running.



Point

A backup destination RAID group must be created at each generation if a backup over multiple generations is performed.



Note

A Snap Data Volume and a Snap Data Pool should not be included in a backup destination RAID group. The RAID group disk drive which includes a Snap Data Volume or a Snap Data Pool does not reach a stop status.

E.1.2 Settings for AdvancedCopy Manager

AdvancedCopy Manager needs to be configured so that the RAID group for the Eco-mode is used.

Using Backup Operation

When assigning a backup volume, use a logical unit or partition (slice) that belongs to a RAID group for the Eco-mode.

Using Replication Operation

For the destination volumes, use a logical unit or partition (slice) that belongs to a RAID group.

E.1.3 Settings for Storage Cruiser

A Storage Cruiser command is used for the power-saving operation. Therefore, Storage Cruiser needs to be set. For details, refer to the Storage Cruiser manual.

E.2 Operation

This section describes the procedure for establishing power-saving operations. Perform all procedures described in this section on the Management Server.

1. [E.2.1 Start Disk Drive](#)
2. [E.2.2 Operate AdvancedCopy Manager](#)
3. [E.2.3 Stop Disk Drive](#)

E.2.1 Start Disk Drive

A RAID group disk drive is started using the storageadm spindle command of Storage Cruiser. It is possible to wait for the start time of the disk drive by specifying the -sync option.

Refer to "Command References" in the *ETERNUS SF Storage Cruiser Operation Guide* for information on the command.

E.2.2 Operate AdvancedCopy Manager

Backup, restoration and replication are performed using AdvancedCopy Manager commands. For details, refer to the description of each operation in this manual.

E.2.3 Stop Disk Drive

Stopping of a RAID group disk drive is scheduled using the disk control command of the Storage Cruiser. This command returns control after the end of the stopping schedule of the disk drive.

Refer to "Command References" in the *ETERNUS SF Storage Cruiser Operation Guide* for information on the command.

Note

Depending on the status of the back up destination RAID group, a disk drive may not stop immediately.

- While OPC physical copying

If the OPC physical copy is in progress on the disk drive, the disk drive stops after a certain period once the physical copy is completed.

- While the disk is being accessed

If the RAID group for backup destination is being accessed, for example when backing up to a tape, the drive stops after a certain period once the access is stopped.

Other than the above, there are cases that the disk does not stop. Refer to "storageadm spindle (Disk Control Command)" in the *ETERNUS SF Storage Cruiser Operation Guide* for information on the command.

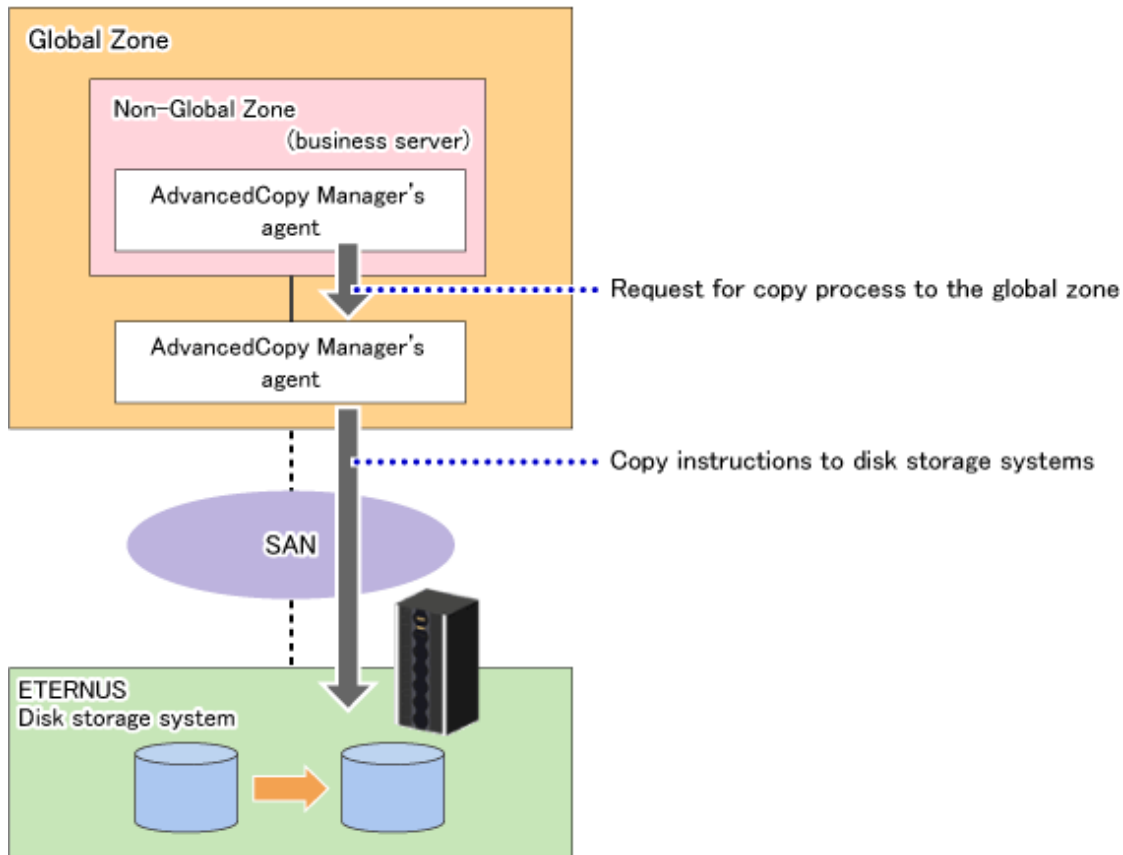
Appendix F Operation in Oracle Solaris Zones

This appendix explains operation using the Oracle Solaris Zones Support function.

F.1 Overview

The Oracle Solaris Zones Support function, via AdvancedCopy Manager's agents in the global zones, implements backup and restoration with transactions in the non-global zones.


Figure F.1 Overview of Oracle Solaris Zones Support Function



F.2 Configuration

This section shows a system configuration example using the Oracle Solaris Zones Support function.

Figure F.2 System Configuration Example Using Oracle Solaris Zones Support Function



Install AdvancedCopy Manager's agents on both of the following:

- Global zone
- Non-global zone performing backup or replication operations

 Note

The version of AdvancedCopy Manager's agent to be installed in the global zone must be equal to or later than the one in the non-global zone.

F.3 Software Requirements

The basic software requirements in the Oracle Solaris Zones are as shown below.

Refer to "1.3 Operating Environment of AdvancedCopy Manager" for the software requirements of the global zone.

- Management Server

Non-global Zone	Global Zone	
	Solaris 10 (*1)	Solaris 11 (11/11 or Later)
Oracle Solaris 10 Zones (*1)	Yes	Yes
Oracle Solaris 11 Zones	-	Yes

*1: Apply patch 141596-06.

- Managed Server

Non-global Zone	Global Zone	
	Solaris 10 (*1)	Solaris 11 (11/11 or Later)
Oracle Solaris 10 Zones (*1)	Yes	Yes
Oracle Solaris 11 Zones	-	Yes

*1: Apply patch 141596-06.

F.4 Available Functions

This section explains the functions that are available through the Oracle Solaris Zones Support function.

F.4.1 Functions of AdvancedCopy Manager

The following functions of AdvancedCopy Manager are available:

- AdvancedCopy functions
- Backup
- Replication

Out of the AdvancedCopy functions, the following types of copying are available:

- OPC
- QuickOPC
- EC
- REC
- SnapOPC
- SnapOPC+



See

.....
 For details on the functions of AdvancedCopy Manager, refer to "AdvancedCopy Manager Functions" in the *ETERNUS SF AdvancedCopy Manager Overview*.

F.4.2 Cluster Operation

Cluster operation using PRIMECLUSTER is supported by Oracle Solaris Zones Support function.

Cluster operation other than with PRIMECLUSTER is currently not supported.

F.5 Supported Devices

This section explains devices that can use the Oracle Solaris Zones Support function.

The following device names can be specified in AdvancedCopy Manager:


Device Type	Specifiable Device Name
Normal device	The sd driver resources, MPxIO driver resources, and MPLB resources (*1) can be specified. For details on device names, refer to "1.4.1 Normal Devices".
SDX objects of PRIMECLUSTER GD	Names of devices that are operating in units of slices can be specified. The format used to specify a device name is as follows: <code>/dev/sfdsk/<class-name>/dsk/<volume-name>:<DEVNAM-value-of-sdxinfo></code> Example: <code>/dev/sfdsk/class1/dsk/volume1:c1t0d10</code> For details on device names, refer to "1.4.2 SDX Objects of PRIMECLUSTER GD/SynfinityDISK". Names of devices that are operating in units of logical volumes are not supported.

*1: MPLB resources can be specified only if Oracle Solaris 11 Zones are operating in the non-global zone.

F.6 Flow of Operations

The following diagram shows the flow of operations using the Oracle Solaris Zones Support function.

Figure F.3 Flow of Operations Using Oracle Solaris Zones Support Function



The following control units and processing servers are used for above procedures:

No.	Procedure	Control Unit	Processing Server
1	"F.7.1 Setting up Non-global Zones"	All global zones	All non-global zones
2	"F.7.2 Starting Daemons"	- Management Server - All global and non-global zones	- Management Server - All global and non-global zones
3	"F.7.3 Starting Web Console"	PC	PC
4	"F.7.4 Registering Managed Server"	- Web Console (when using the GUI) - Management Server (when using the CLI)	All global and non-global zones
5	"F.7.5 Setting Environment for Non-global Zones"	All non-global zones	All non-global zones
6	"F.7.6 Setting of Virtual Environment Support Function of Managed Servers"	All non-global zones	All non-global zones
7	"F.7.7 Fetching Device Information from Managed Servers"	- Web Console (when using the GUI) - Management Server (when using the CLI)	All non-global zones
8	"F.7.8 Other Preparations"	- Web Console (when using the GUI) - Management Server (when using the CLI)	All non-global zones

F.7 Preparations

This section explains the preparations required for using the Oracle Solaris Zones Support function.

F.7.1 Setting up Non-global Zones

Set up non-global zones.



See

For details on configuration procedures, refer to the operating system manuals.

In case of cluster operation, refer to the PRIMECLUSTER manuals.

If you use the MPLB resources, refer to the ETERNUS Multipath Driver manuals.

Configuration of Non-global Zones

Use the `zonecfg` command to configure the non-global zones.

[Example]

Create the configuration of a non-global zone (zone01).

```
# zonecfg -z zone01<Enter>
```

Device Assignment

Use "add device" to assign the block devices and raw devices used by AdvancedCopy Manager to the non-global zone.

[Example]

Assign `/dev/dsk/c0t1d0s0` and `/dev/rdisk/c0t1d0s0` to a non-global zone (zone01).


```

zonecfg :zone01> add device <Enter>
zonecfg :zone01:device> set match=/dev/rdisk/c0t1d0s0 <Enter>
zonecfg :zone01:device> end <Enter>
zonecfg :zone01> add device <Enter>
zonecfg :zone01:device> set match=/dev/dsk/c0t1d0s0 <Enter>
zonecfg :zone01:device> end <Enter>

```

Table F.1 Coverage of Oracle Solaris Zones Support Function

		Devices and File Systems in Non-global Zones				
		ZFS File Systems	ZFS Storage Pools	ZFS Volumes	UFS File Systems	Raw Devices
Devices and file systems in global zones (zonecfg specification method)	ZFS file systems (add fs)	No	-	-	-	-
	ZFS storage pools (add dataset)	No	No (*1)	-	-	-
	ZFS volumes (add device)	-	-	No	-	-
	UFS file systems (add fs)	-	-	-	No	-
	RAW devices (add device)	-	-	-	Yes (*2)	Yes

Yes: Supported
No: Not supported
-: Non-existent combination

*1: Backup and restoration of ZFS is implemented in global zones. For details, refer to "[Appendix G Backup and Restoration of ZFS File Systems](#)".

*2: Raw devices in non-global zones are formatted as UFS file systems.

Note

- For using assigned devices as file systems, do one of the following:
 - Mount file systems in non-global zones manually.
 - Allocate the devices in /etc/vfstab in order to have the file systems mount them when booting the non-global zones.
- If you do not make settings for non-global zones, the display of fetched device information does not reflect devices in non-global zones.
- For using PRIMECLUSTER GD volumes, use the following method for assigning devices. For details, refer to the PRIMECLUSTER GD manuals.
 - Method of using volumes as raw devices from non-global zones

Setting of fs-allowed Properties

Perform the following operation only if the operating system of global zone is Solaris 11. For mounting in non-global zones, set the fs-allowed properties of the file systems that you use when configuring the non-global zones.

[Example]

Set ufs in the fs-allowed properties.

```
zonecfg:zone01> set fs-allowed=ufs
```

Reflection and Saving of Configuration Modification

Reflect and save the configuration modification of non-global zones.

[Example]

```
zonecfg:zone01> verify
zonecfg:zone01> commit
zonecfg:zone01> exit
```

Installation of Non-global Zones

If you install non-global zones by migrating an existing system that already has AdvancedCopy Manager's agents installed, uninstall them temporarily and then install them again.

F.7.2 Starting Daemons

Refer to "2.2 Starting Daemons" to start the daemons.

F.7.3 Starting Web Console

Refer to "How to start Web Console" in the *ETERNUS SF Web Console Guide* to start the Web Console.

F.7.4 Registering Managed Server

Register all global and non-global zones to be used with the Oracle Solaris Zones Support function as Managed Servers.

This process is carried out using the Web Console.



Note

- Change to the value of 800 characters or less before registering the Managed Server when the value of the environment variable PATH of the Management Server is more than 800 characters.
- Perform this task after confirming that it is possible to communicate from the Managed Server to the IP address of the Management Server. Check the Management Server IP address with the [stgxfwcmdisprv](#) command.

F.7.5 Setting Environment for Non-global Zones

Set up the environment for each non-global zone according to the operation.

For Single Operation and Cluster Operation (Except Cold Standby Operation of PRIMECLUSTER)

Create the following settings file in all non-global zones.

Table F.2 Settings File for Oracle Solaris Zones Support Function

Settings File	Description	File Path
Host information settings file	This is the file for setting the IP addresses and port numbers of the communication daemons of the Managed Servers (global zones) to which copying instructions are sent from non-global zones.	/etc/opt/swstorage/vstxapi_hostacm.ini

[Example of host information settings file]

```
ipaddr=10.10.10.22
port=1226
```

For details on the settings file, refer to "F.11 Specifications of Settings File".

For Cluster operation (Cold Standby Operation of PRIMECLUSTER)

Create host information settings files in the global zones for shared use with the non-global zones.

Table F.3 Settings File for Oracle Solaris Zones Support Function

Settings File	Description	File Path
Host information settings file	This is the file for setting the IP addresses and port numbers of the communication daemons of the Managed Servers (global zones) to which copying instructions are sent from non-global zones.	/etc/opt/swstorage/vstxapi_hostacm.ini

[Example]

```
# zonecfg -z zone01<Enter>
zonecfg:zone01> add fs
zonecfg:zone01:fs> set type=lofs
zonecfg:zone01:fs> set options=ro
zonecfg:zone01:fs> set special=/etc/opt/swstorage/vstxapi_hostacm.ini
zonecfg:zone01:fs> set dir=/etc/opt/swstorage/vstxapi_hostacm.ini
zonecfg:zone01:fs> end
zonecfg:zone01> verify
zonecfg:zone01> commit
zonecfg:zone01> exit
```

For details on the settings file, refer to "F.11 Specifications of Settings File".

F.7.6 Setting of Virtual Environment Support Function of Managed Servers

Execute the `stgxfwcmsetmode` command to enable the Oracle Solaris Zones Support function in all non-global zones.

In case of cluster operation, before running the command, specify the logical node name in the environment variable SWSTGNODE.

[Example]

```
# export SWSTGNODE=nodeagt
# /opt/FJSVswstf/bin/stgxfwcmsetmode -virt zone
stgxfwcmsetmode completed (SWSTGNODE=nodeagt)
#
```

F.7.7 Fetching Device Information from Managed Servers

Fetch the device information from the Managed Servers (non-global zones).

This procedure need not be performed for Managed Servers (global zones).

For operational procedures, refer to the following sections, depending on the operation.

- For backing up normal transaction volumes: "3.4.4 Fetching Device Information on Managed Server"
- For replicating: "6.4.4 Fetching Device Information on Managed Server"



After fetching the device information, confirm that the fetched devices are imported as the devices in the ETERNUS Disk storage systems. This operation can be performed with Web Console or the `stgxfwcmdisppat` command. If the devices are not imported, there may be errors in the host information setting file. Check the contents of the file. If errors are found, correct them and then fetch the device information on the Managed Servers again.

F.7.8 Other Preparations

Make the necessary preparations depending on the operation.

For operational procedures, refer to the following sections:

- For backing up normal transaction volumes: "[3.4.5 Configuring Environment for Backup Operation Server](#)" and subsequent sections
- For replicating: "[6.4.5 Configuring Source Volume and Replica Volume](#)" and subsequent sections

F.8 Operation

This section explains operation of AdvancedCopy Manager using the Oracle Solaris Zones Support function.

Refer to the following sections, depending on the method of operation. Note, however, that the preparations described in these sections need not be implemented.

- For backing up normal transaction volumes: "[3.5 Operation](#)"
- For replicating: "[6.5 Operations](#)"



When you use the Oracle Solaris Zones Support function, data communication with the ETERNUS Disk storage system is carried out via the global zones. Therefore, depending on the number of volumes, processing takes more time than in physical server environments.

F.9 Changing Operating Environment

This section explains how to modify the operating environment when using the Oracle Solaris Zones Support function. For other procedures to change the operating environment, refer to "[8.5 Changing Operating Environment](#)".

- "[F.9.1 Changing IP Addresses of Managed Servers](#)"
- "[F.9.2 Changing Port Numbers of Managed Servers](#)"

F.9.1 Changing IP Addresses of Managed Servers

Referring to "[8.5.2 Changing IP Address of Managed Server](#)", modify the IP addresses of the Managed Servers.

For modifying IP addresses of Managed Servers (global zones), refer to "[F.7.5 Setting Environment for Non-global Zones](#)", and also change the IP addresses that are set in the host information settings files.

F.9.2 Changing Port Numbers of Managed Servers

Referring to "[8.5.3 Changing Port Number of Management Server or Managed Server](#)", modify the port numbers of the Managed Servers.

For modifying port numbers of Managed Servers (global zones), refer to "[F.7.5 Setting Environment for Non-global Zones](#)", and also change the port numbers that are set in the host information settings files.

F.10 Available Commands

The commands that are available with the Oracle Solaris Zones Support function are as follows:

- "[9.1 Backup Management Commands](#)"
- "[9.2 Configuration Management Commands](#)"
- "[9.3 Replication Management Commands](#)"
- "[9.4 Start/Stop Daemon Commands](#)"
- "[9.5 SnapOPC/SnapOPC+ Support Commands](#)"

F.11 Specifications of Settings File

This section explains the settings file that is used with the Oracle Solaris Zones Support function.

F.11.1 Host Information Settings File (vstxapi_hostacm.ini)

The host information settings file is the file for setting the IP addresses and port numbers of the communication daemons of the Managed Servers (global zones) to which AdvancedCopy instructions are sent. Creation and editing of the file are carried out by the user.

Create the host information settings file on the following path.

[Path for host information settings file]

```
/etc/opt/swstorage/vstxapi_hostacm.ini
```

[Contents of host information settings file]

```
ipaddr=xxx.xxx.xxx.xxx  
port=xxxxxx
```

Parameter	Value to Be Set
- IPv4 format ipaddr=xxx.xxx.xxx.xxx - IPv6 format ipaddr=n:n:n:n:n:n (*1)	IP address of a communication daemon of the Managed Server (global zone) to which AdvancedCopy instructions are sent One IP address in either IPv4 format or IPv6 format can be specified.
port=xxxxxx	Port number (1 - 65535) of a communication daemon of the Managed Server (global zone) to which AdvancedCopy instructions are sent

*1: Full notation and abbreviated notation are available. Refer to "1.6.5 IPv6 Address Notation" for more information.


Appendix G Backup and Restoration of ZFS File Systems

This appendix explains backup and restoration of ZFS file systems.

G.1 Overview


In AdvancedCopy Manager, backup and replication of ZFS file systems can be implemented with the replication function or AdvancedCopy Manager CCM.

Figure G.1 Scheme of Backup and Restoration of ZFS File Systems



G.1.1 Units of Backing up

AdvancedCopy Manager supports backup and restoration in units of ZFS storage pools (all file systems within a storage pool). Backup and restoration in units of ZFS file systems are not supported.



After registering all devices within a storage pool in AdvancedCopy Manager, you can back up and restore them in a batch.

For details on how to register to AdvancedCopy Manager, refer to "[G.2.3 Preparations](#)", "[G.3.3 Preparations](#)" or "[G.4.3 Preparations](#)".

For details on backup and restoration methods, refer to "[G.2.4 Operation](#)", "[G.3.4 Operation](#)" or "[G.4.4 Operation](#)".

Also, for details on supported configurations of ZFS storage pools, refer to "[G.2.1 Operation Design](#)", "[G.3.1 Operation Design](#)" or "[G.4.1 Operation Design](#)".

Note

- AdvancedCopy Manager is controlled in units of devices constituting a ZFS storage pool. In order to maintain consistency as a ZFS storage pool, control all devices in the storage pool synchronously.
- If you are encrypting your ZFS file systems, you should also preserve (back up) the encryption keys as necessary.
- Before you start to design and operate backup and restoration of ZFS storage pools, make sure to sufficiently understand the design and operation of each device type.

Point

For backing up ZFS with AdvancedCopy Manager, you should divide the storage pool into units to be backed up. For example, you could divide a storage pool by transactions to be backed up with a different timing.

G.1.2 Backup Methods

With AdvancedCopy Manager, the following methods are supported for ZFS backup and restoration.


Select them according to your particular backup purpose.

- Online backup using the ZFS Snapshot function (hereafter referred to as "ZFS Snapshot") and AdvancedCopy Manager function. (recommendation)
- Online backup using the ZFS Snapshot and ZFS uses the zfs send and zfs receive commands.
- Offline backup using the export and import functions of ZFS (system stop or disconnection of storage pool)

Online Backup Using ZFS Snapshot with AdvancedCopy Manager Function

This backup method uses ZFS Snapshot.

Using ZFS Snapshot, backups can even be performed while the ZFS file system is operating, but it is necessary to ensure the consistency of data in each operation when backing up or restoring them. Since there are different ways to ensure consistency, depending on your applications and types of data (DB etc.), you have to investigate them for each type of operation.



Note

- Online backup of the system area is not supported. For the system area, perform offline backup.
- Log device (ZIL) and cache device (L2ARC) are not supported. As for log device and cache device, perform offline backup.

Online Backup Using ZFS Snapshot with ZFS Uses zfs send and zfs receive Commands

This backup method uses ZFS Snapshot.

Backups are performed by combining ZFS Snapshot with the `zfs send` and `zfs receive` commands. With consideration for the effects on transactions, we recommend that you implement a combination of periodical (e.g. daily) differential backups with full backups at constant intervals (weekly or monthly).

Using ZFS Snapshot, backups can even be performed while the ZFS file system is operating, but it is necessary to ensure the consistency of data in each operation when backing up or restoring them. Since there are different ways to ensure consistency, depending on your applications and types of data (DB etc.), you have to investigate them for each type of operation.

- Differences as compared to online backup using the ZFS Snapshot with AdvancedCopy Manager function

Online backup using the ZFS Snapshot with ZFS uses the `zfs send` and `zfs receive` commands of ZFS uses the `zfs send` and `zfs receive` commands to back up the data captured by the snapshot onto the backup server. Therefore, compared to online backup using the

ZFS Snapshot with AdvancedCopy Manager function, it has characteristics as stated below. When online backup, online backup using the ZFS Snapshot with AdvancedCopy Manager function is recommended.

- Depending on the amount of backup data, processing takes more time.
- Depending on the amount of backup data, a higher burden is placed on CPUs and disks (and on the network, if you back up onto another server).

Note

Online backup of the system area is not supported. For the system area, perform offline backup.

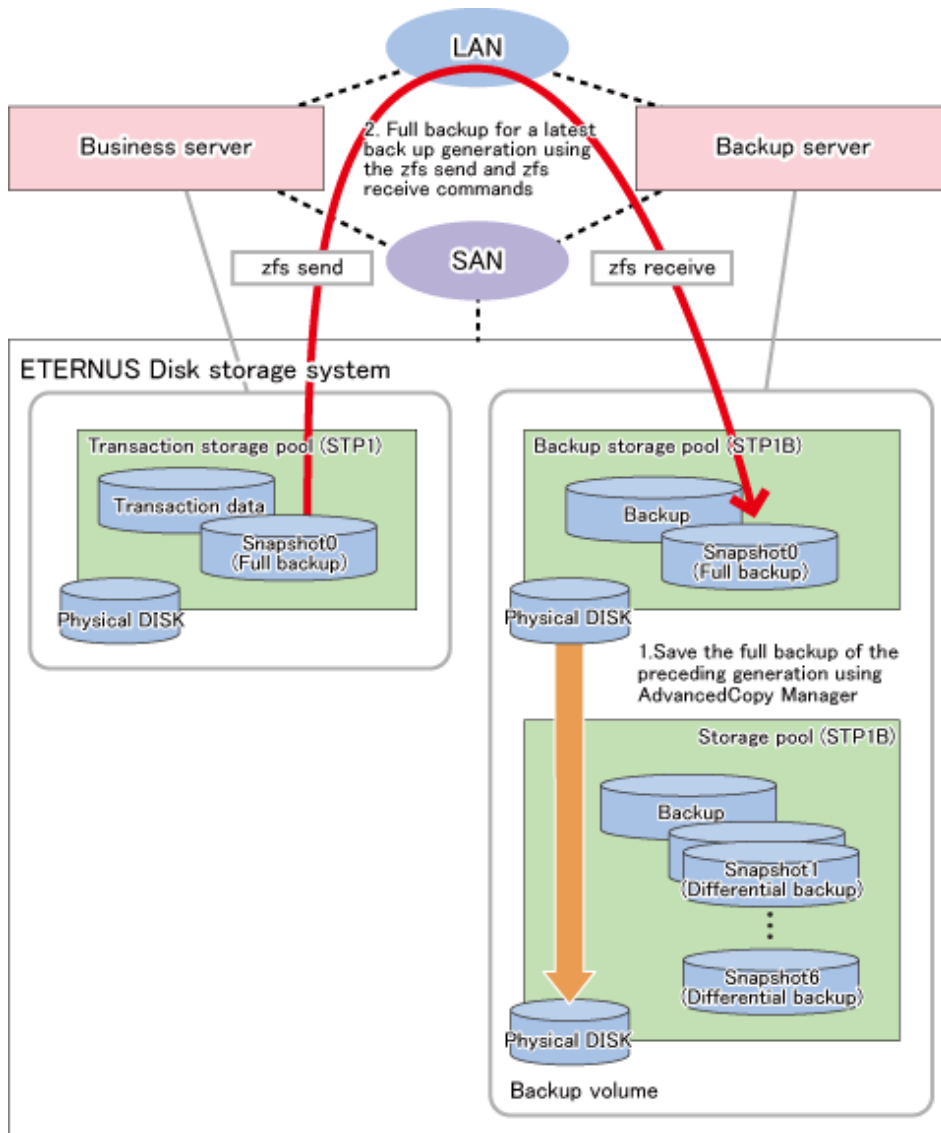
Point

Since differential backup processes only the portions that are different from the previous snapshot, the backup amount depends on the amount of updated data.

Differential backups are backed up to the same backup storage pools as full backups. In the event of failure of a physical disk that constitutes a backup storage pool, all differential backups in the same backup storage pool are affected. Therefore, we recommend you to implement full backups periodically.


Full Backup

This method backs up entire snapshots. It is implemented for the first-time backup and for backups at constant intervals (weekly or monthly).



Differential Backup

This method backs up only the updated data that are different from the previous differential or full backup.



Offline Backup

This backup method uses the export and import functions of ZFS storage pools.

In order to maintain consistency of storage pools, you should implement backups and restorations with AdvancedCopy Manager offline. Depending on the purpose (system area or data area) of the ZFS storage pool to be backed up, implement the following measures:

- For backup sources in the system area
Stop the backup source servers before you start backup or restoration.
For details on backup and restoration procedures, refer to "[Backup and Restoration of System Areas](#)".
- For backup sources in the data area
Disconnect the storage pools before you start backup or restoration.
For details on backup and restoration procedures, refer to "[Backup and Restoration of Data Areas](#)".



Note

AdvancedCopy Manager is controlled in units of devices constituting a ZFS storage pool. In order to maintain consistency as a ZFS storage pool, control all devices in the storage pool synchronously.

G.1.3 Functions of AdvancedCopy Manager

The following functions of AdvancedCopy Manager are available:

- Replication
- AdvancedCopy Manager CCM

The following types of copying are available:

- OPC

- QuickOPC
- EC
- REC



See

For details on the functions of AdvancedCopy Manager, refer to "AdvancedCopy Manager Functions" in the *ETERNUS SF AdvancedCopy Manager Overview*.

G.1.4 Software Requirements

The basic software requirements for backup and restoration of ZFS file systems are as shown below.

Server Type	Basic Software
Business server (*1)	Solaris 10 (*2)
Backup server (*1)	Solaris 11

*1: The versions of the operating systems and the release versions of business servers and backup servers must be the same.

*2: For online backup in Solaris 10 environments, version 9/10 or later is required.

G.2 Online Backup Using ZFS Snapshot with AdvancedCopy Manager Function

This section explains methods for implementing online backup.

G.2.1 Operation Design

When you design your storage pools and backup operation, pay attention to the following aspects:

- "Unit of Backup and Restoration: Storage Pool"
- "Version of ZFS Storage Pools"
- "Types and Configurations of ZFS Storage Pools"
- "Preparing Backup Volume"

Unit of Backup and Restoration: Storage Pool

AdvancedCopy Manager supports backup and restoration only in units of ZFS storage pools (all file systems within a storage pool). Backup and restoration in units of ZFS file systems are not available.

For details, refer to "G.1.1 Units of Backing up".

Version of ZFS Storage Pools

For implementing online backup, ZFS storage pool version of business servers and that of backup servers must be the same. For checking your ZFS storage pool version, refer to the operating system manuals.

[Example]

```
# zpool upgrade -v
This system is currently running ZFS pool version 33.
```

Types and Configurations of ZFS Storage Pools

Purposes of ZFS Storage Pools: System Area or Data Area

With online backup, you can back up and restore data areas (usual storage pools). For backing up system areas (ZFS root pools), you have to use offline backup.

Devices Constituting ZFS Storage Pool

Online backup supports the following as device that constitutes ZFS storage pools (Physical disks).

Device Type

Normal devices are supported.



For details on using replication management function, refer to "[1.4 Managing Device on AdvancedCopy Manager](#)".

Device Unit

Entire disks (LUs) are supported.



As is constituted with the entire disks (LUs), create storage pools.

Preparing Backup Storage Pool for Backup

As backup targets in online backup, you need to arrange for the following two types of backup targets:

- Backup storage pool for backup
- Backup volume

Backup Storage Pool for Backup

A "backup storage pool" serves to store of the latest data generations.

Make sure to arrange for one on the backup server.

The backup storage pool must have the same physical disk configuration (number of disks, sizes, redundant configuration) as the ZFS storage pool that is the backup source. Also, unit of devices that constitute storage pools must be the same.

Backup Volume

A "backup volume" is a backup volume for storing the storage pools.

Make sure to arrange for the number of backup generations for the storage pools to be preserved on the backup server.

A backup volume must fulfill the conditions and have a configuration as described below.

For any conditions that are not stated below, depending on the type of device in the backup source volume, refer to "[1.4 Managing Device on AdvancedCopy Manager](#)", "[Chapter 6 Replication](#)", or the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module*.

- Make sure that the backup volume has the same physical disk configuration (number of disks, sizes, and redundant configuration) as the ZFS storage pool that is the backup source.
If a ZFS storage pool consists of multiple physical disks, you have to use a backup volume whose physical disks all have the same size.
- If you are using AdvancedCopy Manager CCM, do not connect the backup volume to the server.
- When using replication management function, use backup volume as slices.

Combination of Transaction Storage Pools and Backup Volumes

The combinations of transaction storage pools, backup storage pools and backup volumes are described below.

Function	Device		
	Transaction Storage Pools	Backup Volume	
		Backup Storage Pools	Backup Volume
Replication management	LU	LU	Slice (Note)
AdvancedCopy Manager CCM	LU	LU	LU

Note: Define copy source as LU and copy destination as slice, then copy the entire LUs to the slice.

If a backup destination volume is LU, make the disk format of the transaction storage pool the same as that of the backup destination volume as shown below:

- If the transaction storage pool is a VTOC-labeled disk, the backup destination volume should also be a VTOC-labeled disk.
- If the transaction storage pool is an EFI-labeled disk, the backup destination volume should also be an EFI-labeled disk.

Connection Servers of Backup Target Volumes

The servers to which to connect the backup volumes, depending on the copying functions to be used, are as shown below.

Backup Target Volume	Copying Function	
	Replication Management	AdvancedCopy Manager CCM
Backup storage pool	Backup server	Backup server
Backup volume	Backup server	Do not connect to a server.

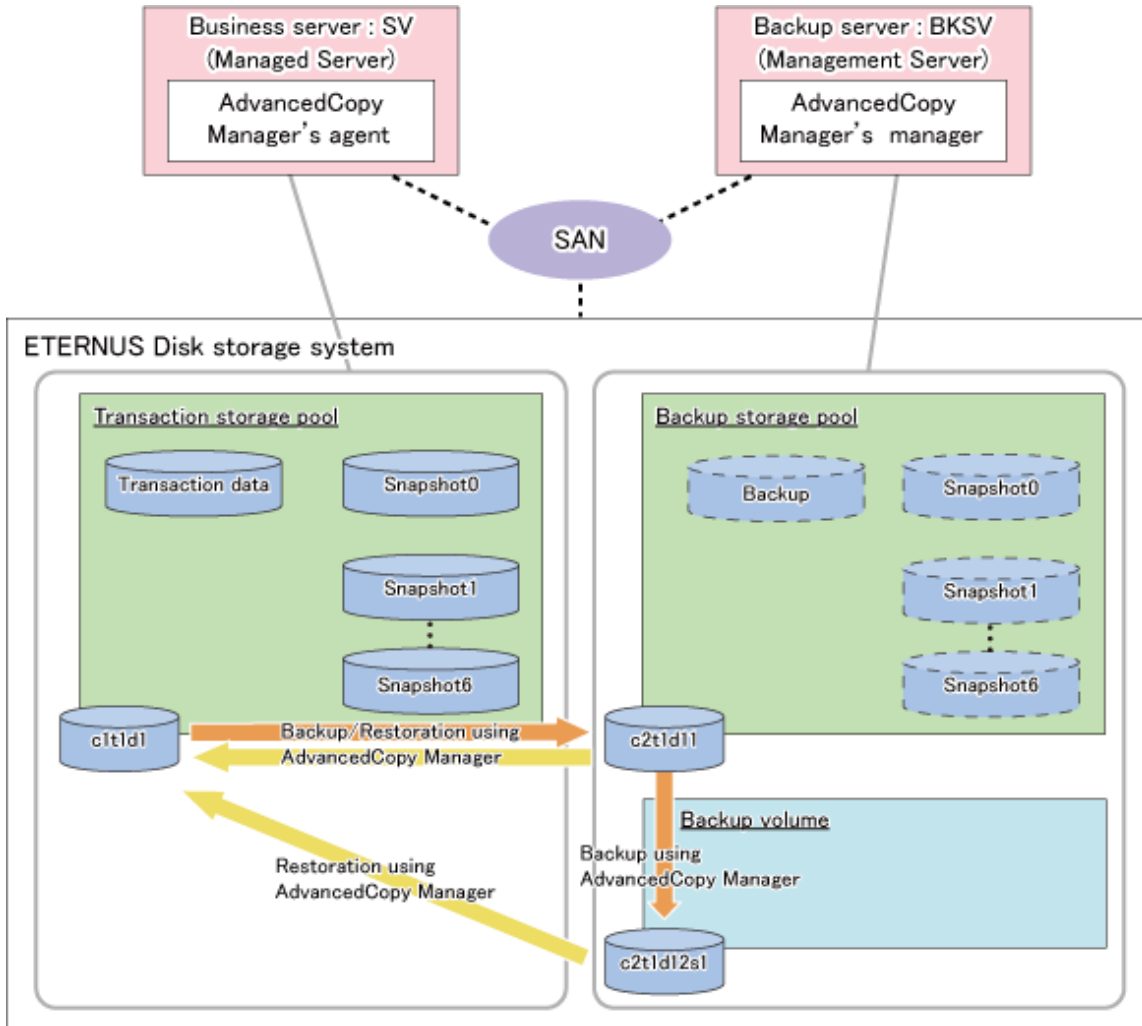
G.2.2 System Configuration

Since, in online backup operations, backup servers are used for backing up, backup servers are required that are separate from business servers. Consequently, you need a server configuration which consists of at least 2 servers.

Replication Operations

This section explains the system configuration for implementing replication operations.

Figure G.2 Sample Configuration for Replication Operations



Details on Sample Configuration

Using one physical disk that constitutes a storage pool for transactions, the entire disk is registered. Including the backup storage pool, backup can preserve a maximum of two generations.


- Business server (Managed Server): SV
- Backup server(Management Server): BKSV (name for registration in AdvancedCopy Manager)
- Transaction storage pool: STP1
- Backup storage pool: STP1
- Disk

Physical Disk	Connection Server	Device Type	Device Name	Purpose (Storage Pool)
Transaction VOL	SV	LU	c1t1d1	Transaction storage pool (STP1)
Backup VOL-1	BKSV	LU	c2t1d11	Backup storage pool (STP1)
Backup VOL-2	BKSV	Slice	c2t1d12s1	Backup volume

AdvancedCopy Manager CCM Operations

This section explains the system configuration for implementing AdvancedCopy Manager CCM operations.

Figure G.3 Sample Configuration for AdvancedCopy Manager CCM Operations



Details on Sample Configuration

Using one physical disk that constitutes a storage pool for transactions, the entire disk is registered.

Including the backup storage pool, backup can preserve a maximum of two generations.

- Business server: SV
- CCM Server: BKSV
- Transaction storage pool: STP1
- Backup storage pool: STP1
- Disk

Physical Disk	Connection Server	LUN	Device Type	Device Name	Purpose (Storage Pool)
Transaction VOL	SV	0x01	LU	c1t1d1	Transaction storage pool (STP1)
Backup VOL-1	BKSV	0x11	LU	c2t1d11	Backup storage pool (STP1)
Backup VOL-2	-	0x12	LU	-	Backup volume

G.2.3 Preparations

This section describes the preparations that are required for backing up and restoring ZFS storage pools.

For details on each procedure and a setup required in advance for using AdvancedCopy Manager, refer to "[Chapter 6 Replication](#)" and the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module*.

Whereas this document states some typical methods for controlling ZFS for your reference, refer to the operating system manuals for more details on usage.

Point

Below examples of command execution are based on the configurations described for replication operations and AdvancedCopy Manager CCM operations in "G.3.2 System Configuration". Adapt any real control actions to your actual configuration.

Implement the following procedures for preparation:

1. [Fetching Lists of Devices in Transaction Storage Pools](#)
2. [Copy Settings for Backup and Restoration](#)

Fetching Lists of Devices in Transaction Storage Pools

On the business servers, check the devices that constitute the transaction storage pools to be backed up.

[Example]

Check the disk information of the storage pool. In below example, storage pool (STP1) consists of 1 physical disk (c1t1d1).

```
# zpool status STP1
pool: STP1
state: ONLINE
scrub: none requested
config:
   NAME      STATE      READ WRITE CKSUM
   STP1      ONLINE      0   0   0
     c1t1d1  ONLINE      0   0   0

errors: No known data errors
#
```

Copy Settings for Backup and Restoration

Make the following three copy settings:

- Copy settings for the transaction storage pool and the backup storage pool (for backup and restoration)
- Copy settings for the backup storage pools and the backup volumes (for backup to backup volume)
- Copy settings for the transaction storage pools and the backup volumes (for restoration from backup volume)

Use the following combinations to register them in AdvancedCopy Manager:

Table G.1 Copy Groups

Used for	Copy Source	Copy Destination	Group Name (ex.)
Backup/Restoration	Transaction storage pool	Backup storage pool	STP1
Backup	Backup storage pool	Backup volume	BKSTP1
Restoration	Backup volume	Transaction storage pool	RSSTP1

Table G.2 Registered Devices

Device Name	Devices to Be Registered	Example (*1)	
		Device Name	LUN
Transaction storage pool	Devices that you confirmed in " Fetching Lists of Devices in Transaction Storage Pools " to be devices constituting the transaction storage pool	c1t1d1	0x01

Device Name	Devices to Be Registered	Example (*1)	
		Device Name	LUN
Backup storage pool	Devices that constitute the backup storage pool	c2t1d11	0x11
Backup volume	Devices you prepared as backup volumes for older generations	c2t1d12s1	0x12

*1: When using replication management function, use device name. When using AdvancedCopy Manager CCM, use LUN number.

For Using Replication Management Function

Register copy sources indicated at "Table G.1 Copy Groups" as copy destination volumes, and copy destinations as copy destination volumes. Then register the devices indicated at "Table G.2 Registered Devices" respectively.

Since these are controlled synchronously in processing for backup and restoration, you have to register all devices that constitute a storage pool as a single group. For details on controlling groups, refer to "6.4.7 Creating Groups".

[Example]

Specify the source volumes and copy destination volumes to create the 3 groups shown in the chart.

```
# /opt/FJSVswsrp/bin/swsrpsetvol -o BOTH -Xgroup STP1 /dev/dsk/c1t1d1@SV /dev/dsk/
c2t1d11@BKSV

swsrpsetvol completed
# /opt/FJSVswsrp/bin/swsrpsetvol -Xgroup BKSTP1 /dev/dsk/c2t1d11 /dev/dsk/c2t1d12s1
swsrpsetvol completed
# /opt/FJSVswsrp/bin/swsrpsetvol -o BOTH -Xgroup RSSTP1 /dev/dsk/c2t1d12s1@BKSV /dev/dsk/
c1t1d1@SV
swsrpsetvol completed
# /opt/FJSVswsrp/bin/swsrpvoinfo -L
Server Original-Volume      Size      Replica-Volume      Size      Copy
Op-Server  Group
BKSV  /dev/dsk/c1t1d1@SV      4.0 Gbyte  /dev/dsk/c2t1d11@BKSV  4.0 Gbyte  bi-
direction both          STP1
BKSV  /dev/dsk/c2t1d11@BKSV  4.0 Gbyte  /dev/dsk/c2t1d12s1@BKSV 4.0 Gbyte  bi-
direction original    BKSTP1
BKSV  /dev/dsk/c2t1d12s1@BKSV 4.0 Gbyte  /dev/dsk/c1t1d1@SV      4.0 Gbyte  bi-
direction both          RSSTP1
#
```

For Using AdvancedCopy Manager CCM

Register each combination of copy sources and copy destinations.

Since these are controlled synchronously in processing for backup and restoration, you have to register all devices that constitute a storage pool as a single group. For details on controlling groups, refer to "Creating Copy Groups" and "Adding Copy Pairs" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module*.

1. Creation of copy groups

Create a copy group for each copy setting.

[Example]

Create the 3 copy groups shown in the chart. (Copying type: OPC)

```
# acgroup create -g STP1 -type OPC -a DX440
Successful completion.
#
# acgroup create -g BKSTP1 -type OPC -a DX440
Successful completion.
#
# acgroup create -g RSSTP1 -type OPC -a DX440
```

```
Successful completion.  
#
```

2. Registration of copy pairs

In the created copy groups, register the combinations of copy sources and copy destinations. Register all devices that constitute a storage pool.

[Example]

In the created copy groups, register each device shown in the chart.

```
# acpair add -g STP1 -p DX440/0x01:DX440/0x11  
Successful completion.  
# acpair add -g BKSTP1 -p DX440/0x11:DX440/0x12  
Successful completion.  
# acpair add -g RSSTP1 -p DX440/0x12:DX440/0x01  
Successful completion.
```

G.2.4 Operation

This section describes procedures for backing up and restoring ZFS storage pools.

For details on the settings (device settings) in advance for using AdvancedCopy Manager, or on the backup and restoration procedures, refer to "[Chapter 6 Replication](#)" or the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module*.

Note that the procedures and commands shown here are examples of snapshot-type copying (OPC). For using other types of copying, replace the procedures and commands as needed for each type of copying.

Whereas this document states some typical methods for controlling ZFS for your reference, refer to the operating system manuals for more details on usage.

Backup

This section shows the procedure for online backup.

No.	Description	Controlling Server	At First Backup (*1)
1	Export the backup storage pool.	Backup server	
2	Using AdvancedCopy Manager, back up data to a backup volume from a device in the backup storage pool.	Backup server	
3	Create a snapshot of the transaction storage pool.	Business server	Yes
4	Using AdvancedCopy Manager, back up data to a device in the backup storage pool from a device in the transaction storage pool.	Business server	Yes
5	Check the status of the backup storage pool.	Backup server	Yes

*1: When you make a backup for the first time, only implement the processes indicated with "Yes".

The procedure for the online backup is as follows:

1. Export the backup storage pool.

For more information on export of ZFS storage pools, refer to the operating system manual.

[Example]

```
# zpool list  
NAME      SIZE  USED  AVAIL  CAP  HEALTH  ALTROOT  
STP1     345G  121M   345G   0%  ONLINE  -  
# zpool export STP1  
#
```

- Using AdvancedCopy Manager, back up the backup data on the backup storage pool to a backup volume.

Before storing backups in the backup storage pool, back up all the devices that compose the backup storage pool to a backup volume. If multiple backup volumes are used (to save backups of 2 or more generations), use any unused or oldest generation backup volumes for backups.

For Using Replication Management Function

Specify backup group "BKSTP1" for the -Xgroup option and execute the `swsrpmake` command.

[Example]

```
# /opt/FJSVswsrp/bin/swsrpmake -Xgroup BKSTP1
GROUP=BKSTP1 swsrpmake completed
#
```

For Using AdvancedCopy Manager CCM

Specify backup group "BKSTP1" for the -g option and execute the `acopc start` command for AdvancedCopy Manager CCM.

[Example]

```
# /opt/FJSVccm/bin/acopc start -g BKSTP1
DX440/0x11:DX440/0x12
# DATE : 2012/01/01 00:00:00 - << OPC Started >>
#
From:BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=11/Adr_high=0/Adr_low=0/size_high=0/size_low=0
#
To:BoxID=303045343030304d3323232323234534353053323041232323234b443430333036333930303
42323/Olu=12/Adr_high=0/Adr_low=0
Succeeded : 1
Failed : 0
#
```



Note

Since processing is carried out for all devices that constitute the storage pool, do not specify any copy pairs.

- Create snapshots of the transaction storage pool.
 - Delete the snapshot that is not restored to the transaction storage pool.
To delete ZFS snapshots, refer to the operating system manual.
 - Execute the `sync` command to synchronize the transaction storage pool with the file system.
 - Create a file system snapshot on the transaction storage pool.
To create ZFS snapshots, refer to the operating system manual.
 - Execute the `sync` command to synchronize the transaction storage pool with the file system.

[Example]

```
# zfs destroy STP1@SNAP0
# sync
# zfs snapshot -r STP1@SNAP0
# sync
#
```

- Using AdvancedCopy Manager, back up data to a device in the backup storage pool from a device in the transaction storage pool.

For Using Replication Management Function

Specify backup/restore group "STP1" for the -Xgroup option and execute the [swsrpmake](#) command.

[Example]

```
# /opt/FJSVswsrp/bin/swsrpmake -Xgroup STP1
GROUP=STP1 swsrpmake completed
#
```

For Using AdvancedCopy Manager CCM

Specify backup/restore group "STP1" for the -g option and execute the `acopc start` command for AdvancedCopy Manager CCM.

[Example]

```
# /opt/FJSVccm/bin/acopc start -g STP1
DX440/0x01:DX440/0x11
# DATE : 2012/01/01 00:00:00 - << OPC Started >>
#
From:BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=01/Adr_high=0/Adr_low=0/size_high=0/size_low=0
#
To:BoxID=303045343030304d3323232323234534353053323041232323234b443430333036333930303
42323/Olu=11/Adr_high=0/Adr_low=0
Succeeded : 1
Failed : 0
#
```

Note

Since processing is carried out for all devices that constitute the storage pool, do not specify any copy pairs.

5. Check the status of the backup storage pool.
 - a. Import the backup storage pool on the backup server.
 - b. Export the backup storage pool on the backup server.

For more information on import/export of the ZFS storage pool, refer to the operating system manual.

[Example]

```
# zpool import
pool: STP1
  id: 15271135801818262407
  state: ONLINE
action: The pool can be imported using its name or numeric identifier.
config:

          STP1          ONLINE
            c2t1d11    ONLINE
# zpool import -f STP1
# zpool list
NAME  SIZE  USED  AVAIL  CAP  HEALTH  ALTROOT
STP1  345G  345M  345G   0%  ONLINE  -
# zpool export STP1
#
```

Note

If import/export has failed, re-execute it from step 3.

Restoration

This section shows the procedure for restoration.

No.	Description	Controlling Server
1	Export the transaction storage pool.	Business server
2	When restoring data from a device in the backup storage pool, export the backup storage pool.	Backup server
3	Using AdvancedCopy Manager, restore data to a device in the transaction storage pool from a device in the backup storage pool or a backup volume.	Business server or Backup server
4	Import the transaction storage pool.	Business server
5	Roll back the transaction storage pool to a restored snapshot.	Business server
6	Mount the target ZFS file systems.	Business server

The procedure for restoration is as follows:

1. If any transaction storage pool exists, export it.

For more information on export of ZFS storage pools, refer to the operating system manual.

[Example]

```
# zpool list
NAME      SIZE  USED    AVAIL    CAP  HEALTH  ALTROOT
STP1     345G  345M    345G     0%   ONLINE  -
# zpool export STP1
#
```

2. When restoring data from a device in the backup storage pool, export the backup storage pool.

For more information on export of ZFS storage pools, refer to the operating system manual.

[Example]

```
# zpool list
NAME      SIZE  USED    AVAIL    CAP  HEALTH  ALTROOT
STP1     345G  121M    345G     0%   ONLINE  -
# zpool export STP1
#
```

3. Using AdvancedCopy Manager, restore data to a device in the transaction storage pool from a device in the backup storage pool or a backup volume.

When Restoring Data from Device in Backup Storage Pool Using Replication Management Function

Specify backup/restore group "STP1" for the `-Xreverse` option and the `-Xgroup` option and execute the `swsrpmake` command.

[Example]

```
# /opt/FJSVswsrp/bin/swsrpmake -Xgroup STP1 -Xreverse
GROUP=STP1 swsrpmake completed
#
```

When Restoring Data from Device in Backup Storage Pool Using AdvancedCopy Manager CCM

Specify backup/restore group "STP1" for the `-r` option and the `-g` option and execute the `acopc start` command for AdvancedCopy Manager CCM.

[Example]

```
# /opt/FJSVccm/bin/acopc start -g STP1 -r
DX440/0x11:DX440/0x01
# DATE : 2012/01/01 00:00:00 - << OPC Started >>
# From:
BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=11/Adr_high=0/Adr_low=0
# To:
BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=01/Adr_high=0/Adr_low=0/size_high=0/size_low=0
Succeeded : 1
Failed : 0
#
```

When Restoring Data from Device in Backup Volume Using Replication Management Function

Specify restore group "RSSTP1" for the -Xgroup option and execute the [swsrpmake](#) command.

[Example]

```
# /opt/FJSVswsrp/bin/swsrpmake -Xgroup RSSTP1
GROUP=RSSTP1 swsrpmake completed
#
```

When Restoring Data from Device in Backup Volume Using AdvancedCopy Manager CCM

Specify restore group "RSSTP1" for the -g option and execute the acopc start command for AdvancedCopy Manager CCM.

[Example]

```
# /opt/FJSVccm/bin/acopc start -g RSSTP1
DX440/0x12:DX440/0x01
# DATE : 2012/01/01 00:00:00 - << OPC Started >>
# From:
BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=12/Adr_high=0/Adr_low=0/size_high=0/size_low=0
# To:
BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=01/Adr_high=0/Adr_low=0
Succeeded : 1
Failed : 0
#
```



Note

Since processing is carried out for all devices that constitute the storage pool, do not specify any copy pairs.

4. In the transaction storage pool, import the device restored at step 3 in the restore procedure.

For information on how to import a ZFS storage pool, refer to the operating system manuals.

[Example]

```
# zpool import STP1
#
```

5. Roll back the transaction storage pool to a restored snapshot.

To see how to roll back the ZFS storage pool to a snapshot, refer to the operating system manual.

[Example]

For rolling back to the "SNAP1"

```
# zfs list -t snapshot
NAME          USED  AVAIL  REFER  MOUNTPOINT
STP1@SNAP1    0      -    19K    -
STP1@SNAP2   16M     -    18K    -
STP1@SNAP3   32M     -    18K    -
# zfs rollback -r STP1@SNAP1
#
```

6. Mount the target ZFS file systems.

If the ZFS file systems have not been mounted automatically by step 5, mount them manually.

For information on how to mount ZFS file systems, refer to the operating system manuals.

[Example]

```
# zfs mount STP1/data1
# zfs mount STP1/data2
#
```

G.3 Online Backup Using ZFS Snapshot with ZFS Uses `zfs send` and `zfs receive` Commands

This section explains methods for implementing online backup.

G.3.1 Operation Design

When you design your storage pools and backup operation, pay attention to the following aspects:

- "[Unit of Backup and Restoration: Storage Pool](#)"
- "[Version of ZFS Storage Pools](#)"
- "[Types and Configurations of ZFS Storage Pools](#)"
- "[Preparing Backup Volume](#)"

Unit of Backup and Restoration: Storage Pool

AdvancedCopy Manager supports backup and restoration only in units of ZFS storage pools (all file systems within a storage pool). Backup and restoration in units of ZFS file systems are not available.

For details, refer to "[G.1.1 Units of Backing up](#)".

Version of ZFS Storage Pools

For implementing online backup, ZFS storage pool version of business servers and that of backup servers must be the same. For checking your ZFS storage pool version, refer to the operating system manuals.

[Example]

```
# zpool upgrade -v
This system is currently running ZFS pool version 33.
```

Types and Configurations of ZFS Storage Pools

Purposes of ZFS Storage Pools: System Area or Data Area

With online backup, you can back up and restore data areas (usual storage pools). For backing up system areas (ZFS root pools), you have to use offline backup.

Devices Constituting ZFS Storage Pool

Online backup supports the following as device that constitutes ZFS storage pools (Physical disks).

Device type

Normal devices are supported.

Point

For details on using replication management function, refer to "1.4 Managing Device on AdvancedCopy Manager".

Device unit

Entire disks (LUs) are supported.

Note

As is constituted with the entire disks (LUs), create storage pools.

Preparing Backup Volume

As backup targets in online backup, you need to arrange for the following 2 types of backup targets:

- Backup storage pool for latest backup generations
- Backup volume for old backup generations

Backup Storage Pool for Latest Backup Generations

A "backup storage pool" serves to store full and differential backups of the latest data generations.

Make sure to arrange for one on the backup server.

The backup storage pool must have the same physical disk configuration (number of disks, sizes, redundant configuration) as the ZFS storage pool that is the backup source. Also, unit of devices that constitute storage pools must be the same.

Backup Volume for Storing Old Backup Generations

A "backup volume for older generations" is a backup volume for storing the storage pools that contain full and differential backups of the older data generations.

Make sure to arrange for the number of backup generations for the storage pools (full and differential backups) to be preserved on the backup server.

A backup volume for older generations must fulfill the conditions and have a configuration as described below.

For any conditions that are not stated below, depending on the type of device in the backup source volume, refer to "1.4 Managing Device on AdvancedCopy Manager", "Chapter 6 Replication", or the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module*.

- Make sure that the backup volume has the same physical disk configuration (number of disks, sizes, and redundant configuration) as the ZFS storage pool that is the backup source.
 - If a ZFS storage pool consists of multiple physical disks, you have to use a backup volume whose physical disks all have the same size.
- If you are using AdvancedCopy Manager CCM, do not connect the backup volume for older generations to the server.
- When using replication management function, use backup volume for older generations as slices.

Combination of Transaction Storage Pools and Backup Volumes

The combinations of transaction storage pools, backup storage pools for the latest generation and backup volumes for older generations are described below.

Function	Device		
	Transaction Storage Pools	Backup Volume	
		Backup Storage Pools for Latest Generation	Backup Volume for Older Generations
Replication management	LU	LU	Slice (*1)
AdvancedCopy Manager CCM	LU	LU	LU

*1: Define copy source as LU and copy destination as slice, then copy the entire LUs to the slice.

If a backup destination volume is LU, make the disk format of the transaction storage pool the same as that of the backup destination volume as shown below:

- If the transaction storage pool is a VTOC-labeled disk, the backup destination volume should also be a VTOC-labeled disk.
- If the transaction storage pool is an EFI-labeled disk, the backup destination volume should also be an EFI-labeled disk.

Connection Servers of Backup Target Volumes

The servers to which to connect the backup volumes, depending on the copying functions to be used, are as shown below.

Backup Target Volume	Copying Function		Remarks
	Replication Management	AdvancedCopy Manager CCM	
Backup storage pool	Backup server	Backup server	Create the backup storage pool on the backup server.
Backup volume for older generations	Backup server	Do not connect to a server.	

G.3.2 System Configuration


Since, in online backup operations, backup servers are used for backing up, backup servers are required that are separate from business servers. Consequently, you need a server configuration which consists of at least 2 servers.

As data communication to backup servers uses ssh, it is necessary to enable ssh connections from business servers to backup servers. For details on setting methods, refer to the operating system manuals.

Replication Operations

This section explains the system configuration for implementing replication operations.

Figure G.4 Sample Configuration for Replication Operations



Details on Sample Configuration

Using 1 physical disk that constitutes a storage pool for transactions, the entire disk is registered.

Including the backup storage pool, backup can preserve a maximum of 2 generations.

- Business server (Managed Server): SV-1
- Backup server (Management Server): BKSV (name for registration in AdvancedCopy Manager)
- Transaction storage pool: STP1
- Backup storage pool: STP1B
- Disk


Physical Disk	Connection Server	Device Type	Device Name	Purpose (Storage Pool)
Transaction VOL	SV-1	LU	c1t1d1	Transaction storage pool (STP1)
Backup VOL-1	BKSV	LU	c2t1d11	Backup storage pool (STP1B)

Physical Disk	Connection Server	Device Type	Device Name	Purpose (Storage Pool)
Backup VOL-2	BKSV	Slice	c2t1d12s1	Backup volume for older generations

AdvancedCopy Manager CCM Operations

This section explains the system configuration for implementing AdvancedCopy Manager CCM operations.

Figure G.5 Sample Configuration for AdvancedCopy Manager CCM Operations



Details on Sample Configuration

Using 1 physical disk that constitutes a storage pool for transactions, the entire disk is registered.

Including the backup storage pool, backup can preserve a maximum of 2 generations.

- Business server: SV-1
- CCM Server: BKSV
- Transaction storage pool: STP1
- Backup storage pool: STP1B

- Disk

Physical Disk	Connection Server	LUN	Device Type	Device Name	Purpose (Storage Pool)
Transaction VOL	SV-1	0x01	LU	c1t1d1	Transaction storage pool (STP1)
Backup VOL-1	BKSV	0x11	LU	c2t1d11	Backup storage pool (STP1B)
Backup VOL-2	-	0x12	LU	-	Backup volume for older generations

G.3.3 Preparations

This section describes the preparations that are required for backing up and restoring ZFS storage pools.

For details on each procedure and a setup required in advance for using AdvancedCopy Manager, refer to "[Chapter 6 Replication](#)" and the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module*.

Whereas this document states some typical methods for controlling ZFS for your reference, refer to the operating system manuals for more details on usage.



Point

Below examples of command execution are based on the configurations described for replication operations and AdvancedCopy Manager CCM operations in "[G.3.2 System Configuration](#)". Adapt any real control actions to your actual configuration.

Implement the following procedures for preparation:

1. [Fetching Lists of Devices in Transaction Storage Pools](#)
2. [Creation of Backup Storage Pools](#)
3. [Copy Settings for Backup and Restoration](#)

Fetching Lists of Devices in Transaction Storage Pools

On the business servers, check the devices that constitute the transaction storage pools to be backed up.

[Example]

Check the disk information of the storage pool. In below example, storage pool (STP1) consists of 1 physical disk (c1t1d1).

```
# zpool status STP1
pool: STP1
state: ONLINE
scrub: none requested
config:
  NAME          STATE      READ WRITE CKSUM
  STP1          ONLINE    0    0    0
    c1t1d1      ONLINE    0    0    0

errors: No known data errors
#
```

Creation of Backup Storage Pools

On the backup servers, create the backup storage pools to which the snapshots of the transaction storage pools are backed up.

[Example]

Create a backup storage pool with the same configuration as transaction storage pool (STP1). In below example, the backup storage pool (STP1B) is created on 1 physical disk (c2t1d11).

```
# zpool create STP1B c2t1d11
#
```

Copy Settings for Backup and Restoration

Make the following 3 copy settings:

- Copy settings for the backup storage pools and the backup volumes for older generations (for backup)
- Copy settings for the backup storage pools and the transaction storage pools (for restoring the latest generations)
- Copy settings for the backup volumes for older generations and the transaction storage pools (for restoring older generations)

Use the following combinations to register them in AdvancedCopy Manager:

Table G.3 Copy Groups

Used for	Copy Source	Copy Destination	Group Name (ex.)
Backup	Backup storage pool	Backup volume for older generations	BKSTP1
Restoration of the latest generations	Backup storage pool	Transaction storage pool	R0STP1
Restoration of older generations	Backup volume for older generations	Transaction storage pool	R1STP1

Table G.4 Registered Devices

Device Name	Devices to Be Registered	Example (*1)	
		Device Name	LUN
Transaction storage pool	Devices that you confirmed in " Fetching Lists of Devices in Transaction Storage Pools " to be devices constituting the transaction storage pool	c1t1d1	0x01
Backup storage pool	Devices that you created in " Creation of Backup Storage Pools " and constitute the backup storage pool	c2t1d11	0x11
Backup volume for older generations	Devices you prepared as backup volumes for older generations	c2t1d12s1	0x12

*1: When using replication management function, use device name. When using AdvancedCopy Manager CCM, use LUN number.

For Using Replication Management Function

Register copy sources indicated at "Copy groups" as copy destination volumes, and copy destinations as copy destination volumes. Then register the devices indicated at "Registered devices" respectively.

Since these are controlled synchronously in processing for backup and restoration, you have to register all devices that constitute a storage pool as a single group. For details on controlling groups, refer to "[6.4.7 Creating Groups](#)".

[Example]

Specify the source volumes and copy destination volumes to create the 3 groups shown in the chart.

```
# /opt/FJSVswsrp/bin/swsrpsetvol -Xgroup BKSTP1 /dev/dsk/c2t1d11 /dev/dsk/c2t1d12s1
swsrpsetvol completed
# /opt/FJSVswsrp/bin/swsrpsetvol -o BOTH -Xgroup R0STP1 /dev/dsk/c2t1d11@BKSV /dev/dsk/
c1t1d1@SV-1
swsrpsetvol completed
# /opt/FJSVswsrp/bin/swsrpsetvol -o BOTH -Xgroup R1STP1 /dev/dsk/c2t1d12s1@BKSV /dev/dsk/
c1t1d1@SV-1
swsrpsetvol completed
# /opt/FJSVswsrp/bin/swsrpvoinfo -L
```

Server	Original-Volume	Size	Replica-Volume	Size	Copy	Op-
BKSV	/dev/dsk/c2t1d11@BKSV	4.0 Gbyte	/dev/dsk/c2t1d12s1@BKSV	4.0 Gbyte	bi-direction	
original	BKSTP1					
BKSV	/dev/dsk/c2t1d11@BKSV	4.0 Gbyte	/dev/dsk/clt1d1@SV-1	4.0 Gbyte	bi-direction	
both	R0STP1					
BKSV	/dev/dsk/c2t1d12s1@BKSV	4.0 Gbyte	/dev/dsk/clt1d1@SV-1	4.0 Gbyte	bi-direction	
both	R1STP1					
#						

For Using AdvancedCopy Manager CCM

Register each combination of copy sources and copy destinations.

Since these are controlled synchronously in processing for backup and restoration, you have to register all devices that constitute a storage pool as a single group. For details on controlling groups, refer to "Creating Copy Groups" and "Adding Copy Pairs" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module*.

1. Creation of copy groups

Create a copy group for each copy setting.

[Example]

Create the 3 copy groups shown in the chart. (Copying type: OPC)

```
# acgroup create -g BKSTP1 -type OPC -a DX440
Successful completion.
#
# acgroup create -g R0STP1 -type OPC -a DX440
Successful completion.
#
# acgroup create -g R1STP1 -type OPC -a DX440
Successful completion.
#
```

2. Registration of copy pairs

In the created copy groups, register the combinations of copy sources and copy destinations. Register all devices that constitute a storage pool.

[Example]

In the created copy groups, register each device shown in the chart.

```
# acpair add -g BKSTP1 -p DX440/0x11:DX440/0x12
Successful completion.
# acpair add -g R0STP1 -p DX440/0x11:DX440/0x01
Successful completion.
# acpair add -g R1STP1 -p DX440/0x12:DX440/0x01
Successful completion.
```

G.3.4 Operation

This section describes procedures for backing up and restoring ZFS storage pools.

- "Backup"
 - "Full Backup"
 - "Differential Backup"
- "Restoration"
 - "Restoration from Latest Generations"

- "Restoration from Older Generations"

For details on the settings (device settings) in advance for using AdvancedCopy Manager, or on the backup and restoration procedures, refer to "Chapter 6 Replication" or the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module*.

Note that the procedures and commands shown here are examples of snapshot-type copying (OPC). For using other types of copying, replace the procedures and commands as needed for each type of copying.

Whereas this document states some typical methods for controlling ZFS for your reference, refer to the operating system manuals for more details on usage.

 **Point**

Below examples of command execution are based on the operating methods for replication operations and AdvancedCopy Manager CCM operations described in "G.3.2 System Configuration". Adapt any real control actions to your actual configuration.

- Backup operations to run a full backup once a week and a differential backup every day, keeping a total of 7 snapshots for each generation per full backup.
- Assign snapshot0 as the name for the snapshot for backups, snapshot1 for the snapshot for differential backup on Day 1, and snapshot2 for that of Day 2.

Backup

Full Backup

This section shows the procedure for a full backup of a transaction storage pool.

No.	Description	Controlling Server	At First Full Backup (*1)
1	Save the full backup of the preceding generation.	Backup server	
2	Create the backup storage pool again.	Backup server	
3	Delete all of the snapshots of the transaction storage pool.	Business server	
4	Create a snapshot of the transaction storage pool.	Business server	Yes
5	Back up the snapshot of the transaction storage pool.	Business server	Yes

*1: When you make a full backup for the first time, only implement the processes indicated with "Yes".

The procedure for the full backup is as follows:

1. Save the full backup of the preceding generation.

Before capturing the full backup of the latest generation, save the data of the preceding full backup that is in the backup storage pool.

- a. Export the backup storage pool.

For details, refer to the operating system manuals.

[Example]

```
# zpool export STP1B
#
```

- b. Use AdvancedCopy Manager to back up the devices in the backup storage pool to a backup volume for older generations.

Back up all devices constituting the backup storage pool to a backup volume for older generations. If you are using multiple backup volumes for older generations (keeping full backups of 2 or more generations), back up to an unused backup volume or to the one with the oldest generation.

For Using Replication Management Function

Specify backup group "BKSTP1" as the group and execute the `swsrpmake` command.

[Example]

```
# /opt/FJSVswsrp/bin/swsrpmake -Xgroup BKSTP1
GROUP=BKSTP1 swsrpmake completed
#
```

For Using AdvancedCopy Manager CCM

Specify backup group "BKSTP1" as the group and execute the `acopc start` command for AdvancedCopy Manager CCM.

[Example]

```
# /opt/FJSVccm/bin/acopc start -g BKSTP1
DX440/0x11:DX440/0x12
# DATE : 2012/01/01 00:00:00 - << OPC Started >>
#
From:BoxID=303045343030304d3323232323234534353053323041232323234b443430333036
33393030342323/
Olu=11/
Adr_high=0/Adr_low=0/size_high=0/size_low=0
#
To :BoxID=303045343030304d3323232323234534353053323041232323234b443430333036
33393030342323/
Olu=12/
Adr_high=0/Adr_low=0
Succeeded : 1
Failed : 0
#
```



Note

Since processing is carried out for all devices that constitute the storage pool, do not specify any copy pairs.

2. Create the backup storage pool again.

In order to back up the snapshot of the latest generation, create the backup storage pool on the backup server again. For information on how to create a ZFS storage pool, refer to the operating system manuals.

[Example]

```
# zpool create -f STP1B c2t1d11
#
```

3. Delete all snapshots of the transaction storage pool.

Delete all snapshots that are in the transaction storage pool. For information on how to delete snapshots, refer to the operating system manuals.

[Example]

```
# zfs destroy STP1@SNAP0
# zfs destroy STP1@SNAP1
# zfs destroy STP1@SNAP2
# zfs destroy STP1@SNAP3
# zfs destroy STP1@SNAP4
# zfs destroy STP1@SNAP5
# zfs destroy STP1@SNAP6
```

4. Create a snapshot of the transaction storage pool.

Create a snapshot of the transaction storage pool for full backup. For information on how to create snapshots, refer to the operating system manuals.

[Example]

```
# zfs snapshot -r STP1@SNAP0
#
```

5. Back up the snapshot of the transaction storage pool.

Execute the zfs send and zfs receive commands of ZFS to back up the snapshot of the transaction storage pool to the backup storage pool (Full backup). For information on how to use the zfs send and zfs receive commands, refer to the operating system manuals.

[Example]

```
# zfs send -R STP1@SNAP0 | ssh [backup-server] zfs receive -F STP1B
#
```

Differential Backup

This section shows the procedure for a differential backup of a transaction storage pool.

No.	Description	Controlling Server
1	Create a snapshot of the transaction storage pool.	Business server
2	Back up the snapshot of the transaction storage pool.	Business server

The procedure for the differential backup is as follows:

1. Create a snapshot of the transaction storage pool.

Create a snapshot of the transaction storage pool for differential backup. For information on how to create snapshots, refer to the operating system manuals.

[Example]

```
# zfs snapshot -r STP1@SNAP4
#
```

2. Back up the snapshot of the transaction storage pool.

Execute the zfs send and zfs receive commands of ZFS to back up the snapshot of the transaction storage pool to the backup storage pool (Differential backup). For information on how to use the zfs send and zfs receive commands, refer to the operating system manuals.


[Example]

Back up the differential data for SNAP4 (differential data between SNAP3 and SNAP4).

```
# zfs send -i STP1@SNAP3 STP1@SNAP4 | ssh [backup-server] zfs receive -F STP1B
#
```

Restoration

Restoration from Latest Generations



This section shows the procedure for restoring a latest-generation full backup and its differential backup generations.

No.	Description	Controlling Server
1	Unmount all ZFS file systems that are in the transaction storage pool.	Business server
2	Delete the transaction storage pool.	Business server
3	Export the backup storage pool.	Backup server
4	Use AdvancedCopy Manager to restore from the devices in the backup storage pool to the devices in the transaction storage pool.	Business server or backup server
5	Change the pool name and import the restored devices.	Business server
6	Roll back to the snapshots to be restored.	Business server
7	Mount the target ZFS file systems.	Business server

The procedure for restoring from the latest generations is as follows:

1. Unmount all ZFS file systems that are in the transaction storage pool.

Unmount all ZFS file systems that are in the transaction storage pool to be restored. For information on how to unmount ZFS file systems, refer to the operating system manuals.

[Example]

```
# zfs mount
STP1/data1 /mnt/data1
STP1/data2 /mnt/data2
#
```

```
# zfs umount /mnt/data1
# zfs umount /mnt/data2
#
```

2. Delete the transaction storage pool.

Delete (destroy) the transaction storage pool that is to be restored. For information on how to delete a ZFS storage pool, refer to the operating system manuals.

[Example]

```
# zpool destroy STP1
#
```

3. Export the backup storage pool.

Export the backup storage pool. For details, refer to the operating system manuals.

[Example]

```
# zpool export STP1B
#
```

4. Use AdvancedCopy Manager to restore from the devices in the backup storage pool to the devices in the transaction storage pool.

Use AdvancedCopy Manager commands to restore from the devices in the backup storage pool to the devices in the transaction storage pool.

For Using Replication Management Function

Specify group (R0STP1) as the group to be restored from the latest generation and execute the [swsrpmake](#) command.

[Example]

```
# /opt/FJSSVswsrp/bin/swsrpmake -Xgroup R0STP1
GROUP=R0STP1 swsrpmake completed
#
```

For Using AdvancedCopy Manager CCM

Specify the group to be restored from the latest generation and execute the `acopc start` command for AdvancedCopy Manager CCM.

[Example]

```
# /opt/FJSSVccm/bin/acopc start -g R0STP1
DX440/0x11:DX440/0x01
# DATE : 2012/01/01 00:00:00 - << OPC Started >>
#
From:BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=11/
Adr_high=0/Adr_low=0/size_high=0/size_low=0
#
To :BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=1/
Adr_high=0/Adr_low=0

Succeeded : 1
Failed : 0
#
```

Note

Since processing is carried out for all devices that constitute the storage pool, do not specify any copy pairs.

5. Change the pool name and import the restored devices.

Import the devices that you restored in step 4. Since the pool name is currently that of the backup storage pool, modify the pool name to that of the original transaction storage pool before you start the import. For information on how to import a ZFS storage pool, refer to the operating system manuals.

[Example]

```
# zpool import STP1B STP1
#
```

6. Roll back to the snapshots to be restored.

Roll back to the snapshots to be restored. For information on how to roll back to the snapshots, refer to the operating system manuals.

[Example]

For rolling back to the differential backup of Day 4 (SNAP4)

```
# zfs rollback STP1@SNAP4
#
```


7. Mount the target ZFS file systems.

Mount the ZFS file systems that are in the restored transaction storage pool. If the ZFS file systems have not been mounted automatically by step 6, mount them manually. For information on how to mount ZFS file systems, refer to the operating system manuals.

[Example]

```
# zfs mount STP1/data1
# zfs mount STP1/data2
#
```

Restoration from Older Generations



This section shows the procedure for restoring from an older-generation full backup and its differential backups.

No.	Description	Controlling Server
1	Unmount all ZFS file systems that are in the transaction storage pool.	Business server
2	Delete the transaction storage pool.	Business server
3	Use AdvancedCopy Manager to restore from the backup volume for older generations to the devices in the transaction storage pool.	Business server or backup server
4	Change the pool name and import the restored devices.	Business server
5	Roll back to the snapshots to be restored.	Business server
6	Mount the target ZFS file systems.	Business server

The procedure for restoring from older generations is as follows:

1. Unmount all ZFS file systems that are in the transaction storage pool.

Unmount all ZFS file systems that are in the transaction storage pool to be restored. For information on how to unmount ZFS file systems, refer to the operating system manuals.

[Example]

```
# zfs mount
STP1/data1 /mnt/data1
STP1/data2 /mnt/data2
#
# zfs umount /mnt/data1
# zfs umount /mnt/data2
#
```

2. Delete the transaction storage pool.

Delete (destroy) the transaction storage pool that is to be restored. For information on how to delete a ZFS storage pool, refer to the operating system manuals.

[Example]

```
# zpool destroy STP1
#
```

3. Use AdvancedCopy Manager to restore from the backup volume for older generations to the devices in the transaction storage pool.

Use AdvancedCopy Manager commands to restore from the backup volume for older generations to the devices in the transaction storage pool.

For Using Replication Management Function

Specify group (R1STP1) as the group to be restored from the older generation and execute the [swsrpmake](#) command for the group.

[Example]

```
# /opt/FJSVswsrp/bin/swsrpmake -Xgroup R1STP1
GROUP=R1STP1 swsrpmake completed
#
```

For Using AdvancedCopy Manager CCM

Specify group (R1STP1) as the group to be restored from the older generation and execute the `acopc start` command for AdvancedCopy Manager CCM.

[Example]

```
# /opt/FJSVccm/bin/acopc start -g R1STP1
DX440/0x12:DX440/0x01
# DATE : 2012/01/01 00:00:00 - << OPC Started >>

From:BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=12/
Adr_high=0/Adr_low=0/size_high=0/size_low=0
#
To :BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=1/
Adr_high=0/Adr_low=0

Succeeded : 1
Failed : 0
#
```



Note

Since processing is carried out for all devices that constitute the storage pool, do not specify any copy pairs.

4. Change the pool name and import the restored devices.

Import the devices that you restored in step 3. Since the pool name is currently that of the backup storage pool, modify the pool name to that of the original transaction storage pool before you start the import.

[Example]

```
# zpool import STP1B STP1
#
```

5. Roll back to the snapshots to be restored.

Roll back to the snapshots to be restored. For information on how to roll back to the snapshots, refer to the operating system manuals.

[Example]

For rolling back to the differential backup of Day 4 (SNAP4) of the older generation (1 week earlier)

```
# zfs rollback STP1@SNAP4
#
```

6. Mount the target ZFS file systems.

Mount the ZFS file systems that are in the restored transaction storage pool. If the ZFS file systems have not been mounted automatically by step 5, mount them manually. For information on how to mount ZFS file systems, refer to the operating system manuals.

[Example]

```
# zfs mount STP1/data1
# zfs mount STP1/data2
#
```

G.4 Offline Backup

This section explains methods for implementing offline backup.

G.4.1 Operation Design

When you design your storage pools and backup operation, pay attention to the following aspects:

- "[Unit of Backup and Restoration: Storage Pool](#)"
- "[Types and Configurations of ZFS Storage Pools](#)"
- "[Preparing Backup Volume](#)"
- "[Notes on Design](#)"

Unit of Backup and Restoration: Storage Pool

AdvancedCopy Manager supports backup and restoration only in units of ZFS storage pools. Backup in units of ZFS file systems is not available. For details, refer to "[G.1.1 Units of Backing up](#)".

Types and Configurations of ZFS Storage Pools

Purposes of ZFS Storage Pools: System Area or Data Area

With AdvancedCopy Manager, you can back up and restore system areas (ZFS root pools) as well as data areas (usual storage pools).

For backing up and restoring system areas (ZFS root pools), you have to stop the system. For details, refer to "[Backup and Restoration of System Areas](#)".

Before backing up and restoring data areas, you have to disconnect them from the system. For details, refer to "[Backup and Restoration of Data Areas](#)".

Devices Constituting ZFS Storage Pool

AdvancedCopy Manager supports the following devices (physical disks) in a ZFS storage pool configuration.

Device Types

The same devices are supported as for replication management function and AdvancedCopy Manager CCM, respectively. For details on supported devices when using replication management function, refer to "[1.4 Managing Device on AdvancedCopy Manager](#)".

If you use SDX object of PRIMECLUSTER GD/SafeDISK, only operation in units of slices is available in offline backup operation where server-to-server replication is performed.

With AdvancedCopy Manager CCM, copying is implemented by specifying disks in the ETERNUS Disk storage system, while there is no processing of the devices from the operating system to which they are connected. Therefore, if you are using, for example, the Logical Volume Manager of PRIMECLUSTER GD, consistency with the Logical Volume Manager can possibly not be guaranteed.

If you are using Logical Volume Manager volumes, use replication management function.

Point

For information on how to create a ZFS storage pool, refer to the manuals of each device.

Device Units

As devices (physical disks) in a ZFS storage pool configuration, entire disks (LUs) and slices are supported.

However, the units by which AdvancedCopy Manager can back up and restore depend on the AdvancedCopy Manager function you are using (replication management function or AdvancedCopy Manager CCM) and the device types (usual device or SDX object of PRIMECLUSTER GD) constituting the ZFS storage pool. For details, refer to the articles on functions and device types in the *ETERNUS SF AdvancedCopy Manager Operation Guide* and the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module*.

Note

- A ZFS storage pool configuration which contains both of entire disks (LUs) and slices is not supported.
- For backup in the system areas, you have to specify the entire disks (LUs) as backup sources.

Note

For Using AdvancedCopy Manager CCM

The only available unit of backup in AdvancedCopy Manager CCM is "entire disk" (LU). If you want to use AdvancedCopy Manager CCM to back up storage pools that consist of slices, you have to specify the LUs that contain these slices as backup sources.

In this case, however, since the entire LU is copied, all other slices in the LU are copied at the same time. Therefore, in your storage pool configuration and backup operation, pay attention to the following points:

- Define the entire LU as a single slice. Alternatively, embed all slices from the same LU into the same storage pool.

- If multiple storage pools consist of different slices that are in the same LU, back up or restore all of these storage pools simultaneously.

Preparing Backup Volume

A backup volume must fulfill the conditions and have a configuration as described below.

For any conditions that are not stated below, depending on the type of device in the backup source volume, refer to "[1.4 Managing Device on AdvancedCopy Manager](#)", "[Chapter 6 Replication](#)", or the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module*.

- Combinations of backup sources and targets

Backup sources and targets can be combined as shown below. Note, however, that available combinations depend on the device types, so, depending on the device types you are using, some combinations are not available.

Backup Source	Backup Target	
	Slice	LUN
Slice	Yes	No
LUN	Yes	Yes

- If both backup source and backup destination volumes are LU, make both disk types the same as shown below:
 - If the backup source volume is a VTOC-labeled disk, the backup destination volume should also be a VTOC-labeled disk.
 - If the backup source volume is an EFI-labeled disk, the backup destination volume should also be an EFI-labeled disk.
- Make sure that the backup volume has the same physical disk configuration (number of disks, sizes, and redundant configuration) as the ZFS storage pool that is the backup source.

If a ZFS storage pool consists of multiple physical disks, you have to use a backup volume whose physical disks all have the same size.

Likewise, if the backup source is a slice, arrange for a backup target that has the same number and sizes of slices.

Moreover, if the backup source is a LU and the backup target a slice, arrange for the same number and sizes for each.



Point

For Backing up Multiple Generations

Refer to "[Handling of Backup Volumes \(Destination Devices\)](#)".

Notes on Design

Handling of Backup Volumes (Destination Devices)

For operation, pay attention to the points below, so that devices (*) in the same storage pool on the same server do not appear as multiple ones.

- Do not connect backup target volumes (copy destination devices) to backup source servers.
- Do not connect multiple backup target volumes (copy destination devices) to the same server.

(*) Backup source volumes (copy source devices) and backup target volumes (copy destination devices) in the same storage pool

Figure G.6 Example for Unsupported Connection of Backup Target Volume: Backup Target Volume Connected to Backup Source Server





Figure G.7 Example for Unsupported Connection of Backup Target Volumes: Multiple Backup Target Volumes Connected to Same Server




Depending on the function you are using, respond as follows:

For Using Replication Management Function


- Connect copy destination volumes to servers (e.g. backup servers) other than the copy source volumes. (Perform replications from server to server.)

- For backing up multiple generations to the same backup server, the device combination should be "LU" for the copy source and "slice" for the copy destination.



For Using AdvancedCopy Manager CCM

- Connect copy destination volumes to servers other than the copy source volumes, or, alternatively, operate without any server connection.
- For copying multiple generations, do not connect the copy destinations to servers.



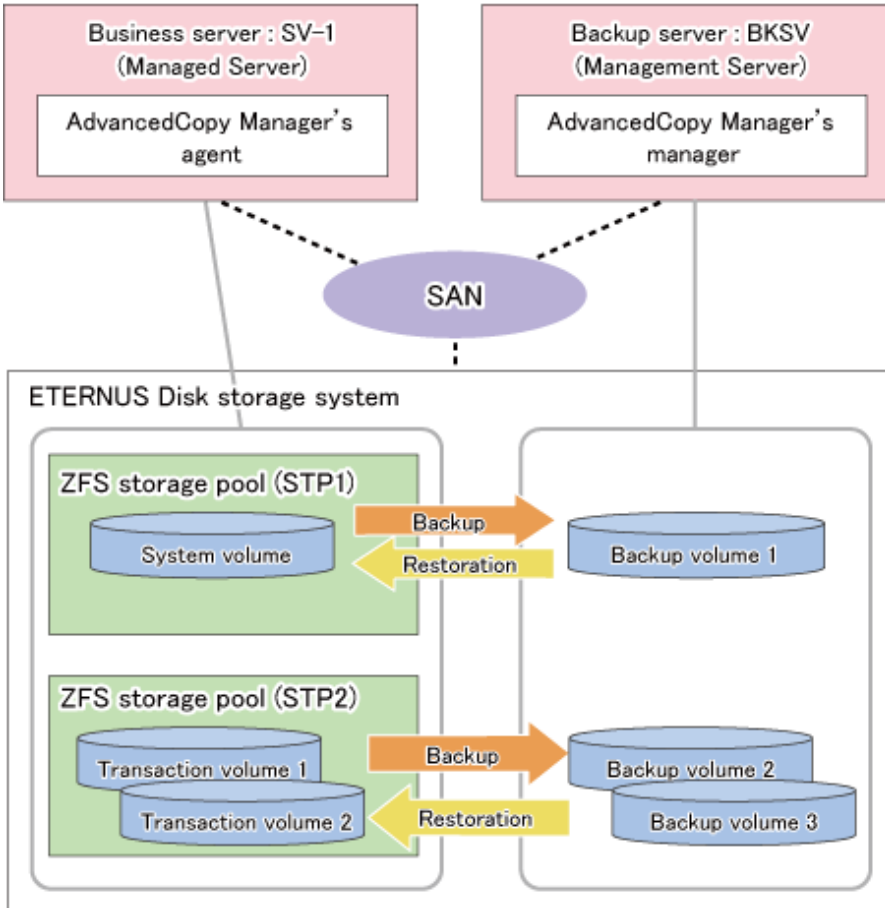
G.4.2 System Configuration

Replication Operations

This section explains the system configuration for implementing replication operations.

Since replications are implemented as server-to-server operations, a backup control server is required that is separate from the servers to be backed up. Consequently, you need a server configuration which consists of at least 2 servers.

Figure G.8 Sample Configuration for Replication Operations



Details on Sample Configuration

- Business server (Managed Server): SV-1 (name for registration in AdvancedCopy Manager)
- Backup server (Management Server): BKS (name for registration in AdvancedCopy Manager)
- System storage pool: STP1
- Transaction storage pool: STP2
- Disk


Physical Disk	Connection Server	Device Name	Purpose (Storage Pool)
System VOL	SV-1	c1t1d1s0	STP1
Transaction VOL-1	SV-1	c1t1d2	STP2
Transaction VOL-2	SV-1	c1t1d3	STP2
Backup VOL-1	BKS	c1t1d11s2	
Backup VOL-2	BKS	c1t1d12	
Backup VOL-3	BKS	c1t1d13	

AdvancedCopy Manager CCM Operations

This section explains the system configuration for implementing AdvancedCopy Manager CCM operations.

A backup control server is required that is separate from the servers to be backed up. Consequently, you need a server configuration which consists of at least 2 servers.

Figure G.9 Sample Configuration for AdvancedCopy Manager CCM Operations



Details on Sample Configuration

- Business server: SV-1
- CCM Server: BKSV
- System storage pool: STP1
- Transaction storage pool: STP2
- Disk

Physical Disk	LUN Number	Connection Server	Device Name	Purpose (Storage Pool)
System VOL	0x01	SV-1	c1t1d1s0	STP1
Transaction VOL-1	0x02	SV-1	c1t1d2	STP2
Transaction VOL-2	0x03	SV-1	c1t1d3	STP2
Backup VOL-1	0x11	-	-	-
Backup VOL-2	0x12	-	-	-
Backup VOL-3	0x13	-	-	-

G.4.3 Preparations

This section describes the preparations that are required for backing up and restoring ZFS storage pools.

For details on each procedure and a setup required in advance for using AdvancedCopy Manager, refer to "[Chapter 6 Replication](#)" and the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module*.

Whereas this document states some typical methods for controlling ZFS for your reference, refer to the operating system manuals for more details on usage.

Point

Below examples of command execution are based on the configurations described for replication operations and AdvancedCopy Manager CCM operations in "[G.4.2 System Configuration](#)". Adapt any real control actions to your actual configuration.

Implement the following procedures for preparation:

1. [Fetching Lists of Devices in Storage Pools](#)
2. [Copy Settings of AdvancedCopy Manager](#)

Fetching Lists of Devices in Storage Pools

Check the devices that constitute the storage pools to be backed up.

[Example]

Check the disk information of the storage pool. In below example, storage pool (STP1) consists of 1 physical disk (c1t1d1s0), and storage pool (STP2) consists of 2 physical disks (c1t1d2, c1t1d3).

```
# zpool status STP1
pool: STP1
state: ONLINE
scrub: none requested
config:
  NAME          STATE          READ WRITE CKSUM
  STP1          ONLINE         0    0    0
    c1t1d1s0    ONLINE         0    0    0

errors: No known data errors
#
# zpool status STP2
pool: STP2
state: ONLINE
scrub: none requested
config:
  NAME          STATE          READ WRITE CKSUM
  STP2          ONLINE         0    0    0
    c1t1d2      ONLINE         0    0    0
    c1t1d3      ONLINE         0    0    0

errors: No known data errors
```

Copy Settings of AdvancedCopy Manager

Register the combinations of backup sources and targets in AdvancedCopy Manager.

For Using Replication Management Function

Register the devices you checked in above procedure as copy source volumes, and the devices you prepared as backup target volumes as copy destination volumes.

Since these are controlled synchronously in processing for backup and restoration, you have to register all devices that constitute a storage pool as 1 group. For details on controlling groups, refer to "[6.4.7 Creating Groups](#)".

[Example]

Specify the copy source and destination volumes to create the groups for storage pools STP1 and STP2.

- With the disks constituting storage pool STP1 as the copy source volume, create group "STP1".

- With the disks constituting storage pool STP2 as the copy source volume, create group "STP2".

```
# /opt/FJSVswsrp/bin/swsrpsetvol -o BOTH -Xgroup STP1 -Xvtoc /dev/dsk/clt1d1s2@SV-1 /dev/dsk/
clt1d1s2@BKSV
swsrpsetvol completed
# /opt/FJSVswsrp/bin/swsrpsetvol -o BOTH -Xgroup STP2 /dev/dsk/clt1d2@SV-1 /dev/dsk/
clt1d12@BKSV
swsrpsetvol completed
# /opt/FJSVswsrp/bin/swsrpsetvol -o BOTH -Xgroup STP2 /dev/dsk/clt1d3@SV-1 /dev/dsk/
clt1d13@BKSV
swsrpsetvol completed
# /opt/FJSVswsrp/bin/swsrpvoinfo -L
Server Original-Volume      Size      Replica-Volume      Size      Copy      Op-Server
Group
SV-1 /dev/dsk/clt1d1s2@SV-1 4.0 Gbyte /dev/dsk/clt1d1s2@BKSV 4.0 Gbyte bi-direction
both STP1
SV-1 /dev/dsk/clt1d2@SV-1 4.0 Gbyte /dev/dsk/clt1d12@BKSV 4.0 Gbyte bi-direction
both STP2
SV-1 /dev/dsk/clt1d3@SV-1 4.0 Gbyte /dev/dsk/clt1d13@BKSV 4.0 Gbyte bi-direction
both STP2
```

 **Point**

For backup of system areas, you have to specify the entire disks (clt1d1s2) which contain system areas. Specify the -Xvtoc option for including the VTOC area.

For Using AdvancedCopy Manager CCM

Register the devices you checked in above procedure as copy source, and the devices you prepared as backup targets as copy destinations.

Since these are controlled synchronously in processing for backup and restoration, you have to register all devices that constitute a storage pool as 1 group. For details on controlling groups, refer to "Creating Copy Groups" and "Adding Copy Pairs" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module*.

1. Creation of copy groups

Create a copy group for each storage pool that is to be backed up.

[Example]

Create group "STP1" for storage pool STP1, and group "STP2" for STP2. (Copying type: OPC)

```
# acgroup create -g STP1 -type OPC -a DX440
Successful completion.
#
# acgroup create -g STP2 -type OPC -a DX440
Successful completion.
```

2. Registration of copy pairs

In the created copy groups, register the combinations of copy sources and destinations. Register all devices that constitute a storage pool.

[Example]

- With the disks constituting storage pool STP1 as the copy source, register the pairs of copy sources and targets in group "STP1".
- With the disks constituting storage pool STP2 as the copy source, register the pairs of copy sources and targets in group "STP2".


```
# acpair add -g STP1 -p DX440/0x01:DX440/0x11
Successful completion.
# acpair add -g STP2 -p DX440/0x02:DX440/0x12
Successful completion.
# acpair add -g STP2 -p DX440/0x03:DX440/0x13
Successful completion.
```

G.4.4 Operation

This section describes procedures for backing up and restoring ZFS storage pools.

- "Backup and Restoration of System Areas"
 - "Backup"
 - "Restoration"
- "Backup and Restoration of Data Areas"
 - "Backup"
 - "Restoration"

For details on the settings (device settings) to be made in advance for using AdvancedCopy Manager, or on the backup and restoration procedures, refer to "Chapter 6 Replication" and the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module*.

Note that the procedures and commands shown here are examples of snapshot-type copying (OPC). For using other types of copying, replace the procedures and commands as needed for each type of copying.

Whereas this document states some typical methods for controlling ZFS for your reference, refer to the operating system manuals for more details on usage.

Point

Below examples of command execution are based on the configurations described for replication operations and AdvancedCopy Manager CCM operations in "G.4.2 System Configuration". Adapt any real control actions to your actual configuration.

Backup and Restoration of System Areas

Backup

This section shows the procedure for backing up a system area (root pool).

No.	Description	Controlling Server
1	Stop the backup source server.	Business server
2	Use AdvancedCopy Manager to back up the devices in the storage pool that is to be backed up.	Backup server
3	Start the backup source server.	Business server

The procedure for backup is as follows:

1. Stop the backup source server.
 - For information on how to stop a server, refer to the operating system manuals.
 - After stopping the server, make sure that the server's power supply is turned OFF.
2. Use AdvancedCopy Manager to back up the devices in the storage pool that is to be backed up.
 - From the backup server, use AdvancedCopy Manager commands to perform the backup.

For Using Replication Management Function

Specify the group and execute the [swsrpmake](#) command. Specify the -m option in order to prevent data communication to the backup source server.

[Example]

```
# /opt/FJSVswsrp/bin/swsrpmake -m -Xgroup STP1
GROUP=STP1 swsrpmake completed
#
```

For Using AdvancedCopy Manager CCM

Execute the acopc start command for AdvancedCopy Manager CCM.

[Example]

```
# /opt/FJSVccm/bin/acopc start -g STP1
DX440/0x01:DX440/0x11
# DATE : 2012/01/01 00:00:00 - << OPC Started >>
#
From:BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=1/
Adr_high=0/Adr_low=0/size_high=0/size_low=0
#
To :BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=11/
Adr_high=0/Adr_low=0

Succeeded : 1
Failed : 0
#
```

 **Note**

Since processing is carried out for all devices that constitute the storage pool, do not specify any copy pairs.

3. Start the backup source server.

Restoration

This section shows the procedure for restoring a system area (root pool).

No.	Description	Controlling Server
1	Stop the server that is to be restored.	Business server
2	Use AdvancedCopy Manager to restore the devices in the storage pool that is to be restored.	Backup server
3	Start the server that is to be restored.	Business server

1. Stop the server that is to be restored.
For information on how to stop a server, refer to the operating system manuals. After stopping the server, make sure that the server's power supply is turned OFF.
2. Use AdvancedCopy Manager to restore the devices in the storage pool that is to be restored.
From the backup server, use AdvancedCopy Manager commands to perform the restore.

For Using Replication Management Function

Perform copying in the opposite direction (from copy destination to source).

Specify the -Xreverse option and execute the [swsrpmake](#) command. Also specify the -m option in order to prevent data communication to the server that is being restored.

[Example]

```
# /opt/FJSVswsrp/bin/swsrpmake -m -Xreverse -Xgroup STP1
GROUP=STP1 swsrpmake completed
#
```

For Using AdvancedCopy Manager CCM

Perform copying in the opposite direction (from copy destination to source).

Specify the -r option and execute the acopc start command for AdvancedCopy Manager CCM.

[Example]

```
# /opt/FJSVccm/bin/acopc start -r -g STP1
DX440/0x01:DX440/0x11
# DATE : 2012/01/01 00:00:00 - << OPC Started >>
#
From:BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=11/
Adr_high=0/Adr_low=0/size_high=0/size_low=0
#
To :BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=1/
Adr_high=0/Adr_low=0

Succeeded : 1
Failed : 0
#
```

Note

Since processing is carried out for all devices that constitute the storage pool, do not specify any copy pairs.

3. Start the server that is to be restored.

Backup and Restoration of Data Areas

Backup

This section shows the procedure for backing up a data area.

No.	Description	Controlling Server
1	Unmount all ZFS file systems that are in the storage pool to be backed up.	Business server
2	Export the storage pool to be backed up.	Business server
3	Use AdvancedCopy Manager to back up the devices in the storage pool that is to be backed up.	Business server or backup server
4	Import the storage pool to be backed up.	Business server
5	Mount the target ZFS file systems.	Business server

The procedure for backup is as follows:

1. Unmount all ZFS file systems that are in the storage pool to be backed up.

Unmount all ZFS file systems that are on the business server (backup source server). For information on how to unmount ZFS file systems, refer to the operating system manuals.

[Example]

```
# zfs mount
STP2/data1 /mnt/data1
STP2/data2 /mnt/data2
#
# zfs umount /mnt/data1
# zfs umount /mnt/data2
#
```

2. Export the storage pool to be backed up.

Export the ZFS storage pool on the business server (backup source server). For details, refer to the operating system manuals.

[Example]

```
# zpool export STP2
#
```

3. Use AdvancedCopy Manager to back up the devices in the storage pool that is to be backed up.

From the business server or backup server, use AdvancedCopy Manager commands to perform the backup.

For Using Replication Management Function

Specify the group and execute the [swsrpmake](#) command.

[Example]

```
# /opt/FJSVswsrp/bin/swsrpmake -Xgroup STP2
GROUP=STP2 swsrpmake completed
#
```

For Using AdvancedCopy Manager CCM

Execute the `acopc start` command for AdvancedCopy Manager CCM.

[Example]

```
# /opt/FJSVccm/bin/acopc start -g STP2
DX440/0x02:DX440/0x12
# DATE : 2012/01/01 00:00:00 - << OPC Started >>
#
From:BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=2/
Adr_high=0/Adr_low=0/size_high=0/size_low=0
#
To :BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=12/
Adr_high=0/Adr_low=0

DX440/0x03:DX440/0x13
# DATE : 2012/01/01 00:00:02 - << OPC Started >>
#
From:BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=3/
Adr_high=0/Adr_low=0/size_high=0/size_low=0
#
To :BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=13/
Adr_high=0/Adr_low=0

Succeeded : 2
Failed : 0
#
```

Note

Since processing is carried out for all devices that constitute the storage pool, do not specify any copy pairs.

4. Import the storage pool to be backed up.

Import the ZFS storage pool on the business server (backup source server). For details, refer to the operating system manuals.

[Example]

```
# zpool import STP2
#
```

5. Mount the target ZFS file systems.

Mount all ZFS file systems that are on the business server (backup source server). If the ZFS file systems have not been mounted automatically at step 4 (importing storage pool), mount them manually. For information on how to mount ZFS file systems, refer to the operating system manuals.

[Example]

```
# zfs mount STP2/data1
# zfs mount STP2/data2
#
```

Restoration

This section shows the procedure for restoring a data area.

No.	Description	Controlling Server
1	Unmount all ZFS file systems that are in the storage pool to be restored.	Business server
2	Delete the storage pool that is to be restored.	Business server
3	Use AdvancedCopy Manager to restore the devices in the storage pool that is to be restored.	Business server or backup server
4	Import the storage pool to be backed up.	Business server
5	Mount the target ZFS file systems.	Business server

The procedure for restoration is as follows:

1. Unmount all ZFS file systems that are in the storage pool to be restored.

Unmount all ZFS file systems that are on the business server (backup source server). For information on how to unmount ZFS file systems, refer to the operating system manuals.

[Example]

```
# zfs mount
STP2/data1 /mnt/data1
STP2/data2 /mnt/data2
#
# zfs umount /mnt/data1
# zfs umount /mnt/data2
#
```

2. Delete the storage pool that is to be restored.

Delete (destroy) the ZFS storage pool on the business server (backup source server). For information on how to delete a ZFS storage pool, refer to the operating system manuals.

[Example]

```
# zpool destroy STP2
#
```

3. Use AdvancedCopy Manager to restore the devices in the storage pool that is to be restored.

From the business server or the backup server, use AdvancedCopy Manager commands to perform the restore.

For Using Replication Management Function

Perform copying in the opposite direction (from copy destination to source).

Specify the `-Xreverse` option and execute the `swsrpmake` command.

[Example]

```
# /opt/FJVSwsrp/bin/swsrpmake -Xreverse -Xgroup STP2
GROUP=STP2 swsrpmake completed
#
```

For using AdvancedCopy Manager CCM

Perform copying in the opposite direction (from copy destination to source).

Specify the `-r` option and execute the `acopc start` command for AdvancedCopy Manager CCM.

[Example]

```
# /opt/FJSVccm/bin/acopc start -r -g STP2
DX440/0x02:DX440/0x12
# DATE : 2012/01/01 00:00:00 - << OPC Started >>
#
From:BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=12/
Adr_high=0/Adr_low=0/size_high=0/size_low=0
#
To :BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=2/
Adr_high=0/Adr_low=0

DX440/0x03:DX440/0x13
# DATE : 2012/01/01 00:00:02 - << OPC Started >>
#
From:BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=13/
Adr_high=0/Adr_low=0/size_high=0/size_low=0
#
To :BoxID=303045343030304d3323232323234534353053323041232323234b4434303330363339303
0342323/Olu=3/
Adr_high=0/Adr_low=0

Succeeded : 2
Failed : 0
#
```

4. Import the storage pool to be backed up.

Import the ZFS storage pool on the business server (backup source server). For details, refer to the operating system manuals.

[Example]

```
# zpool import STP2
#
```

5. Mount the target ZFS file systems.

Mount all ZFS file systems that are on the business server (backup source server). If the ZFS file systems have not been mounted automatically at step 4 (importing storage pool), mount them manually. For information on how to mount ZFS file systems, refer to the operating system manuals.

[Example]

```
# zfs mount STP2/data1
# zfs mount STP2/data2
#
```

G.5 Maintenance

This section explains maintenance methods to be applied during backup and restoration of ZFS file systems.

G.5.1 Changing Devices in Storage Pool Configuration

If you change the configuration of a storage pool, you have to reintegrate the configuration information in the AdvancedCopy Manager definitions and reimplement a backup. This section explains some typical procedures for changing configurations.

Changing a storage pool configuration means to make the following modifications to devices that constitute the storage pool:

- Adding/replacing/deleting a device
- Changing the capacity of a device (e.g. by recreating a LUN)

For changing a configuration, implement below procedure.

The steps 1, 2, 4, and 5 are implemented in AdvancedCopy Manager, step 3 in the operating system. For operational details, refer to the corresponding manuals.

1. For replacing a device, deleting a device, or changing volumes, cancel any EC/REC or QuickOPC copying sessions, or OPC physical copying for that device.
2. For replacing a device, deleting a device, or changing volumes, delete the devices and the pairing of the copy destination device to be replaced or deleted from the group definitions of AdvancedCopy Manager.
3. Change the device configuration in the storage pool.
4. Add devices and copy destination pairs, which have been added, replaced or changed their volumes, to the group definitions of AdvancedCopy Manager.
5. In order to maintain consistency of the entire storage pool, perform a backup of the reconfigured storage pool.

Appendix H Operation in Oracle VM Server for SPARC

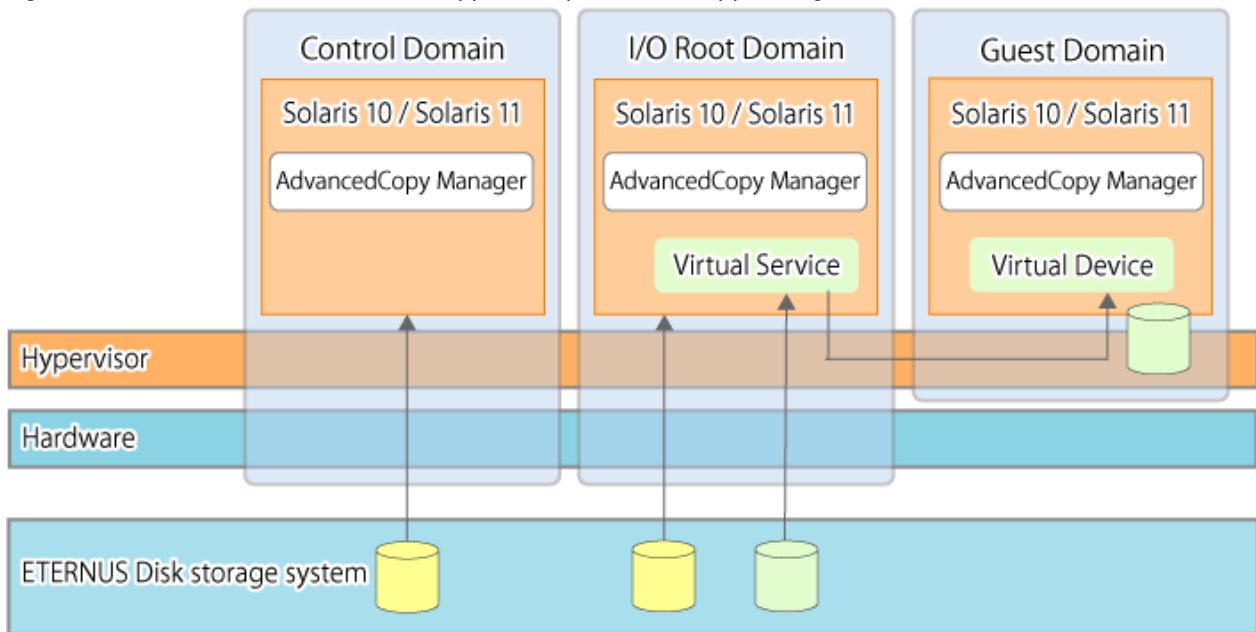
This appendix explains the operation in Oracle VM Server for SPARC.

H.1 Overview

In AdvancedCopy Manager, the virtual environment in Oracle VM Server for SPARC can be backed up and restored. The roles of the virtual environment in Oracle VM Server for SPARC to support backup and restoration are as follows:

- Control Domain
- I/O Root Domain (Root Domain)
- Guest Domain

Figure H.1 Oracle VM Server for SPARC Supported by AdvancedCopy Manager



H.2 Software Requirements

The basic software requirements for using Oracle VM Server for SPARC are as shown below.

Roles of Virtual Environment in Oracle VM Server for SPARC	Server Type	Basic Software
Control Domain I/O Root Domain Guest Domain	Management Server	Solaris 10 Solaris 11
	Managed Server	Solaris 10 Solaris 11
	CCM Server	Solaris 10 (*1) Solaris 11 (*1)

*1: The operation in Oracle Solaris Zones is not supported.

H.3 Available Functions

This section explains the functions that are available when using Oracle VM Server for SPARC.

H.3.1 Functions of AdvancedCopy Manager

The following functions of AdvancedCopy Manager are available:

- Backup
- Replication
- AdvancedCopy Manager CCM



-
- For the functions that are available through the Oracle Solaris Zones, refer to "[F.4.1 Functions of AdvancedCopy Manager](#)".
 - For details on the functions of AdvancedCopy Manager, refer to "AdvancedCopy Manager Functions" in the *ETERNUS SF AdvancedCopy Manager Overview*.
-

H.3.2 Cluster Operation

Cluster operation using PRIMECLUSTER is supported by Oracle VM Server for SPARC.

Cluster operation other than with PRIMECLUSTER is currently not supported.

H.4 Supported Devices

This section explains devices that can use when using Oracle VM Server for SPARC.

H.4.1 Supported Devices for Control Domain and I/O Root Domain

The following device names can be specified in AdvancedCopy Manager for Control Domain and I/O Root Domain:

Device Type	Specifiable Device Name
Normal device	<ul style="list-style-type: none">- sd driver resources- MPxIO driver resources- MPLB resources <p>The unit of management is a disk partition (slice) unit or a Logical Unit (disk unit). For details on device names, refer to "1.4.1 Normal Devices".</p>
SDX objects of PRIMECLUSTER GD	<ul style="list-style-type: none">- Names of devices that are operating in units of slices <p>The format used to specify a device name is as follows: <code>/dev/sfdsk/<class-name>/dsk/<volume-name>:<DEVNAM-value-of-sdxinfo></code> Example: <code>/dev/sfdsk/class1/dsk/volume1:c1t0d10</code></p> <p>For details on device names, refer to "1.4.2 SDX Objects of PRIMECLUSTER GD/ SynfinityDISK".</p> <p>Names of devices that are operating in units of logical volumes are not supported.</p>

H.4.2 Supported Device for Guest Domain

The following virtual device only can be specified in AdvancedCopy Manager for Guest Domain:

Point

Connect a virtual device specified for AdvancedCopy Manager to the Guest Domain using the Virtual Disk Services in the Control Domain or the I/O Root Domain. Allocate the entire LU to Virtual Disk Services.

Device Type	Specifiable Device Name
Virtual device	<p>Specify the device name that has been registered as a virtual disk using Virtual Disk Services (vds).</p> <p>The format used to specify a device name is as follows:</p> <pre>/dev/dsk/cXdXsX</pre> <p>Example: /dev/dsk/c0d2s0</p> <p>The format used to specify the LUN with EFI label is as follows:</p> <pre>/dev/dsk/cXdX</pre> <p>Example: /dev/dsk/c0d2</p>
Virtual device (MPLB resources)	<p>Specify the device name that has been set as an MPLB resource on the Guest Domain.</p> <p>The format used to specify a device name is as follows:</p> <pre>/dev/FJSVmplb/[r]dsk/mplb*s*</pre> <p>Example: /dev/FJSVmplb/dsk/mplb0s6</p> <p>Refer to the <i>ETERNUS Multipath Driver (Oracle Solaris Version) Installation Information</i> for information on how to specify MPLB resources for the Guest Domain.</p>
SDX objects of PRIMECLUSTER GD	<p>Names of devices that are operating in units of slices can be specified.</p> <p>The format used to specify a device name is as follows:</p> <pre>/dev/sfdsk/<class-name>/dsk/<volume-name>:<DEVNAM-value-of-sdxinfo></pre> <p>Example: /dev/sfdsk/class1/dsk/volume1:mplb0</p> <p>Refer to "1.4.2 SDX Objects of PRIMECLUSTER GD/SynfinityDISK" for details on device names.</p> <p>Names of devices that are operating in units of logical volumes are not supported.</p>

H.5 Operation

This section explains the operation of AdvancedCopy Manager in Oracle VM Server for SPARC.

For details on the backup operation, refer to the following sections:

- For backing up normal transaction volumes: "[Chapter 3 Backup Operation](#)".

For details on the replication operation, refer to the following section:

- For replicating: "[Chapter 6 Replication](#)".

For operating AdvancedCopy Manager CCM, use the following procedure:

1. Configure the environment for AdvancedCopy Manager CCM.

For details, refer to "Setup" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module*.

Note

For Guest Domain, the virtual device must be specified as the Access volume.

For details, refer to "[H.4.2 Supported Device for Guest Domain](#)".

2. Perform the operations for AdvancedCopy Manager CCM.

For details, refer to "Operation" in the *ETERNUS SF AdvancedCopy Manager Operation Guide for Copy Control Module*.